Accounting for Adaptation in the Economics of Happiness

Miles Kimball* Ryan Nunn† Dan Silverman‡

July 7, 2015

Abstract

Reported happiness provides a potentially useful way to evaluate unpriced goods and events; but measures of subjective well-being (SWB) often revert to the mean after responding to events, and this hedonic adaptation creates challenges for interpretation. Previous work tends to estimate time-invariant effects of events on happiness. In the presence of hedonic adaptation, this restriction can lead to biases, especially when comparing events to which people adapt at different rates. Our paper provides a flexible, extensible econometric framework that accommodates adaptation and permits the comparison of happiness-relevant life events with dissimilar hedonic adaptation paths. We present a method that is robust to individual fixed effects, imprecisely-dated data, and permanent consequences. The method is used to analyze a variety of events in the Health and Retirement Study panel. Many of the variables studied have substantial consequences for subjective well-being - consequences that differ greatly in their time profiles.

*Department of Economics, University of Michigan, 611 Tappan St., Ann Arbor, MI 48109. We are grateful for financial support from NIH/NIA grants R01-AG040787 and P01 AG026571 to the University of Michigan.

†U.S. Department of the Treasury, 1500 Pennsylvania Ave. NW, Washington, DC 20002. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the U.S. Treasury Department or of any other person associated with the U.S. Treasury Department.

‡Corresponding author. Department of Economics, Arizona State University, PO Box 879801, Tempe, AZ 85287. Email: Daniel.Silverman.1@asu.edu.

1 Introduction

In recent years, researchers have increasingly turned to survey data on subjective well-being (happiness) to investigate a wide range of economic questions. Happiness data have been used to study preferences over inflation and unemployment (Di Tella et al., 2001), the consequences of excise taxes (Gruber and Mullainathan, 2005), the externalities of neighbors' higher earnings (Luttmer, 2005), the progress of women (Stevenson and Wolfers, 2009), and the effect of health on preferences for consumption (Finkelstein et al., 2008), to take just a few examples. While the use of happiness data is often controversial, measures of an individual's recent mood or life-satisfaction have important potential for social science. Taken simply as an outcome, an individual's emotional state or level of satisfaction may be an important aspect of their well-being. Interpreted as information about utility, happiness data provide an opportunity to infer an individual's preferences under circumstances, such as the presence of externalities, non-market goods, or cognitive biases, where choices and prices alone will typically be inadequate. The potential value of these measures motivates the development of tools for rigorous economic analysis of happiness data.

This paper provides analytic tools to account for "hedonic adaptation," a feature of happiness data that makes their economic interpretation more difficult. Hedonic adaptation refers to an aspect of the dynamic response of happiness to changes in life-circumstances; the magnitude of the response decreases as the change fades into the past. In a variety of studies, there is evidence that happiness responds in expected ways to the arrival of both good and bad events, but individuals return to their prior levels of mood with surprising thoroughness and speed.[1] In a canonical example, people are much less happy upon the arrival of a serious health problem (paraplegia, renal failure) but eventually appear to adapt and reveal measured happiness at or near normal levels (Riis et al., 2005). The idea that happiness quickly adapts to changes in income has been central to investigations of the Easterlin Paradox, but it is common for economic analyses of happiness data to leave hedonic adaptation unmodeled.

We develop and estimate, using panel data from the Health and Retirement Study (HRS), a parsimonious model of hedonic adaptation that allows life events to have both transitory and permanent effects on subjective well-being (SWB). The model is based on a theory of happiness,

[1] Frederick and Loewenstein (1999) offer a review of this evidence. Diener et al. (2006) review evidence on the limits of hedonic adaptation.

Kimball and Willis (2006), that interprets the impulse response of SWB to an event as indicating the importance of that event for lifetime utility. That theory suggests that the rate of hedonic adaptation should depend on the particular type of event, so we estimate an event-specific rate at which the transitory effect decays. While the formulation is flexible, it provides a single, interpretable statistic by which to compare the overall SWB effects of life events with very different paths of hedonic adaptation.

The estimates show that the model fits the HRS data well and that accounting for hedonic adaptation can have qualitative effects on the inference drawn from happiness data. We find, for example, that estimates of the relative rank of important life events change substantially depending on whether adaptation affects are modeled. For instance, a commonly-used regression specification implies that unemployment is worse for SWB than either a heart attack or cancer. By contrast, when adaptation is taken into account, both health events are revealed to be more serious than unemployment. Likewise, accounting for adaptation downgrades the significance of life insurance for happiness in the aftermath of widowhood.

The methods presented here also reveal important differences across events in the paths of SWB that follow them. As an example, we investigate parental death, which has negligible, statistically insignificant SWB effects when estimated using a common method that does not account for adaptation. This masks a very large, but short-lived, initial drop in happiness after a parent's death. More generally, existing work that characterizes events as having only a single, permanent effect on SWB (what we term a "pooled regression" approach) does not fully exploit the panel nature of available data. This is especially relevant in light of our finding that life events have widely differing hedonic consequences, both in magnitude and in composition (e.g., temporary vs. permanent effects).

The figures below preview how the method characterizes the dynamics of SWB, and how it allows inference about the transitory effects of life events. Proceeding clockwise from the top left, the figure shows the estimated dynamics of SWB in response to widowing when the deceased had life insurance, widowing when the deceased did not, death of a child, and death of a mother. In these figures, the method quantifies how events differ in the size and persistence of their effects on SWB. It also allows inference about the ability of financial resources to ease the SWB effects of widowing (seemingly little) and about how well-being recovers from the loss of a parent (quickly

and thoroughly) versus the loss of a child (slowly, if at all).

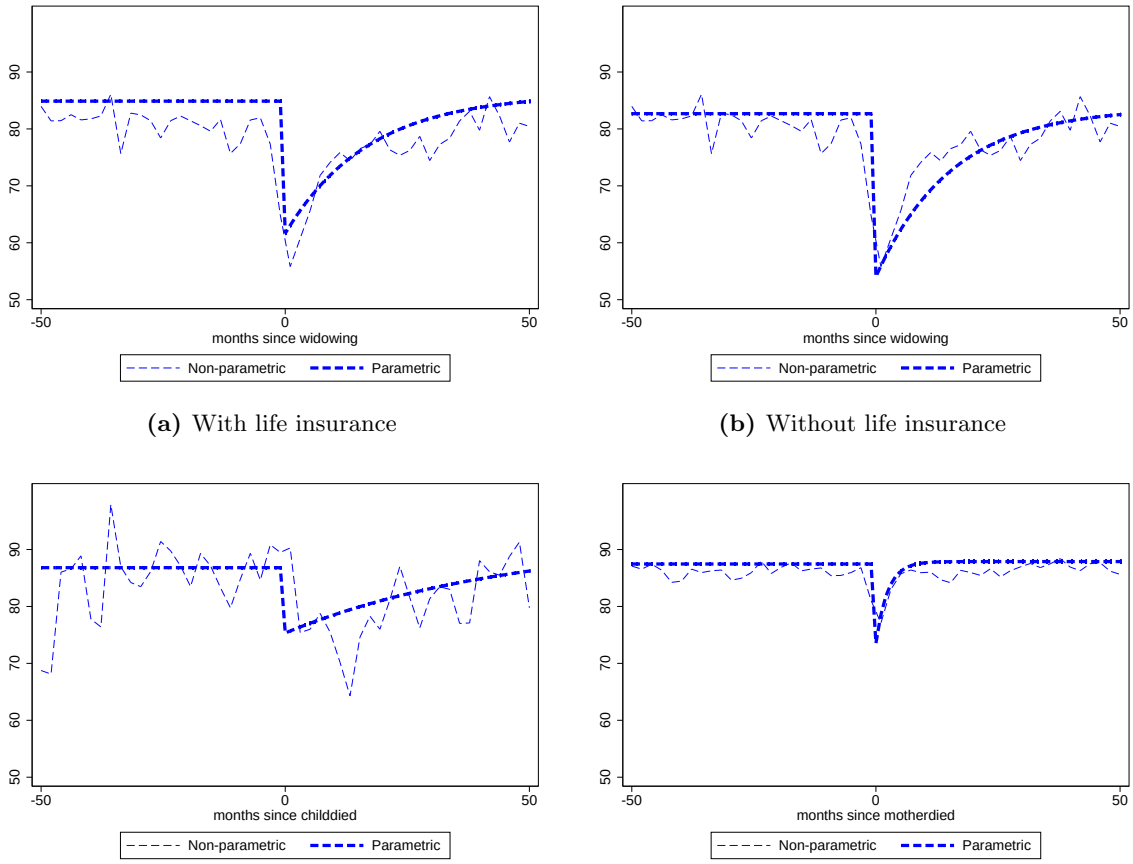

(a) With life insurance

(b) Without life insurance

(c) Child died

(d) Mother died

Figure 1: Baseline and non-parametric results for selected events

As in the rest of the SWB literature, the interpretation of these figures will depend on whether happiness data are viewed as providing information about utility or as representing outcomes of intrinsic interest analogous to health or income. In either case, we offer a convenient and tractable econometric method for making use of these data in the presence of hedonic adaptation.

The remainder of the paper proceeds as follows. First, in the next section, we discuss the related literature in greater detail, with special emphasis on the econometric approach taken in previous work. Section 3, presents the model and multiple variants of the baseline specification. That section also develops a procedure for comparing our results to previous work. In Section 4 and Section 5, we discuss data and results, with a discussion and conclusion following.

2 Related Literature

The SWB literature was advanced by early papers like Andrews and Withey (1976), Bradburn (1969), Campbell et al. (1976) that explored the potential usefulness of survey questions on happiness. As the literature matured, some of challenges of interpreting SWB data became evident. Schwarz and Strack (1999) and Schwarz (1987), for example describe some of the difficulties of interpreting reported happiness data, including the temporary and context-dependent nature of the reports. Subsequent research has also argued that SWB is not always best conceived of as a straightforward proxy for utility; Benjamin et al. (2012) show evidence for systematic discrepancies between SWB and revealed preference utility. Kimball and Willis (2006) develop a theory that reconciles the two.

Within the literature on SWB, hedonic adaptation has been the focus of substantial recent work.[2] If adaptation and relative income are important in the determination of SWB, a number of implications follow for cross-sectional individual and national data, as well as optimal taxation, consumption, and other topics (Clark et al., 2008b). Concerns about "set points" (SWB levels to which individuals eventually return after a disturbance) and the measurement of happiness are explored in Diener et al. (2006). Clark et al. (2008a) looks for permanent SWB responses to various events, generally not finding significant effects; Lucas (2007), by contrast, finds evidence of incomplete adaptation. Bottan and Truglia (2011) present evidence indicating that happiness itself is positively autocorrelated. Oswald and Powdthavee (2008) and Powdthavee (2009) investigate adaptation in the wake of disability and widowing. They acknowledge the importance of hedonic adaptation but emphasize the mutability and heterogeneity of set points.

Furthermore, data on adaptation is likely to be more useful than data on SWB levels because levels are multi-dimensional (incorporating both life evaluation and temporary mood, for example) and difficult to compare across individuals. Dolan and Kahneman (2008) and Kimball and Willis (2006) offer theories attentive to dynamics and consistent with hedonic adaptation that can further structure the analysis. In Kimball and Willis (2006), intra-person variation in affect, as opposed to life evaluation, provides comprehensive information about the reaction to life events. Our data permit analysis consistent with this theory, facilitating the comparison of various life events. Unlike

[2] For a review of some of the literature, see Frederick and Loewenstein (1999).

many other sources of SWB data, the Health and Retirement Study panel allows for examination of intra-person changes in affect, abstracting from reporting idiosyncrasies and differences in baseline mood.

The choice of SWB measure may matter as well. The German Socioeconomic Panel, for instance, provides information on life satisfaction (e.g., Headey et al., 2010, Lucas et al., 2004) rather than the affect measure we use from the Health and Retirement Study. Benjamin et al. (2014) provide evidence that individuals value a variety of SWB dimensions, including autonomy, health, social status, etc. The authors point to the difficulty of aggregating across these dimensions and across individuals. Luhmann et al. (2012) conduct a meta-analysis of SWB studies, finding that adaptation and effect sizes differ across both life events and the SWB measure used. Specifically, "cognitive well-being" (overall life evaluation) is generally more affected by life events than "affective well-being" (mood).

2.1 Leading Econometric Models

Several authors estimate the impact of life events on happiness using a cross-sectional regression that includes indicator variables for various events. In particular, Deaton et al. (2008) use the following specification with individual-level data from the Gallup World Poll:

$$H_{it} = \alpha_c + X_{it}\beta_X + Y_{it}\beta_Y + \mathbf{1}_{it}\beta_{\mathbf{1}} + \epsilon_{it}, \tag{1}$$

where H is a SWB measure, α_c is a vector of country fixed effects, X is a vector of demographic variables, Y is log income, and $\mathbf{1}$ is a vector of dummies for the death of a family member from a given disease in the last twelve months. Without the possibility of individual fixed effects, the authors rely on the assumption that "baseline mood" (i.e., the SWB level reported before an event) is not systematically different across respondents in a way that is endogenous to the specification.

Other papers exploit country-level data. As mentioned previously, Di Tella et al. (2001) run regressions of the following form.

$$LS_{ct} = \alpha_c + U_{ct}\beta_X + \Pi_{ct}\beta_I + \epsilon_{ct},$$

where LS is a modified life satisfaction measure, U is the unemployment rate, Π is the inflation

rate, and α_c is a vector of country fixed effects. Finkelstein et al. (2008) utilize individual panel data and estimate the following equation, among others:

$$H_{it} = \alpha_i + X_{it}\beta_X + Y_{it}\beta_Y + \mathbf{1}_{it}\beta_{\mathbf{1}} + (\mathbf{1}_{it} * Y_{it})\beta_{\mathbf{int}} + \epsilon_{it}, \tag{2}$$

where H is a SWB measure, α_c is a vector of country fixed effects, X is a vector of demographic variables, Y is log income, and $\mathbf{1}$ is a vector of dummies for whether a respondent has ever had a particular disease. The effect of the interaction between income and health events is given by $\beta_{\mathbf{int}}$. However, this panel approach still fails to distinguish immediate consequences from enduring effects. Because there is evidence that adaptation occurs, we prefer a method that allows for dynamic happiness responses. Further, we find that life events vary significantly in the ratio of temporary to permanent happiness effects and dynamic methods are necessary to properly compare them.

3 Model

The model formulated here is based on a theory of happiness, Kimball and Willis (2006), that interprets the impulse response of SWB to an event as indicating the importance of that event for lifetime utility. Kimball and Willis (2006) are skeptical about the level of happiness as a contemporaneous indicator of utility; feeling happy is only one of many commodities that people care about. In this view, the transitory response of happiness to events is something of interest in its own right, not just a nuisance.

The specific formulation of the model aims to capture the dynamic response of happiness to events in a flexible, tractable, and parsimonious way. This motivates modeling hedonic adaptation by exponential decay – which will prove to fit the HRS data well.[3] The decay rate is estimated simultaneously with the intensity of the initial response of happiness.

3.1 Baseline specification

The simplest form of our dynamic equation is given below. It decomposes the cumulative happiness response into an immediate, temporary effect that decays, and a permanent effect that persists

[3] A nonparametric approach makes excessive demands on the (often sparse and incomplete) data, but a pooled cross-sectional or fixed effects regression fails to recognize the substantial and nonlinear adaptation characteristic of the data. Our specification is nonlinear but quick to estimate.

indefinitely. The temporary effect is assumed to vanish exponentially at rate δ.[4] Since many life events have important consequences for income and wealth or are related to income and wealth levels, we include the log levels of household income and wealth as covariates. The estimating equation is given by

$$H_{it} = \underbrace{\alpha_i}_{\text{fixed effect}} + \underbrace{Y_{it}\beta_Y}_{\text{income effect}} + \underbrace{W_{it}\beta_W}_{\text{wealth effect}} + \chi_i \mathbf{1}(t >= t_0)[\underbrace{\beta_P}_{\text{permanent effect}} + \underbrace{\beta_T e^{-\delta(t-t_0)}}_{\text{temporary effect}}] + \epsilon_{it}, \quad (3)$$

where H is a SWB measure ranging $0 - 100$, χ is a dummy for respondents who have experienced a given event, t is the time that happiness is observed, t_0 is the time the event occurs, β_Y is the income effect, β_W is the wealth effect, β_P is the permanent effect, β_T is the temporary effect, α_i is the person fixed effect, and δ is the rate of decay of the shock.[5] Since δ enters the equation nonlinearly, we use nonlinear least squares (NLS) estimation. The NLS estimator is given by

$$\hat{\theta} = \arg\min_{\theta} \sum_{i=1}^{N} [y_i - f(x_i; \theta)]^2,$$

where $f(x_t; \theta)$ is the nonlinear model, y is the endogenous variable, N is the number of observations, and θ is the parameter vector.

With some SWB-relevant events, it may be the case that the likelihood of occurrence is related to baseline mood. This may be the case even if the event is unanticipated by the respondent. For instance, health problems such as heart attacks may be induced by stress, which could itself imply lower baseline mood. Since the healthy respondents report higher SWB, the effect of a heart attack will be over-estimated with a specification that fails to account for the already-lower baseline mood of respondents who are about to have heart attacks. To account for these differences in baseline mood, we include individual-specific fixed effects.

3.2 Comparability with previous estimates

For some purposes, it may be interesting to consider the cumulative impact of an event. In particular, some of the estimates in previous work (along the lines of equations 1 and 2) have an interpretation as the cumulative SWB consequence of an event. For comparability with these re-

[4] We experimented with less restrictive specifications but found no evidence of non-exponential decay.

[5] All specifications also contains a quadratic in respondent age, not shown.

sults, we adapt our baseline specification, equation 3, by integrating to find the total effect. For an individual with d annual mortality risk and interest rate r, the "area under the curve" will have the following form:

$$\int_{t_0}^{\infty} (\beta_P e^{-(d+r)(s-t_0)} + \beta_T e^{-(\delta+d+r)(s-t_0)}) ds = \frac{\beta_P}{d+r} + \frac{\beta_T}{d+r+\delta} \tag{4}$$

This formulation gives a single statistic that can be used to rank events by their hedonic importance. It also allows for comparison of our dynamic results with the previous literature's static estimates, since both are measures of a cumulative hedonic effect. Following previous work, the fixed effects regression

$$H_{it} = \alpha_i + Y_{it}\beta_Y + \mathbf{1}_{it}\beta_{\mathbf{1}} + \epsilon_{it}, \tag{5}$$

is conducted to make this comparison of our results with the usual linear specification. $\mathbf{1}_{it}$ is an "absorbing state" indicator, which means that it is set to 1 for all observations after the initial event occurrence. This is consistent with some of the previous SWB literature (e.g., Finkelstein et al., 2008) and aims to capture a cumulative SWB effect. Were $\mathbf{1}_{it}$ to equal 1 only in the initial event occurrence observation, $\beta_{\mathbf{1}}$ would capture only a portion of the temporary SWB consequences, and would not be comparable to our cumulative results.

The chief virtue of our baseline specification is its ability to separately identify temporary and permanent SWB effects. While important in its own right, mean reversion also complicates estimation of the cumulative happiness effect. In equation 5, $\beta_{\mathbf{1}}$ is only identical to β_P in equation 3 if $\beta_T = 0$. For $\beta_T > 0$, equation 5 may imply a different estimate of the cumulative SWB effect than equation 3. Consider the ordinal ranking of SWB-relevant events by $\beta_{\mathbf{1}}$. Since the specification of equation 5 makes no use of time since occurrence, this parameter will depend on the probability that a respondent subsequently leaves the sample, which may differ across events. For example, an event that is substantially mean-reverting may be associated with subsequent SWB observations over many years. Another event, with the same balance of temporary and permanent effects, may have relatively few subsequent SWB observations. Assume, for specificity, that both temporary and permanent effects are negative in sign. For the event with few post-occurrence observations, $\beta_{\mathbf{1}}$ will be estimated to be larger in magnitude, because the temporarily depressed SWB values

just after event occurrence will dominate the data. By contrast, β_T, β_P, and by implication $\beta_{cumulative} = \frac{\beta_T}{d+r+\delta} + \frac{\beta_P}{d+r}$ do not depend on these factors, but are estimated consistently if the underlying model is correct. Rankings of SWB-relevant events based on the latter expression will then be different, in general, than a ranking based on $\beta_{\mathbf{1}}$.

When there is a permanent component to the change in happiness after an event, two different interpretations of these results correspond to two different views on happiness. If happiness is a contemporaneous indicator of utility, then the total area under the curve relative to the initial baseline is of interest. By contrast, in the Kimball-Willis (2006) theory, the extra lifetime utility from the permanent component is already accounted for in the magnitude and duration of the transitory component. So, to avoid double-counting, only the area under the transitory component should be counted in gauging the importance of an event. We report both measures.

3.3 Allows for various sorts of decompositions

One interesting extension of this method involves subjective life expectancy. Some life events, in particular health events like heart attacks and strokes, will in general have important consequences for life expectancy. The original dynamic specification, equation 3, can be modified to examine the role of subjective life expectancy changes in creating the observed hedonic effects. A particular implementation is given below:

$$H_{it} = \alpha_i + Y_{it}\beta_Y + \chi_i \mathbf{1}(t >= t_0)[\beta_P + (\beta_T + \eta \Delta SLE)e^{-\delta(t-t_0)}] + \epsilon_{it}, \tag{6}$$

where SLE is a constructed measure of subjective life expectancy, denominated in years, and η is the temporary effect of changes in subjective life expectancy.

3.4 Anticipation

In order for our interpretation of the baseline specification to be correct, it must be the case that the "news" component of event happiness consequences does not occur prior to the event itself. In other words, respondents must not learn of and (in terms of subjective well-being) react to an event that has yet to occur. Since this assumption is likely violated in a number of cases, we include as a control a dummy for SWB in the six-month period prior to an event. This is done for all the

events that are precisely dated. If news about the event generates a change in SWB prior to the event date, our modified specification will correctly capture the consequences of the event itself (as opposed to news of the event) in the parameters β_P, β_T, δ.

3.5 Sparse data

Another advantage of our method is that it can be easily modified to handle infrequently-measured data. With all the HRS data, we posit an underlying continuous-time data generating process, with the happiness data only observed periodically (every two years in the core of the HRS). Some of the SWB-relevant events are dated precisely to the month in which they occur, others are known only to the calendar year in which they occurred (with extra information for same-year events coming from the fact that they must be before the survey date), while still others can only be dated as occurring sometime between waves of survey data. To deal with this, we assume a uniform distribution of the logically possible interval of time in which an event could have occurred given the data. We time-aggregate the equations for the continuous-time data generating process to obtain a nonlinear estimating equation. The key identifying assumption is that there are no important transitory movements in baseline mood after an event that are correlated with the event itself.

For instance, events that are dated only to the year are estimated by the following equation.

$$H_{it} = \alpha_i + Y_{it}\beta_Y + \chi_i \mathbf{1}(y >= y_0)[\beta_P + \beta_T e^{-\delta(t-t_0)} \cdot (\underbrace{\mathbf{1}(y > y_0) \cdot \frac{e^{\delta-1}}{\delta}}_{\text{post-year}} + \underbrace{\mathbf{1}(y = y_0) \cdot \frac{e^{\delta * \frac{m}{12} - 1}}{\delta}}_{\text{same-year}})] + \epsilon_{it}, \quad (7)$$

where y is the year happiness is measured, y_0 is the year an event occurred, and m is the interview month in which happiness is measured.[6] The "post-year" component isolates the possibility that happiness is observed in a year that follows the year in which the event occurred, integrating over the entire distribution of possible months, while the "same-year" component integrates over the distribution that is possible given that an event must have occurred prior to the interview month in which happiness was recorded.

For events that are dated only to the wave,

[6] By contrast to the precisely-dated case, t_0 is now constructed with the assumption that the event occurred in January of a given year. The subsequent "post-year" and "same-year" components make adjustments consistent with an underlying assumption of a uniform distribution of the possible event occurence interval.

$$H_{it} = \alpha_i + Y_{it}\beta_Y + \chi_i \mathbf{1}(y \geq y_0)[\beta_P + \beta_T e^{-\delta(t-t_0)} * \frac{e^{(t_{w_1}-t_{w_0})\delta-1}}{\delta}] + \epsilon_{it}, \tag{8}$$

where t_{w_1} is the time of the interview directly after an event occurred and t_{w_0} is the time of the interview directly before.

4 Data

We use data from the Health and Retirement Study (HRS), which conducts a biennial representative survey of Americans over the age of 50. The resulting panel data spanning the years 1992 through 2012 includes detailed happiness reports and information on a rich set of important life events. Although there are important panel data sets for subjective well-being for other countries (most notably the German Socioeconomic Panel and the British Household Panel Study), for the U.S., the HRS is the only survey with a long panel of repeated observations on subjective well-being. In addition to the core HRS waves, we use the Asset and Health Dynamics Among the Oldest Old (AHEAD), the Children of the Depression Age (CODA), and the War Babies cohorts[7]. For some of the variables, we use a version of this data provided by the RAND Center for the Study of Aging that includes some additional imputations.

Each wave of the HRS asks respondents the following questions: "Now think about the past week and the feelings you have experienced. Please tell me if each of the following was true for you much of the time this past week: a) You felt you were happy b) You felt sad c) You enjoyed life d) You felt depressed." We treat the binary variables "happy" and "enjoylife", along with the reverse-coded "notsad" and "notdepressed", as four indicators of the underlying latent value of happiness at the time of the interview. Thus, we treat the probability of answering in the positive direction for each of these variables as an increasing function of latent happiness.

HRS respondents are questioned about many health and other important life events in each wave of the survey. Some of these events are precisely dated to the month but some are only known to have occurred at some point between waves. Examples of the latter include episodes of incontinence, congestive heart failure, hip fractures, cataract surgery, births of grandchildren, and

[7]The AHEAD cohort was initially part of a distinct study and includes respondents born before 1924. CODA and the War Babies cohorts were added in 1998 and includes respondents born 1924-1930, and 1942-1947, respectively.

changes in social isolation. Widowing, heart attacks, strokes, cancer, retirement, unemployment, and entry into nursing homes are dated to the precise month. The Psychosocial Leave-Behind (PLB) component of the HRS provides retrospective information on a number of other events with dating only to the year: death of a child, serious illness of a family member, drug and alcohol addiction of family members, physical assault, labor market discrimination, police discrimination, job search, changes in neighborhood safety, and others.

Measures of household income and wealth are available for all waves, which we include as controls throughout. The HRS includes life insurance status, allowing us to partition the hedonic response to widowing. Interestingly, there are also questions about life expectancy for many of the respondents. We use these to construct a measure of subjective life expectancy, then decompose the temporary response to an event into a component related to changes in subjective life expectancy and a residual, which becomes the standard β_T coefficient. The subjective life expectancy (SLE) measure is constructed in the following way. For certain waves, respondents are asked what probability they assign to their living to a particular age, where said age depends on the current age of the respondent. We linearly interpolate the survival probability for all future years, then calculate SLE as the expectation of years remaining.

5 Results

Because the HRS is a panel, we are able to conduct fixed effects estimation. Without individual fixed effects, any heterogeneity in baseline happiness might bias estimation. For example, if heart attacks tend to happen to people with lower baseline happiness, our approach (modified to exclude fixed effects) would recover an exaggerated estimate of the permanent effect. With fixed effects, variation in baseline happiness is not confounded with the enduring consequence of a life event.

As discussed in Section 4, the HRS provides four closely-related variables pertaining to a respondent's mood at the time of interview. Our preferred approach uses the sum of the (appropriately-coded) variables as the dependent variable in all our specifications. In preliminary work, we compared multiple approaches to the use of HRS subjective well-being variables. We obtained broadly similar results when conducting probit estimation with individual SWB variables.

Table 1 gives estimates for β_P, β_T, and δ using our dynamic method for events dated to the

month (widowing, heart attack, stroke, cancer, unemployment, nursing home entry, and retirement). Standard errors are all heteroskedasticity-robust. Figure 2a-3c illustrates the results graphically, showing the predicted values of both a non-parametric regression and an "impulse response" corresponding to our baseline specification for events dated to the month.[8] The impulse responses are a graphical depiction of the estimated parameters β_T, β_P, and δ. In other words,

$$\widehat{H}_t = \overline{H}_b + \chi_i \mathbf{1}(t >= 0)[\widehat{\beta_P} + \widehat{\beta_T} e^{-\hat{\delta} \cdot t}], \tag{9}$$

where $\overline{H}_b$ is the mean pre-event SWB level and $\widehat{H}_t$ is the predicted SWB value. This yields a graphical representation of the temporary and permanent effects at work in the regressions.

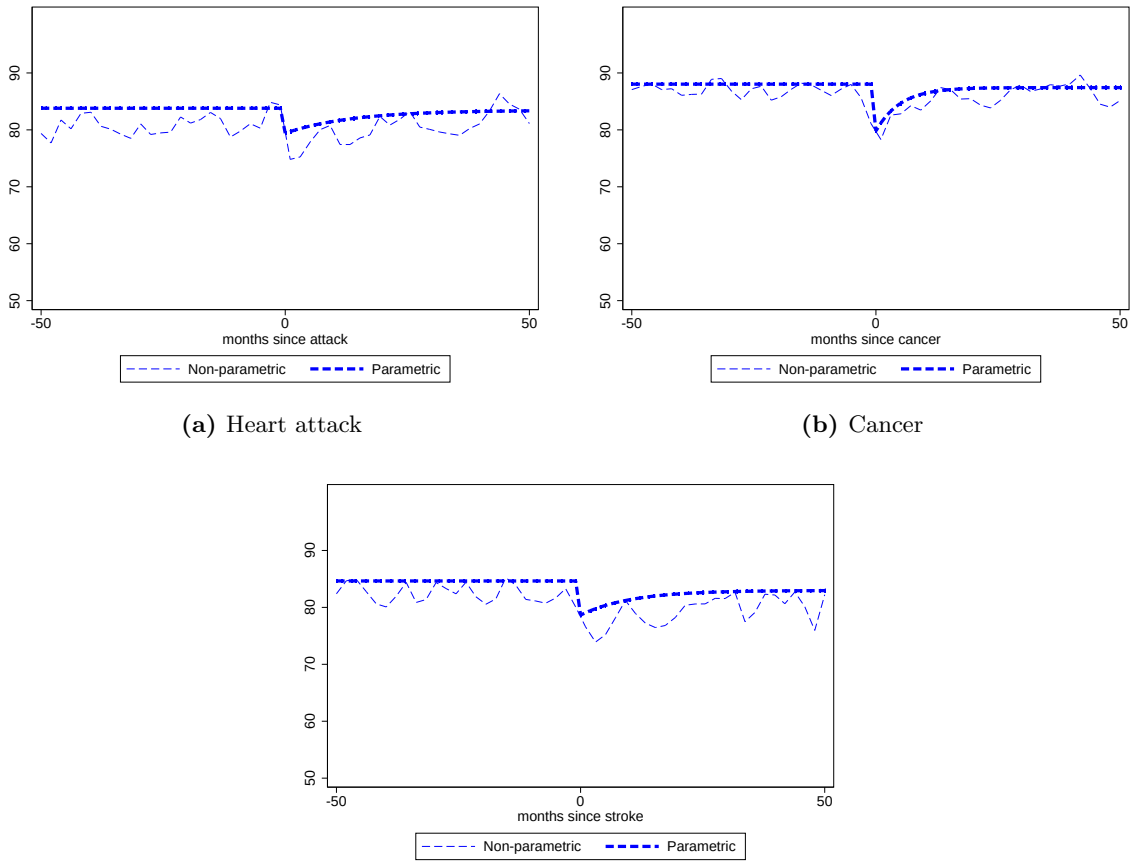

(a) Heart attack

(b) Cancer

(c) Stroke

Figure 2: Baseline and non-parametric results for precisely-dated events

Cancer, heart attacks, and strokes each follow a somewhat different pattern. We estimate a

[8] We use a kernel-weighted local polynomial regression for the non-parametric graph.

negative permanent effect in all three cases, with somewhat larger temporary than permanent effects. Interestingly, strokes have both the largest permanent effect, relative to temporary (as well as the largest permanent effect in absolute value), which is plausibly consistent with the enduring disability often suffered by stroke victims. The temporary effect of cancer is less than one-third as large as the temporary effect of widowing.

Regressions concerning the other events for which we have precise timing information - unemployment, nursing home entry, and retirement - yield reasonable estimates. Unemployment has a negative temporary and effect that is roughly in line with the magnitude of the aforementioned health problems. Entry into a nursing home produces even larger temporary and permanent reduction in SWB. Involuntary retirement produces a smaller negative SWB effect that dissipates almost entirely. Both nursing home entry and retirement are perhaps less plausibly exogenous than the other events, and it may be that correlated factors are driving those estimates.

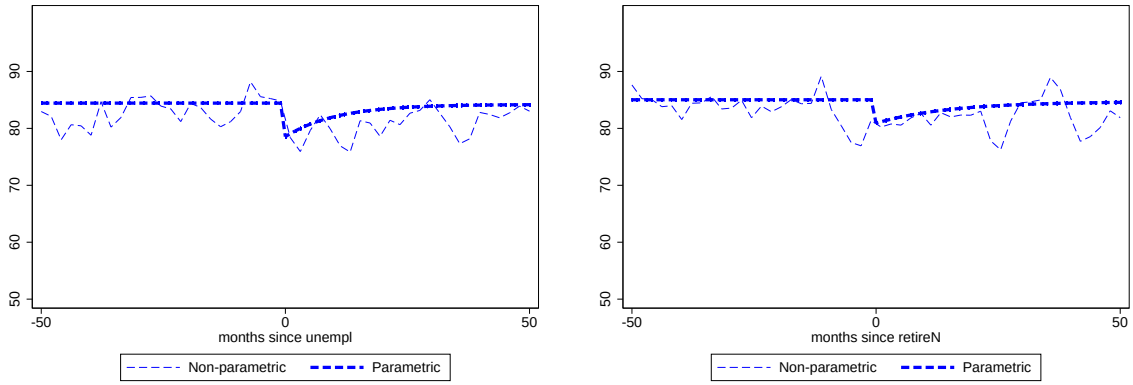

(a) Unemployment **(b)** Involuntary retirement

Familial death and dissolution are also quite relevant to happiness, especially in the short-term. More generally, it is notable how short-lived the happiness consequences of some events appear: half the temporary effect of a parent's death has disappeared by two months, while widowing requires about a year for the same recovery. The estimated depreciation rates are themselves of interest. For all the precisely-dated events, estimated depreciation rates are such that the half-life of an event ranges from about 2 to 13 months. That is, half of the temporary shock associated with an event has disappeared by this time.

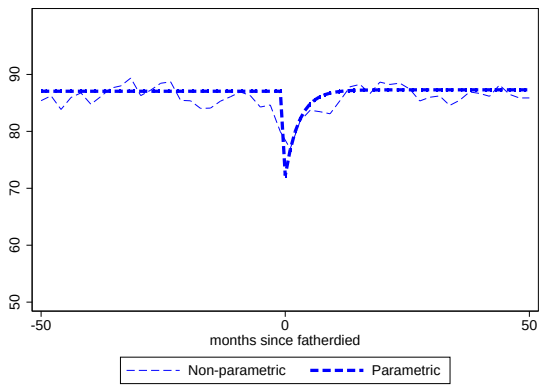

(a) Father died

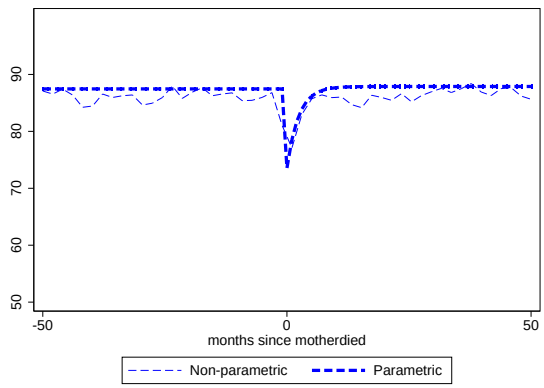

(b) Mother died

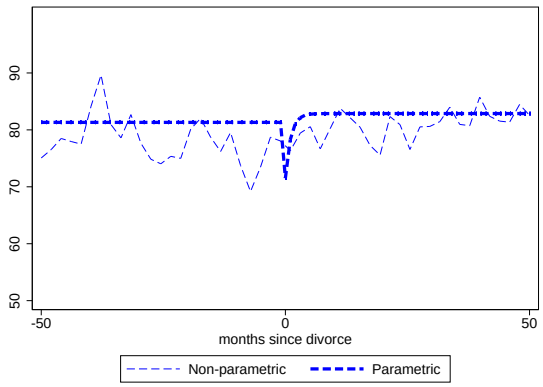

(c) Divorce

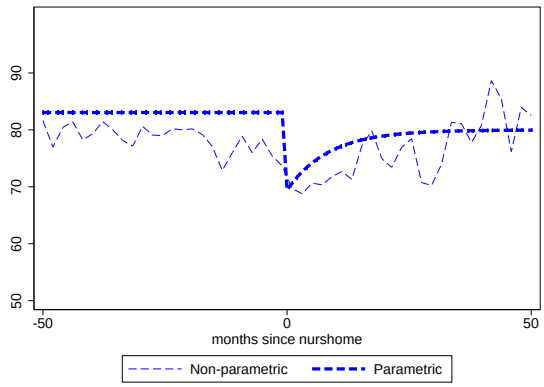

(d) Entered nursing home

Widowing is of particular interest, as widows experience an unusually large reduction in SWB, all of which appears to be temporary. Individuals with life insurance experience roughly 80% of the temporary drop in SWB suffered by those without insurance, though the latter are estimated to recover slightly more quickly from their (larger) fall. This is depicted, both parametrically and non-parametrically, in Figure 3. In both cases, we identify a very large effect that is temporary.

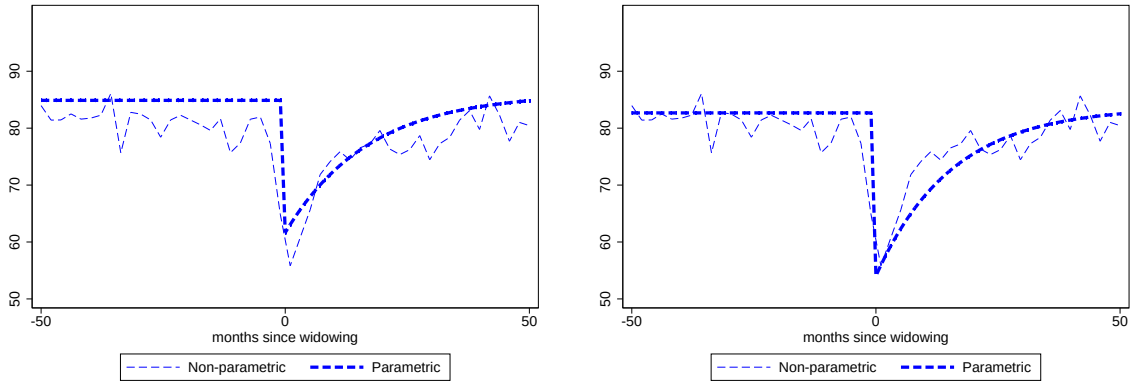

(a) With life insurance **(b)** Without life insurance

Figure 3: Widowing with and without life insurance

Table 2 compares results from equation 5 and our dynamic method. The first three columns show estimates (identical to Table 1) for β_T, β_P, and δ, respectively. The fourth, fifth, and sixth columns give the cumulative temporary, permanent, and total effects, whose construction was described previously. The mortality rate d is projected using Social Security actuarial tables and the age-sex composition of our HRS sample. It is 0.057, while the interest rate stipulated is 0.05.

The "Lost Area" columns, 4, 5, and 6 (so-called because they show the "area under the curve" associated with the hedonic response to an event), condense the information provided in the baseline specification, giving measures of cumulative SWB consequences. A high depreciation rate renders the cumulative temporary effect of cancer relatively small in magnitude, while magnifying the permanent effect by comparison. Strokes have a roughly similar cumulative effect compared to widowing, while heart attacks and cancer are substantially less damaging.

We also implement equation 5, the linear regression described in Section 2.1. The final column displays $\beta_{\mathbf{1}}$ from that regression. These results are generally signed consistently with the Total Lost Area quantities, but relative magnitudes differ. Recall that $\mathbf{1}$ is an indicator that toggles on

permanently after an event. $\beta_{\mathbf{1}}$ registers larger magnitude effects for widowing than for most health events, with the most negative result coming from nursing home entry at -5.9.

Table 3 presents results that utilize data on subjective life expectancy for the same life events. This is an example of the flexibility of our approach, which facilitates any decomposition of SWB effects permitted by the data. The first four columns provide the usual parameter estimates plus η, the effect of a one-year increase in subjective life expectancy on the temporary hedonic effect. The final three columns give estimates from the previous specification for comparison. In all cases where we find a statistically significant result, the effect of η is as expected: an event that increases life expectancy will increase SWB through this channel. The magnitude of these life expectancy effects is small, however, and the available HRS variables pertaining to life expectancy are crude. Respondents are asked for their subjective probability of living to a given age, which gives only limited information about the distribution of longevity expectations.

Table 4 gives estimates of β_P, β_T, and δ for events that are dated relatively imprecisely: either to the year or to the wave. In the former case, these are retrospective data from the Psychosocial Leave-Behind survey. Because there are often long gaps between event occurrence and recollection, the estimates produced from this data are somewhat less reliable, which is reflected in the occasional inability of the NLS procedure to identify a δ significantly above zero. In these cases, the data does not permit the separation of temporary and permanent SWB effects, and consequently β_T and β_P values should likely be disregarded.

6 Discussion and Conclusion

In Section 3.2, conditions were described that may produce discrepancies in the SWB ranking (i.e., the ordering by magnitude of the "Total" cumulative effect) based on our baseline specification, and on a regression that does not incorporate SWB mean reversion. Table 2 suggests some concrete examples. Interestingly, within the three major health events we consider (heart attacks, strokes, and cancer), our cumulative SWB ranking mirrors the ranking generated by equation 5. However, some discrepancies exist for other events: for example, according to the conventional approach, the SWB effect of unemployment is larger than that of cancer or heart attacks. With our approach, the total cumulative effect of unemployment is smaller.

The time path of happiness can be informative in other ways. Applied to parental death, the pooled specification yields negligible effects. Our decomposition provides some insight into what is likely driving this result: in the pooled regression, large but very fleeting negative temporary effects are obscured by slightly positive permanent effects. Without such a decomposition, it would appear that parental death has negligible or even positive consequences for respondent happiness - a result that does not appear particularly plausible. Similarly, the large, but quite fleeting, negative temporary effect of divorce is concealed (in a pooled regression) by the positive permanent effect. Our decomposition shows that divorce is not unambiguously positive for SWB.

This illustrates that the temporary/permanent decomposition of SWB data is sometimes a first-order concern. The overall utility consequence of an event, good, or experience is quite sensitive to this decomposition, since temporary effects are generally fairly quick to decay. Without arbitrarily-frequent and indefinitely-extended panels of SWB measurements, a regression specification that is not sensitive to hedonic dynamics will typically generate errors, particularly when comparing events of dissimilar dynamic profiles. If one of the aims of happiness economics is to inform public policy about relative valuations of events and non-market goods, insensitivity to dynamic effects will compromise that project.

The approach presented here will be useful in a variety of ways. Work that aims at pricing non-market goods, many of which take on characteristics of durable goods and produce a changing flow of utility, will benefit from the explicit treatment of dynamic effects. When data is infrequently-collected, retrospective, or otherwise limited, our approach will facilitate the extraction of usable information. Since SWB variables have only recently been added to some datasets, the ability to handle retrospective data is likely to be useful.

References

Andrews, F. and S. Withey (1976): *Social indicators of well-being: Americans' perceptions of life quality*, New York: Plenum.

Benjamin, D., O. Heffetz, M. Kimball, and A. Rees-Jones (2012): "What do you think would make you happier? What do you think you would choose?" *American Economic Review*, 102, no. 5, 2083–110.

Benjamin, D., O. Heffetz, M. Kimball, and N. Szembrot (2014): "Beyond happiness and satisfaction: toward well-being indices based on stated preference," *American Economic Review*, forthcoming.

Bottan, N. and R. Truglia (2011): "Deconstructing the hedonic treadmill: is happiness autoregressive?" *Journal of Socio-Economics*, 40, 224–36.

Bradburn, N. (1969): *The structure of psychological well-being*, Chicago: Aldine.

Campbell, A., P. Converse, and W. Rodgers (1976): *The quality of American life*, New York: Russell Sage.

Clark, A., E. Diener, Y. Georgellis, and R. Lucas (2008a): "Lags and leads in life satisfaction: a test of the baseline hypothesis," *The Economic Journal*, 118, F222–F243.

Clark, A., P. Frijters, and M. Shields (2008b): "Relative income, happiness, and utility: an explanation for the Easterlin aradox and other puzzles," *Journal of Economic Literature*, 46, 95–144.

Deaton, A., J. Fortson, and R. Tortora (2008): "Life (evaluation), HIV/AIDS, and death in Africa," *Working paper*.

Di Tella, R., R. MacCulloch, and A. Oswald (2001): "Preferences over inflation and unemployment: evidence from surveys of happiness," *The American Economic Review*, 91, 335–41.

Diener, E., R. E. Lucas, and C. Napa (2006): "Beyond the hedonic treadmill: revising the adaptation theory of well-being," *American Psychologist*, May-June, 305–14.

Dolan, P. and D. Kahneman (2008): "Interpretations of utility and their implications for the valuation of health," *The Economic Journal*, 118, 215–34.

Finkelstein, A., E. Luttmer, and M. Notowidigdo (2008): "What good is wealth without health? The effect of health on the marginal utility of consumption," *NBER Working Paper*, 14089.

Frederick, S. and G. Loewenstein (1999): *Well-being: the foundations of hedonic psychology*, Russell Sage, chap. Hedonic adaptation.

Gruber, J. and S. Mullainathan (2005): "Do cigarette taxes make smokers happier," *Advances in Economic Analysis and Policy*, 5.

Headey, B., R. Muffels, and G. Wagner (2010): "Long-running German panel survey shows that personal and economic choices, not just genes, matter for happiness," *Proceedings of the National Academy of Sciences, Early Edition.*

Kimball, M. and R. Willis (2006): "Utility and Happiness," *University of Michigan manuscript.*

Lucas, R., A. Clark, Y. Georgellis, and E. Diener (2004): "Unemployment alters the set point for life satisfaction," *Psychological Science*, 15, 8–13.

Lucas, R. E. (2007): "Adaptation and the set-point model of subjective well-being: does happiness change after major life events?" *Current Directions in Psychological Science*, 16, 75–79.

Luhmann, M., W. Hofmann, M. Eid, and R. E. Lucas (2012): "Subjective well-being and adaptation to life events: a meta-analysis," *Journal of Personality and Social Psychology*, 102, 592–615.

Luttmer, E. F. (2005): "Neighbors as negatives: relative earnings and well-being," *Quarterly Journal of Economics*, 120, 963–1002.

Oswald, A. and N. Powdthavee (2008): "Does happiness adapt? A longitudinal study of disability with implications for economists and judges," *Journal of Economics*, 92, 1061–77.

Powdthavee, N. (2009): "What happens to people before and after disability? Focusing effects, lead effects, and adaptation in different areas of life," *Social Science & Medicine*, 69, 1834–44.

Riis, J., G. Loewenstein, J. Baron, C. Jepson, A. Fagerlin, and P. Ubel (2005): "Ignorance of hedonic adaptation to hemodialysis: a study using ecological momentary assessment," *Journal of Experimental Psychology*, 134, 3–9.

Schwarz, N. (1987): *Mood as Information: On the Impact of Moods on Evaluation of Ones Life*, Heidelberg: Springer-Verlag.

Schwarz, N. and F. Strack (1999): *Well-being: The foundations of hedonic psychology*, New York: Russell-Sage, chap. Reports of subjective well-being: Judgmental processes and their methodological implications, 61–84.

Stevenson, B. and J. Wolfers (2009): "The Paradox of Declining Female Happiness," *American Economic Journal: Economic Policy*, 1, 190–225.

Table 1: Baseline results

	β_T	β_P	δ	N
Widowing (w/ insurance)	-24.48	1.253	0.704	4476
	(1.825)	(0.713)	(0.0923)	
Widowing (w/o insurance)	-29.65	1.023	0.773	2708
	(2.561)	(0.813)	(0.113)	
Cancer	-7.615	-0.603	2.414	25997
	(1.232)	(0.309)	(0.651)	
Heart attack	-4.034	-0.351	0.883	12866
	(1.070)	(0.447)	(0.470)	
Stroke	-4.329	-1.660	1.168	14192
	(1.261)	(0.525)	(0.643)	
Unemployment	-5.660	-0.228	1.158	12859
	(0.961)	(0.274)	(0.434)	
Entered nursing home	-10.54	-3.071	1.408	9334
	(2.204)	(1.145)	(0.639)	
Involuntary retirement	-3.787	-0.332	0.806	14566
	(1.192)	(0.446)	(0.446)	
Father died	-15.07	0.254	4.245	52503
	(1.936)	(0.267)	(0.810)	
Mother died	-14.33	0.427	5.043	62360
	(1.374)	(0.228)	(0.759)	
Divorce	-11.33	1.551	10.88	5488
	(3.380)	(0.614)	(6.649)	

Note: The dependent variable is the (0-100) index of happiness equal to 25*(sum of the four indicators of recent mood). See the text for a description of the indicators. δ is expressed as an annual rate. Standard errors are in parentheses. All events are dated to the month.

Table 2: Comparison of event study results with pooled

	Parameter Estimates			Lost Area			
	β_T	β_P	δ	Temp	Perm	Total	Pooled coefficient
Widowing (w/ insurance)	-24.48	1.253	0.704	-30.18	11.71	-18.47	-2.718
	(1.825)	(0.713)	(0.0923)	(2.252	(6.848)		(0.790)
Widowing (w/o insurance)	-29.65	1.023	0.773	-33.71	9.562	-24.15	-4.150
	(2.561)	(0.813)	(0.113)	(2.910)	(7.909)		(1.104)
Cancer	-7.615	-0.603	2.414	-3.021	-5.636	-8.657	-1.473
	(1.232)	(0.309)	(0.651)	(0.483)	(2.814)		(0.310)
Heart attack	-4.034	-0.351	0.883	-4.075	-3.280	-7.355	-1.277
	(1.070)	(0.447)	(0.470)	(1.070)	(4.078)		(0.419)
Stroke	-4.329	-1.660	1.168	-3.397	-15.52	-18.91	-2.563
	(1.261)	(0.525)	(0.643)	(0.976)	(4.900)		(0.441)
Unemployment	-5.660	-0.228	1.158	-4.474	-2.130	-6.605	-1.531
	(0.961)	(0.274)	(0.434)	(0.753)	(2.627)		(0.493)
Entered nursing home	-10.54	-3.071	1.408	-6.953	-28.70	-35.66	-5.880
	(2.204)	(1.145)	(0.639)	(1.445)	(10.63)		(0.806)
Involuntary retirement	-3.787	-0.332	0.806	-4.149	-3.098	-7.248	-0.786
	(1.192)	(0.446)	(0.446)	(1.298)	(4.138)		(0.437)
Father died	-15.07	0.254	4.245	-3.463	2.377	-1.086	-0.101
	(1.936)	(0.267)	(0.810)	(0.433)	(2.407)		(0.346)
Mother died	-14.33	0.427	5.043	-2.783	3.987	1.204	-0.133
	(1.374)	(0.228)	(0.759)	(0.258)	(2.037)		(0.257)
Divorce	-11.33	1.551	10.88	-1.031	14.50	13.46	1.435
	(3.380)	(0.614)	(6.649)	(0.307)	(5.710)		(0.663)

Note: The dependent variable is the (0-100) index of happiness equal to 25*(sum of the four indicators of recent mood). See the text for a description of the indicators. δ is expressed as an annual rate. Standard errors are in parentheses. All events are dated to the month. Pooled coefficients are from a fixed effects regression of happiness on log household income, log household wealth, and an indicator for whether the event has ever occurred previous to or in the current wave. Area measures are based on an interest rate of .05 and a constant mortality rate of .057, the latter of which is based on Social Security actuarial projections.

Table 3: Results with subjective life expectancy

	β_T	β_P	δ	η
Widowing (w/ insurance)	-26.48	1.733	0.654	-0.0734
	(3.693)	(2.000)	(0.182)	(0.0830)
Widowing (w/o insurance)	-32.09	1.604	1.230	-0.405
	(6.680)	(2.216)	(0.429)	(0.322)
Cancer	-8.654	0.0190	2.407	-0.00273
	(1.743)	(0.125)	(0.766)	(0.0227)
Heart attack	-6.129	1.028	1.574	0.0835
	(2.267)	(1.127)	(0.943)	(0.0377)
Stroke	-5.252	-2.113	1.166	0.0735
	(2.071)	(1.194)	(0.702)	(0.0390)
Unemployment	-9.085	3.117	0.729	0.00359
	(1.699)	(1.665)	(0.326)	(0.0205)
Entered nursing home	-12.42	0.837	1.447	0.391
	(4.735)	(2.589)	(1.041)	(0.264)
Involuntary retirement	-4.460	0.185	0.687	-0.00415
	(1.840)	(0.956)	(0.534)	(0.0162)
Father died	-15.61	-0.910	5.586	-0.00322
	(3.041)	(0.625)	(1.610)	(0.0291)
Mother died	-14.19	-0.0458	4.897	0.0676
	(1.812)	(0.144)	(0.944)	(0.0338)
Divorce	-11.53	4.890	0.817	0.0681
	(3.175)	(2.397)	(0.415)	(0.0562)

Note: The dependent variable is the (0-100) index of happiness equal to 25*(sum of the four indicators of recent mood). See the text for a description of the indicators. δ is expressed as an annual rate. Standard errors are in parentheses. All events are dated to the month. η is the effect of an additional year of subjective life expectancy.

Table 4: Baseline results with imprecisely-dated events

	Parameter Estimates			
	β_T	β_P	δ	N
Events dated to the year				
Death of child	-14.17	4.433	0.277	7302
	(1.458)	(1.195)	(0.0592)	
Family illness	-9.490	3.466	0.233	10448
	(1.167)	(1.062)	(0.0588)	
Family member addiction	-5.841	3.709	0.0363	5163
	(3.736)	(3.691)	(0.0406)	
Fired	-5.865	1.798	0.531	2038
	(2.885)	(2.510)	(0.521)	
Moved to worse neighborhood	11.00	-3.679	0.476	524
	(7.805)	(6.955)	(0.669)	
Respondent illness	-6.251	5.822	0.0471	9912
	(2.138)	(2.221)	(0.0228)	
Unemployed more than 3 months	-5.222	1.538	1.093	1680
	(2.822)	(2.221)	(1.147)	
Serious physical assault	-10.87	11.34	0.642	1780
	(3.842)	(2.369)	(0.387)	
Unfairly denied promotion	-8.135	5.039	0.0619	2414
	(3.554)	(3.682)	(0.0474)	
Unfairly not hired	-2.807	0.739	0.902	1896
	(3.357)	(2.199)	(2.105)	
Unfairly treated by police	1.181	3.278	1.102	1314
	(3.870)	(2.399)	(4.643)	
Combat experience	-42.55	4.360	0.586	1586
	(10.73)	(2.545)	(0.281)	

.

Baseline results with imprecisely-dated events, continued

	Parameter Estimates			
	β_T	β_P	δ	N
Events dated to the wave				
Cataract surgery	-0.142	211.7	0.0002	40971
	(0.140)	(2529)	(0.00296)	
Child moved within 10 miles	-0.404	0.875	0.161	44622
	(0.205)	(1.091)	(0.244)	
Hip fracture	-2.332	9.841	0.0683	5249
	(0.665)	(17.34)	(0.105)	
Incontinence	-1.649	1.865	0.247	76492
	(0.405)	(0.909)	(0.136)	

Note: The dependent variable is the (0-100) index of happiness equal to 25*(sum of the four indicators of recent mood). See the text for a description of the indicators. δ is expressed as an annual rate. Standard errors are in parentheses.