

Valuing Public Goods Using Happiness Data: The Case of Air Quality

Arik Levinson
Georgetown University and NBER
AML6@georgetown.edu

May 9, 2012

Abstract

This paper describes and implements a method for valuing a time-varying local public good: air quality. It models survey respondents' self-reported happiness as a function of their demographic characteristics, incomes, and the air pollution and weather on the date and in the place they were surveyed. People with higher incomes report higher levels of happiness, and people interviewed on days with worse local air pollution report lower levels of happiness. Combining these two concepts, I derive the average marginal rate of substitution between income and current air quality – a compensating differential for short-term changes in air pollution.

JEL codes: Q51, Q53, H41

Key words: willingness to pay, stated well-being, pollution, compensating differential

This research is part of a project funded by the National Science Foundation grant #0617839. I am grateful to Resources for the Future for its hospitality during part of the time I worked on the paper, to Emma Nicholson and Jamie O'Brien for superb research assistance, to Sarah Aldy for editorial support, to the research staff at the General Social Survey for assisting me with matching confidential aspects of their data to geographic information, and to John Helliwell, Chris Barrington-Leigh, Simon Luechinger, Erzo Luttmer, Karl Scholz and Heinz Welsch for helpful suggestions.

Valuing Public Goods Using Happiness Data: The Case of Air Quality

1. Introduction

Valuing local public amenities and other non-market goods is one of the greatest challenges facing applied economics. Existing methods, often applied to environmental quality, include travel-cost models, hedonic regressions of property values, and contingent valuation surveys in which people are asked directly their willingness to pay for public goods. In this paper, I describe and test an alternative method for estimating the economic benefit of a local public good. The fundamental idea is extraordinarily simple. I combine survey data with air quality and weather information to model individuals' self-reported levels of "happiness," or "subjective well-being," as a function of their demographic characteristics, incomes, and the air quality and weather at the date and place they were surveyed. I then use the estimated function to calculate a marginal willingness to pay, or compensating differential, for air pollution: the average marginal rate of substitution between annual household income and current air quality that leaves respondents equally happy.

This happiness-based methodology has a number of advantages over existing tools for valuing environmental quality. The people most averse to air pollution choose to visit and live in clean locales; as a result, travel-cost and many hedonic models may underestimate the value of air quality. But because I include fixed effects and interactions by time and place, coefficients are identified from daily fluctuations in pollution within a location and are not subject to these sorting biases. Because I estimate marginal rates of substitution between income and pollution directly, income effects do not confound the approach, nor do large gaps between measures of willingness to pay and willingness to accept. And because I do not rely on asking people directly about environmental issues, the methodology is not susceptible to the strategic biases and framing problems of the contingent valuation approach.

Furthermore, although happiness studies have recently been used to estimate tradeoffs made by public policies, including valuations of public goods and bads, all of the previous work has relied on annual average values across regions or countries.¹ If the public goods are

¹ Public policy issues studied have included price inflation (Di Tella et al. 2001), state cigarette taxes (Gruber and Mullainathan 2005), airport noise (van Praag and Baarsma 2005), inequality (Alesina et al. 2004), terrorism (Frey et

simultaneously determined by regional characteristics also associated with happiness, studies using annual regional differences in public goods will yield biased estimates of their value. Air quality, on the other hand, varies daily within each location for reasons less likely to be connected to any particular respondent's situation. The results here are identified entirely from short-term changes in air quality at a given location, and consequently they mitigate concerns about unobserved local characteristics correlated with both happiness and air quality.

Because air quality is a public good that fluctuates day-to-day, the results here will need to be interpreted somewhat differently from typical valuations of public goods. If people become habituated to levels of public goods, estimates based on daily fluctuations will yield higher values than estimates based on longer-term levels, because air quality presumably changes more quickly than people become habituated. Any valuation derived from daily fluctuations will omit the effect of habituation, but also will omit any long-term effects of poor air quality because those will be absorbed by time and place fixed effects. In the extreme case of perfect habituation, there might well be significant differences in happiness in any one place on days with low or high air quality, but no happiness differences across otherwise similar places with different average levels of air quality.² That is in fact what I find.

Naturally, this approach also has disadvantages. It treats responses to questions about happiness as a proxy for utility and then makes interpersonal comparisons among respondents. It relies on a vague question about how "things are these days." It identifies the relevant compensating differential based on trade-offs between fluctuations in daily pollution and differences among respondents' annual incomes. And it takes household income to be an exogenous determinant of happiness, rather than potentially determined by happiness. The reason to pursue this line of research, therefore, is not that it is without shortcomings. Instead, the attractive feature of this approach is that its shortcomings differ so markedly from those of standard approaches to valuing public goods, and therefore it serves as a useful point of comparison.

I present two main results. First, I show that happiness is related in sensible ways to daily local air pollution. After accounting for respondents' demographics, daily local weather

al. 2009), and even air pollution (Welsch 2007; Di Tella and MacCulloch 2008; Ferreira et al. 2006; Luechinger 2009).

² Habituation could also be relevant for hedonic estimates of compensating differentials. If owners of homes in polluted regions become habituated, those houses may have smaller measured compensating differentials.

conditions, as well as temporal and geographic fixed effects and interactions, individuals surveyed when the current local levels of airborne particulates are higher report lower levels of happiness. This first step is a straightforward empirical exercise. It requires no strong assumptions except the empirical specification, and I show that the results are robust to a variety of those. I also show that reported happiness is *not* sensitive to local levels of undetectable pollutants, such as carbon monoxide.

The second result uses the estimates from the first part to calculate marginal rates of substitution between pollution and income, and then computes respondents' implicit willingness to pay for improved air quality. This step does involve several strong assumptions, but I describe those in detail and argue they are no stronger than the assumptions underlying travel cost, hedonic, or contingent valuation estimates of willingness to pay for air quality. Moreover, because the assumptions I make differ entirely from the standard set, at a minimum, the results serve as an alternative to the usual approaches.

The analysis here yields two important lessons. For the growing literature on happiness and economics, the results provide yet another demonstration that subjective well-being varies in sensible ways with respondents' observable circumstances. For environmentalists and environmental economists, the results provide evidence that air pollution, in addition to detrimentally affecting health and property, has a direct negative effect on people's stated well-being, as well as evidence that the monetary value of that effect may be quite large. Using my preferred specification, I show that people appear willing to sacrifice about \$35 for an improvement of one standard deviation in air quality for one day, a figure about twice as large as the highest recent hedonic valuations of air quality (Bayer et al., 2009) or the U.S. Environmental Protection Agency's (EPA) assessment of the economic benefits of the 1970 and 1977 Clean Air Act Amendments (EPA 1999, 2011).

2. Happiness in economics

Happiness, as defined by respondents' answers to simple survey questions, has received a recent surge of serious attention from economists. Much of the academic and popular happiness literature addresses the decades-old findings of Easterlin (1974): stated happiness does not increase with income across countries or within a country over time, but it does increase with

income across individuals within a country at any given point in time. Some recent work challenges this Easterlin Paradox, showing that happiness increases with GDP per capita across countries in expected ways (Stevenson and Wolfers 2008; Deaton 2008; Helliwell et al. 2010). But in other recent work the paradox remains, and stated happiness appears unchanged over time even as per capita incomes have increased (Oswald 1997; Layard 2006). If true, the paradox has two obvious interpretations. One is that people become habituated to their situations and change their reference level of well-being.³ Another is that happiness depends on relative income – the richest man in a poor town may be happier than the poorest man in a rich town, even if the rich man is poorer in absolute terms.⁴

Under either interpretation, the Easterlin Paradox has implications for using happiness to measure willingness to pay for public goods. If happiness does not increase with income across regions or over time, it may also be invariant to the level of any particular public good, for similar reasons. For income, happiness does increase relative to other people in the same locale at the same time. The analog for pollution is that happiness may increase with air quality relative to the current regional norm, but not relative to other regions or within regions over long periods of time. That is why a key feature of this analysis identifies happiness as a function of the place-specific, date-specific air quality, at the place and date where the happiness question was asked. I compare stated happiness by statistically similar respondents, at the same locale, during the same season of the same year, who just happen to have been surveyed on days when the air quality differed.

While much of the economics literature on happiness focuses on deep questions about the rationality of economic actors, interpersonal comparisons of ordinal utility functions, and links between economics and psychology, economists are also attempting practical, policy-relevant applications. Recent work uses happiness surveys to evaluate people's willingness to trade unemployment for inflation and argue that central bankers place too much emphasis on combating inflation (Di Tella et al. 2001), examine the welfare consequences of German reunification on different groups (Frijters et al. 2004), assess the degree to which state cigarette taxes make

³ Kahneman (2000) writes about individuals having a base level of stated well-being, which major life events (divorce, injury) perturb at most for a few years. Others, such as Oswald and Powdthavee (2008), show incomplete recovery of happiness after such events. Graham (2009) provides evidence that people become habituated to crime, corruption, democracy, and health.

⁴ See Luttmer (2005). Also, recent work suggests this relative interpretation may be optimal from an evolutionary standpoint (Rayo and Becker 2007).

smokers better off by helping them quit (Gruber and Mullainathan 2005), and estimate the degree to which the marginal utility of consumption increases or decreases when people become ill (Finkelstein et al. 2009). Happiness measures have also been used to try to place a monetary value on airport noise (van Praag and Baarsma 2005), flood disasters (Luechinger and Raschky 2009), terrorism (Frey et al. 2009), and weather and climate (Rehdanz and Maddison 2005; Barrington-Leigh 2008).⁵ All use annual average measures of the public good (or bad), raising the possibility that endogeneity or omitted variables bias their answers.

Several papers close in spirit to this one use happiness measures to value air quality. Welsch (2002, 2006, 2007) estimates values of willingness to pay for air quality using various cross-sections and panels of country-level data. The 2006 paper, for example, estimates that the reductions in nitrogen dioxide and lead pollution in Europe from 1990 to 1997 were worth \$1,200 per capita and \$2,200 per capita, respectively, in 2008 dollars. Di Tella and MacCulloch (2008) regress happiness on income and the national, annual, per capita emissions of sulfur dioxide (SO_2), and show that an increase of one standard deviation in SO_2 correlates with a decline in happiness equivalent to a 17 percent reduction in income. As first uses of happiness data to estimate willingness to pay for air quality, these works break new ground. However, they also share a drawback common to this literature – they use average annual national measures of air quality. Aggregating environmental quality across entire countries masks much of its heterogeneity. The standard deviation of particulate air pollution in the U.S. is twice as large if we look at daily observations within states instead of averages across states or years.

One recent paper (Luechinger, 2009) avoids the problems associated with inter-country comparisons of happiness by looking across regions within Germany, using annual mean concentrations of SO_2 at 533 monitoring over a 19-year period. To control for sorting by individuals into different locales within Germany, he cleverly instruments for air quality using respondents' locations upwind and downwind of large power plants that installed SO_2 emissions control equipment. Luechinger finds a marginal willingness to pay of \$232 for a one microgram per cubic meter ($\mu\text{g}/\text{m}^3$) reduction in SO_2 , while average SO_2 concentrations fell by $38 \mu\text{g}/\text{m}^3$ over the time period.⁶

⁵ These applications raise concerns among critics. Smith (2008) writes, "[T]he [happiness economics] train is precipitously close to leaving the station and heading for use in full-scale policy evaluation."

⁶ \$232 is €183 in 2002, converted to 2008 dollars using the average 2002 exchange rate and the CPI-U-RS.

Two final issues contrast most prior attempts to value air quality using happiness data. First, work based on cross-country pollution differences faces the problem that it compares survey questions asked in various languages and cultures, where notions of happiness may differ. Second, air pollution and weather are correlated. Studies of happiness and weather omit pollution (Rehdanz and Maddison 2005; Barrington-Leigh 2008), while studies of happiness and pollution omit weather (Welsch 2007; Luechinger 2009; Di Tella and MacCulloch 2008). To date, none have included both, a potentially important source of omitted variable bias.

This paper addresses these problems. It focuses entirely on the United States, so fewer language and cultural differences complicate the responses to questions about happiness. It controls for the current local temperature and precipitation, both of which are correlated with both happiness and pollution. Instead of aggregate national or yearly measures of pollution, it uses the environmental quality at the time and in the location where the happiness survey question was asked. Fixed effects and interactions by time and place mean that the measured effect of pollution on happiness will be relative to similar respondents who were interviewed in the same place during the same month, but happen to have been interviewed on a day when the air quality differed.

3. Data and methodology

For happiness measures, I rely on the General Social Survey (GSS), which the National Opinion Research Center conducts annually.⁷ Several thousand U.S. respondents are interviewed in person each year, usually in March. The key GSS question asks, "Taken all together, how would you say things are these days? Would you say that you are very happy, pretty happy, or not too happy?" This question forms the basis for the dependent variable. In addition to asking about happiness, the GSS contains the usual demographic information, including age, household income, race, education, sex, and marital status.⁸

⁷ See www.norc.org/GSS+Website/.

⁸ The GSS income variable is categorical, and the categories change periodically, raising concerns about attenuation bias. But the survey includes numerous categories each year (21 in 1993), mitigating that concern somewhat. I use the GSS reported real income (Ligon 1989), which converts the categories into real values by taking the midpoints of the ranges and adjusting for inflation and top coding. I also attempt to control for attenuation bias by instrumenting for income using average incomes by respondents' and spouses' industries and occupations.

Importantly for this purpose, the GSS contains the date each respondent was questioned. I have obtained from the GSS staff the confidential codes identifying the county or city in which each respondent was surveyed. Knowing the date and place allows me to match the GSS to the particular air quality on the day and in the place where the survey was administered.

For pollution information, I turn to the EPA's Air Quality System (AQS). The AQS contains the raw, hourly, and daily data from thousands of ambient air quality monitors throughout the United States. The data include the latitude and longitude of each monitor, the types of pollutants monitored, and the hourly observations.⁹ For current local weather conditions, I use data from the National Climate Data Center, which reports daily temperature and rainfall at each of the thousands of weather monitoring stations throughout the United States.

To merge the survey data with the weather and air quality data, I take the population-weighted centroid of the GSS respondent's county and draw an imaginary 25-mile circle around it. I then take a weighted average of all the air quality and weather monitors within the circle, where the weights are equal to the inverse of the square root of their distance to the population-weighted centroids.¹⁰ The mean number of monitor readings used in the spatial interpolation is 3.9, ranging from 1 to 22, with a standard deviation of 3.3. Currie and Neidell (2005) confirm the accuracy of a similar weighted-distance measure by predicting pollution levels at the location of actual monitors using readings from nearby monitors. Moreover, they note the measurement error introduced by the procedure will only tend to bias the pollution effect towards zero.

The air quality monitors contain data on ambient concentrations of criteria air pollutants, but not all data are available in all places or during all time periods. Carbon monoxide (CO), for example, does have consistently measured data in many locations going back to the early 1970s. However, CO is odorless and invisible at the current ambient concentrations in these pollution data, and I would not expect it to affect happiness responses in the survey data. Airborne particulates, on the other hand, cause physical discomfort, especially particles smaller than 10 micrometers (PM10). In addition, small particles form visible haze that reduces visibility and may

⁹ Recent years are available on the AQS Web site, earlier years by special request to the EPA. More information about the AQS can be found at www.epa.gov/ttn/airs/airsaqs.

¹⁰ Other weights, such as a simple average of all the monitors in a county, yield similar results. The GSS has surveyed about 275 areas, and the names given to these areas do not typically correspond to U.S. Census or U.S. Postal Service names. The GSS geographic codes sometimes correspond to individual cities, sometimes to counties, and occasionally to multi-county areas. (This last group is dropped). I first translated the GSS place names to Census county codes by hand, then assigned each county its population centroid, and merged those with the data from the weather and pollution stations within 25 miles.

affect people aesthetically. Chay and Greenstone (2003) and Chay et al. (2003) show that particulates have adverse effects on adult and infant mortality, and Neidell and Zivin (2009) show that people avoid outdoor activities when local newspapers report poor air quality. The AQS contains PM10 readings beginning in the mid-1980s, so I begin this analysis in 1984.

For particulates, monitoring stations only record ambient concentrations every six days. As a result, many of the happiness survey questions were asked on days when no nearby air quality monitors recorded data. Moreover, in any given location, different days may be recorded by different sets of nearby monitoring stations. To smooth out this variation and use as many of the happiness survey responses as possible, I interpolate linearly between six-day observations for each monitoring station. In the robustness checks below, I also report results for the subset of observations with actual, uninterpolated values.

The GSS has 19,491 observations between 1984 and 1996, of which 10,193 had identifiable counties and could be matched to PM10 readings from the AQS. Of these, 994 were missing household incomes, another 606 could not be matched to local weather, 2498 were missing self-reported health status, which I worry may be correlated with pollution and happiness, and another 26 were missing one of the other household demographics. The resulting dataset has 6035 complete observations.

Methodology

I estimate versions of the following function:

$$H_{ijt} = \alpha P_{jt} + \gamma \ln Y_i + X'_{ijt} \beta + \delta_j + \eta_t + \delta_j \times \text{year}_t + \epsilon_{ijt}, \quad (1)$$

where H_{ijt} is the stated happiness of respondent i in location j at date t . The variable P_{jt} is the air pollution at location j at date t . The log of income ($\ln Y_i$) conveniently captures the declining marginal effect of income on happiness, consistent with typical papers estimating happiness functions, and it translates directly into an increasing marginal willingness to pay for air quality.¹¹ Below I show that the estimated trade-offs between pollution and income are unchanged if I substitute the log of pollution, the level of income, or ordered probit versions of those; or estimate a binomial probability that $H_{ijt} > H^*$ for an arbitrary H^* . This robustness to empirical specification is especially important given the limited reporting categories for the happiness

¹¹ If happiness successfully proxies for utility, we would expect diminishing marginal happiness/utility. Happiness as reported to the GSS, in three discrete categories, may not follow the same distribution. In what follows I show that the estimates are robust to a variety of functional form assumptions for equation (1).

variable in the GSS. The vector X_{ijt} contains a set of other demographic and local characteristics, δ_j is a location-specific fixed effect, η_t is a month and year fixed effect, and $\delta_j \times \text{year}_t$ captures location-specific trends.

Once estimated, I can totally differentiate the function, set $dH=0$, and solve for the average marginal rate of substitution between pollution and income, $\partial Y/\partial P$:

$$\left. \frac{\partial Y}{\partial P} \right|_{dH=0} = -Y \frac{\hat{\alpha}}{\hat{\gamma}}, \quad (2)$$

the amount of annual income necessary to compensate for a one-unit increase in air pollution on the survey date.¹² To avoid the cumbersome phrase "average marginal rate of substitution," henceforth I will use the term "willingness to pay" (WTP), fully recognizing that equation (2) represents no one person's stated willingness. Rather, it represents an estimate of the trade-offs between income and air quality that will leave people, on average, equally happy.

Some theoretical and practical concerns

Using equation (2) to measure marginal rates of substitution involves placing some strong assumptions on the underlying utility functions. We typically assume individuals make choices as though they are maximizing some unobserved utility function, observe market prices and the choices people make, and infer from those prices and choices properties of their utility functions, such as risk aversion, impatience, and altruism. The fundamental challenge facing economists valuing public goods is that we do not observe market prices or choices. Public goods such as air quality have no markets, and individuals cannot "choose" their own level of public goods directly, except by voting or relocating. So instead, this analysis proposes turning the typical economics around. We will observe utility, or a proxy for utility, and infer what choices people would be willing to make and what prices would therefore be optimal.

The first problem with this approach is that "happiness" as recorded by questions on surveys is not utility. Kahneman (2000) addresses this, distinguishing between "decision utility," which is economists' notion of the individual welfare function that drives economic choices, and "experience utility," something closer to stated happiness, experienced moment to moment. We do not observe either type of utility directly. Perhaps the easiest way to think about this

¹² Naturally, alternative formulations of (1) lead to different expressions for willingness to pay in (2). For example, using the level of income instead of its log means that, conveniently, $\partial Y/\partial P$ is simply the ratio of the coefficients on pollution and income, $\hat{\alpha}/\hat{\gamma}$.

methodology is that it uses respondents' stated happiness as a proxy for their utility, or as an observable manifestation of latent utility. As long as respondents with higher latent utility are more likely to say they are happier, this approach is consistent with a wide variety of discrete choice models in economics.

A second potential concern with the proposed approach is that the GSS happiness question is unclear about what length of time it covers, asking only how happy people are "these days." Ideally the GSS would have asked people *two* happiness questions: one about their overall life satisfaction and one about their happiness at the moment the question is asked. If "these days" refers to several months or years, the happiness response should not be influenced by temporary changes, such as the current daily level of air pollution relative to a regional seasonal norm. Psychologists and economists have found, however, that responses to life satisfaction questions differ based on short-term situations. Schwarz and Strack (1991) describe how people interviewed after making a photocopy were significantly more satisfied with their lives if they found a dime on top of the copy machine. Clark and Georgellis (2004) test whether reported "job satisfaction" proxies for "experience utility." They find that both current and lagged values of reported job satisfaction predict the likelihood British laborers will quit, suggesting that reported satisfaction has a current component. In other words, if people who are asked about their overall satisfaction with life in general respond in a way that is sensitive to current conditions, it may not matter that the GSS question has a vague time horizon.

On a related note, Loewenstein et al. (2003) develop a behavioral theory of "projection bias" wherein people misestimate their future preferences based on current circumstances – buying too much food at the grocery store if they shop while hungry. And Conlin et al. (2007) provide empirical support for projection bias, showing that people are more likely to return cold-weather gear purchased from catalogs if they made those purchase orders on colder days – overestimating their future demand for parkas based on current temperatures. Projection bias could conceivably bias hedonic estimates of WTP if people bid too much for houses on unpolluted or sunny days. One appeal of valuing air quality using happiness responses to daily pollution changes is that the valuations do not rely on people assessing their future preferences based on current circumstances, as they might when deciding where to live. Instead I measure WTP using current tradeoffs based on current circumstances, an approach closer in spirit to experience utility than decision utility.

A third likely objection to this approach is that economists normally assume utility is ordinal rather than cardinal, and that interpersonal comparisons based on stated happiness are impossible. If an unpolluted day moves person #1 from "not happy" to "very happy," and person #2 from "not happy" to "pretty happy," that does not mean that person #1 gets more utility from clean air than person #2, or that person #1 would be willing to pay more for clean air. Put differently, we could alter some people's happiness functions by a positive monotonic transformation while leaving others' unchanged, and it would yield the same rank ordering of outcomes for each individual. It would not, however, yield the same estimates of equation (1).

Economists studying happiness have responded in several ways. Some, like Ng (1997), have argued that ordinal utility is an overly restrictive assumption, and that ample evidence shows people's utilities are interpersonally comparable and cardinal. Others have implicitly assumed that happiness is ordinal but interpersonally comparable. If the latent utility of person #1 is higher than that of person #2, then the stated happiness of person #1 will also be higher. This allows researchers to estimate an ordered discrete choice model such as an ordered logit or probit. Alesina et al. (2004), Blanchflower and Oswald (2004), and Finkelstein et al. (2009) follow this empirical approach. Most researchers who have applied both approaches have found little difference between the results of a linear regression and an ordered logit or probit (Ferrer-i-Carbonell and Frijters 2004).¹³ Since I am not interested in the marginal utility of income or air quality separately, but only the ratio of the two as in equation (2), my analysis is less sensitive to these issues. I show below that the estimates of equations (1) and (2) are robust to a wide variety of empirical specifications.

Finally, economists should be concerned that income may be endogenous with respect to happiness. While more income may make people happier, inherently happier people may earn higher incomes. Very few papers address this. Luttmer (2005) instruments for household income using interactions between the respondents' and spouses' industry, occupation, and location. Powdthavee (2009) uses time series data on the number of household members working.¹⁴ Both find that the income coefficient in IV specifications is larger than in OLS specifications – three times larger in Luttmer's case. This suggests that equation (2) will overstate the marginal WTP for air quality. In contrast, Pischke (2011) finds that industry-based IV estimates of the effect of

¹³ One key advantage of the regression approach over the ordered probit is that the former can easily include fixed effects, so any individual or region-specific norms for happiness can be differenced out.

¹⁴ Gardner and Oswald (2007) circumvent the endogeneity by examining the mental wellbeing of lottery winners.

income on happiness are similar to the OLS estimates, suggesting that the causal relationship runs primarily from income to happiness and not the other way around, in which case equation (2) will not overstate the marginal WTP. To address this issue I report results from one specification where I instrument for income using a version of Luttmer's occupation and industry-based prediction of income, yielding somewhat smaller estimates of WTP.

In the end, my focus is on obtaining convincing evidence for the effect of pollution on happiness, based on local daily variation, and then using that cautiously to infer a marginal WTP. All I can do is remain cognizant of these strong assumptions, remind readers that standard approaches to valuing environmental quality – travel costs, hedonics, contingent valuation – have their own sets of strong assumptions, and demonstrate that the results obtained from this approach yield plausible valuations that are robust to different samples of the data and different empirical specifications.

4. Results

Table 1 begins by estimating versions of equation (1). The first column contains the means and standard deviations of the right-hand-side variables. Column 2 estimates equation (1) but excludes every right-hand side variable except income and daily local pollution, measured using particulates (PM10). Happiness decreases with pollution on the day of the interview and increases with annual household income. The coefficients suggest that a $10 \mu\text{g}/\text{m}^3$ increase in local daily particulates is associated with a decrease in happiness of 0.014, on a three-point scale. The log income coefficient suggests that a 10 percent increase in annual income is associated with an increase of happiness of 0.013. Since happiness may be regarded as only ordinal (or a proxy for utility which is ordinal), I do not want to overemphasize the absolute magnitudes. More important is the ratio of the two coefficients, or the trade-off between pollution and income that leaves people at the same level of happiness.

To place a dollar value on air pollution, we need to calculate equation (2). Plugging in -0.0014 for $\hat{\alpha}$, 0.133 for $\hat{\gamma}$, and 42.5 for the mean income (in \$1,000s), the WTP is $\partial Y/\partial P = \459 , as reported at the bottom of Table 1. A one $\mu\text{g}/\text{m}^3$ increase in PM10, on the day of the interview, reduces an average person's stated happiness by an amount equal to a \$459 decline in annual income. What does this mean? This \$459 figure represents an estimate of the amount of annual

income that increases happiness (at the mean log income in the sample) by the same amount as a one $\mu\text{g}/\text{m}^3$ reduction in PM10 pollution, but the PM10 coefficient is identified from daily fluctuations in air quality. If we divide the \$459 by 365 days, we get an estimate of \$1.26 per day. To put this into context, note that the standard deviation of PM10 is 14.4 $\mu\text{g}/\text{m}^3$. Our estimate, then, corresponds to a WTP of \$18 ($14.4 \times \1.26) for a one-standard-deviation improvement in air quality, for one day.¹⁵ Or, a one-standard-deviation decline in air quality makes people feel worse off by an amount equivalent to a decline in annual income resulting in having \$18 less to spend per day.

Column (3) of Table 1 adds to the regression the average particulate count for each respondent's location for the month in which the survey was taken.¹⁶ The income coefficient remains unchanged, the daily pollution coefficient increases in absolute value to -0.0017, and the monthly pollution level is insignificant and even wrong-signed. The implied WTP for a one-standard-deviation daily change would be \$21 rather than \$18. One interpretation is that the local monthly values are merely imprecise measures of the daily values, which is what people really care about. Another is that people become habituated to their environments and respond only to daily departures from the local norm.¹⁷

Column (4) of Table 1 drops the average local pollution levels, and adds instead year, month, and county fixed effects, and county-specific trends. Now the daily PM10 measure is identified from the difference between air quality on the day of the survey and the local, seasonal, trend-adjusted average air quality. None of the year or month fixed-effect coefficients, and only three of the county and year \times county coefficients, are statistically significant. The daily pollution coefficient decreases slightly, and is only marginally statistically significant, suggesting a WTP of \$15 rather than \$18. In sum, controlling for local conditions, either with a measure of local monthly air pollution or with a set of fixed effects including location-specific trends, does not change the basic findings. Local pollution on a given day appears to diminish the probability that people report high levels of happiness.

¹⁵ Note that this standard deviation of 14.4 $\mu\text{g}/\text{m}^3$ represents variation both across and within year-month-county "cells." The average standard deviation within cells is 5.7 $\mu\text{g}/\text{m}^3$. The sample includes an average of 774 observations per year, 2,298 per month, and 142 per county. The average year-month-county cell has 10 observations, ranging from 1 to 59.

¹⁶ The correlation between daily and monthly pollution levels is 0.74.

¹⁷ The standard errors on the monthly values are large, meaning we cannot differentiate between these interpretations. Monthly fixed effects, added next, also account for seasonal effects. If people are happier in spring and particulates are lower in the spring, that would bias the results absent monthly fixed effects.

Finally, column (5) adds a battery of demographic and local covariates. Happiness decreases and then increases with age, falling to a minimum at about age 40. Women and people who are married, not unemployed, and healthy are happier. All these results conform with standard findings in this literature. If anything, adding the demographic variables halves the coefficient on income, thereby doubling the estimate of WTP to \$35 for a one-standard-deviation change in PM10. This raises the possibility that other unobserved respondent characteristics may also be correlated with both income and happiness, biasing the estimated income coefficient $\hat{\gamma}$ and therefore the calculation of MWTP. On the other hand, including weather, day-of-week, and respondents' characteristics has no effect on the estimated pollution coefficient, $\hat{\alpha}$, supporting the claim that the coefficient on local daily pollution does not suffer from omitted variable bias.

The weather variables are included because pollution levels are positively correlated with temperatures and negatively correlated with rainfall, and because happiness has been shown to be affected by weather. Happiness rises with temperature at low temperatures, falls with temperature at high temperatures, and rises in the difference between the daily maximum and minimum, which proxies for clear skies and low humidity. The two temperature coefficients in column (5) imply that a 10 degree rise in temperature from 30 to 40 degrees Fahrenheit makes people happier by an amount equivalent to having an extra \$17 per day, while a rise from 80 to 90 makes people less happy by \$38. The rainfall coefficients are highly correlated with the other variables, and not statistically significant, but the point estimate implies that a rainy day makes people worse off by \$3 per day. More importantly, the additional demographic and location characteristics do not change the basic result that happiness increases with income and decreases with local daily pollution.

After including multiple fixed effects and interactions, standard household demographics, and five measures of the current local weather, the pollution coefficient remains approximately the same magnitude. The remaining pollution variation in column (5) could result from wind direction, local or upwind construction, traffic, fuel changes at factories or utilities, road paving, or other unmeasured activities. I cannot rule out that some of those might be correlated with both happiness and pollution levels, imparting an omitted variable bias to the models in Table 1. All I can do is include as many local covariates as possible, and point out that their inclusion does not dramatically change the pollution coefficient from the bare-bones specifications in column (2).

Table 2 presents a sample of some alternative specifications. First, the results so far use air quality measures that interpolate between readings that occur every six days. As an alternative, I tried using only the 40 percent of cases where uninterpolated daily readings were available for a nearby station. Those results are summarized in column (1) of Table 2. The effects of pollution and income on happiness are both slightly larger than in the basic specification, leading on balance to a nearly identical estimate of WTP for a one $\mu\text{g}/\text{m}^3$ reduction in PM10 (\$838). Because the variance across the uninterpolated values is higher than for the interpolated values ($18.2 \mu\text{g}/\text{m}^3$ rather than $14.4 \mu\text{g}/\text{m}^3$), the WTP for a one-day change of one standard deviation is slightly higher at \$42. Column (2) uses the level of income rather than its log. Nothing changes except the formula for calculating WTP. (See footnote 12). Column (3) uses both the log of income and the log of PM10, again with no meaningful change in the calculated WTP. Column (4) estimates equation (1) as an ordered probit.¹⁸ Respondents' stated happiness varies systematically with their incomes and the local daily air quality in ways that are robust to a variety of empirical specifications.

Table 3 addresses some deeper issues with the approach. Column (1) includes a control variable for the PM10 count the previous day, to account for the possibility that the effects of pollution on happiness may be cumulative. Here I limit the sample to the 25 percent of cases where uninterpolated readings were available two days in a row. The coefficient on yesterday's pollution is positive and insignificant, but its inclusion increases the negative effect of the current day's air pollution on happiness, resulting in a larger measured WTP. However, given the high degree of correlation between the two air quality measures, the point estimate of WTP over the two-day period is about the same as for the basic specification in Table 1.¹⁹

Columns (2) and (3) of Table 3 address two concerns about the measure of respondents' incomes. First, the GSS asks respondents to place their household incomes into categories representing income ranges, rather than asking them to report their actual incomes. It then takes the midpoint of each range and adjusts for inflation and top coding to report intertemporally consistent income values (Ligon (1989)). Although the survey has more than 20 income categories each year, the procedure raises the possibility of measurement error and attenuation bias, which would reduce the income coefficient and inflate the calculated WTP. A second, deeper issue

¹⁸ Estimates of equation (1) as linear probabilities and probits that $H>1$ or $H>2$ yield the similar results.

¹⁹ For the 1588 observations in column (1) of Table 3, the standard deviation of PM10 is 18.5.

involves the endogeneity of income. Happiness and household incomes are correlated, but we do not know if that is because income causes happiness, or because happy people earn higher incomes.

The solution to both problems – attenuation bias from mismeasuring income and endogeneity of incomes – is to find an instrument for household income, something that is correlated with income but with no independent effect on happiness. Powdthavee (2009) uses panel data to instrument for household incomes using changes over time in the number of household members working. His approach approximately doubles the coefficient on household income. Luttmer (2005) instruments for household incomes using the respondents' and spouses' industry, occupation, and location. Respondents who work in occupations and industries with high wages, or whose spouses do so, are likely to have higher household incomes and are therefore more likely to report higher levels of happiness. Using this instrument, Luttmer finds the coefficient on happiness is three times as large as when he uses household income directly, which suggests I should divide the estimated WTP of \$35 per day by three. On the other hand, Pischke's (2011) OLS and IV estimates are similar, suggesting little concern about income endogeneity. And Ferrer-i-Carbonell and Frijters (2004) use panel data on West German workers to show that including individual fixed effects reduces the coefficient on log-income by two thirds, suggesting the cross-section income coefficients in Table 1 here might be overstated leading to underestimates of willingness to pay.

To address both the possible mismeasurement and endogeneity of respondent's incomes, I estimate a version of Luttmer's (2005) instrumental variables approach. First, I used the Consumer Population Survey (CPS) to calculate the average annual earnings by year, state, industry, and occupation. I then matched each GSS respondent and spouse to the relevant CPS earnings. Finally, I used the respondents' and spouses' matched CPS earnings as instruments for the GSS reported household income. The underlying assumption is that industry and occupation do not predict happiness independently of the average incomes received in those occupations, and that innately happier people are not disproportionately represented in higher-paying industries or occupations.

Column (2) of Table 3 reports the first-stage prediction of log real household income from the GSS on the other right-hand side variables plus the log average wage for the respondents' and spouses' year, state, industry and occupation. The sample size shrinks due to the number of GSS

respondents with missing or mismatched industry or occupation codes. The regression fit is good, and the excluded instruments are jointly and individually statistically significant. Column (3) reports the second stage. The instrumented income coefficient (0.126) is twice as large as in the baseline specification in column (5) of Table 1, consistent with Luttmer (2005) and Powdthavee (2009), resulting in a smaller WTP for air quality. The doubling of the income coefficient would cut the point estimate of WTP in half except for the fact that the coefficient on daily pollution is also a bit higher in this smaller sample. As a result, the estimate of WTP falls to \$29 per day. Column (4) runs the baseline specification without instrumenting for income, but using this smaller sample. A slightly lower income coefficient and higher pollution coefficient lead to a larger WTP estimate of \$76 per day.

Columns (5) and (6) of Table 3 address concerns about respondents' health, reported as four categories ranging from "excellent" to "poor". If health varies daily as a consequence of pollution, then the specifications that control for health will fail to capture variations in well-being due to the health outcomes of pollution, and will only capture variations from other consequences of pollution, such as aesthetics. That is certainly not the goal here. To check whether health varies with daily pollution, in column (5) I estimate a version of equation (1) with health status as the dependent variable, rather than happiness. Here, the coefficient on daily PM10 is small and insignificant, suggesting that responses to the GSS health question are not driven by daily pollution levels, but are more likely based on long-term health characteristics having nothing to do with air quality. To be certain that the inclusion of health controls does not absorb some of the effect of pollution on happiness, in column (6) I estimate a version of equation (1) without any health controls. The pollution coefficient is unaffected, but omitting the health variables, which are negatively correlated with both income and happiness, leads to a higher income coefficient (0.090) and correspondingly a somewhat smaller estimated WTP.

Table 4 estimates the basic specification from column (5) of Table 1 for alternative measures of air quality. Column (1) estimates equation (1) for ozone. Here, the coefficient on pollution is negative but small and statistically insignificant. My initial expectation was that the ozone coefficient would be significant, since ozone is associated with aesthetically unpleasant brown skies. However, because the GSS is collected mostly in March, when ozone is not typically a problem, I may be unable to identify an ozone effect with these data.

Column (2) reports results for SO₂. This is the pollutant Luechinger (2009) studied, using annual averages for SO₂ upwind and downwind from power plants. In my case, the SO₂ coefficient is statistically insignificant, and the point estimate leads to a WTP of \$4, much less than the WTP for reductions in PM10. The different result may stem from the fact that SO₂ is less ubiquitous than PM10. SO₂ poses a particular problem downwind of coal-fired electric power plants. By focusing on respondents in the neighborhood of such plants, Luechinger was able to identify an SO₂ effect. My study covers many areas without significant SO₂ problems.

Column (3) of Table 5 reports results for carbon monoxide. Again the coefficient on CO is statistically insignificant. Unlike particulate matter, CO is odorless and colorless, and will be unnoticeable to survey respondents. Symptoms of CO exposure, including headaches, nausea, and fatigue, only appear after prolonged exposure above 70 parts per million (ppm).²⁰ In my sample, the mean and standard deviation of CO concentrations are both below 2 ppm. Currie and Neidell (2005) do find significant effects of low levels of CO on infant mortality, an important and previously overlooked result. But unnoticeable CO may still cause thousands of infant fatalities while not affecting surveyed happiness because families with ailing infants will be infrequently sampled and unlikely to respond. Any significant effect of CO on surveyed well-being would more likely be the result of its correlation with omitted covariates than the few families with affected infants.

Finally, columns (4) through (6) of Table 5 run the basic specification for PM10, but also include daily measures of Ozone, SO₂, and CO, respectively. In each case, the PM10 and income coefficients are essentially unaffected, the additional variable is statistically insignificant, and the WTP for a one-standard-deviation change in PM10 (ignoring the coefficients on the other pollutants) stays within the same range – between \$29 and \$49.

Magnitudes

So far, I have been discussing WTP for a one-standard-deviation change in pollution, which amounts to 14.4 $\mu\text{g}/\text{m}^3$ for the interpolated PM10 measurements. How large is this change? The average PM10 reading in the sample is 30.4 $\mu\text{g}/\text{m}^3$, so one standard deviation constitutes a 47 percent change in pollution. For comparison, the EPA (1999) publication *Benefits and Costs of*

²⁰ See, for example, the US Consumer Product Safety Commission's FAQ sheet (www.cpsc.gov/cpscpub/pubs/466.html).

the Clean Air Act estimates that the 1970 and 1977 Clean Air Act Amendments reduced ambient particulate matter by an average of 45 percent nationally. Though comparable in magnitude, those Clean Air Act improvements represent long-run changes, whereas the WTP calculations here are identified from short-term fluctuations. Empirical work to date suggests those can have quite different outcomes. Infant mortality, for example, has been shown to be associated with long-term changes in particulates (Chay and Greenstone, 2003) but not with short-term changes (Currie and Neidell, 2005). Still, for context it is worth comparing the valuations I get using happiness data and short-term pollution changes to existing valuations using other methodologies.

Start with the EPA's valuation of the 45 percent reduction in particulates they attribute to the 1970 and 1977 Clean Air Acts. The EPA estimates that those air quality improvements reduced premature mortality, chronic bronchitis, days with respiratory symptoms, and lost work days, each of which they assigned a monetary value based on the existing economics literature valuing health costs and statistical lives. Focusing solely on the reduction in particulates, the estimated total benefit is slightly more than 1.6 trillion 2008 dollars, or \$6,880 per capita, or \$19 per day per person.²¹ By comparison, the value of \$35 per day in Table 1 does not seem out of the question. On one hand, the Table 1 estimates omit any effects of air quality that are only noticeable over long periods. But they include many effects omitted from the EPA study, such as aesthetic values, ecological effects, non-monetized short-term health effects, altruism, and any immediately observable consequences of multiple pollutants correlated with PM10. And, because this approach only examines short term changes it does not include any dampening effects of habituation on willingness to pay.

In 2011 the EPA released its second comprehensive study of the benefits and costs of the Clean Air Act (EPA, 2011). That study estimates that the 1990 Clean Air Act Amendments reduced population-weighted average exposure to particulates smaller than 2.5 micrometers – a slightly different measure of air quality than the PM10 used in this analysis – from $17.7 \mu\text{g}/\text{m}^3$ to 10.9 in 2010. And the study estimates that this improvement prevents 160,000 annual premature fatalities that would have been caused by higher levels of air pollution. Using the EPA standard monetization of the value of a statistical life (VSL) of \$7.9 million in 2008 dollars, this amounts to \$1.2 trillion annually, or \$11 per person per day. Although this estimate involves a different

²¹ Calculations based on Tables ES-1 and ES-3 in EPA (1999), adjusted for inflation using the CPI-U-RS, and a 1990 48-state U.S. population of 247 million.

measure of particulates, the valuation is not wildly different from the ones using this happiness approach.

An alternative to using health and mortality would be the hedonic method, regressing house prices on housing characteristics including air quality. Smith and Huang (1995) conduct a meta-analysis of this literature and find an average marginal WTP for a one $\mu\text{g}/\text{m}^3$ reduction in total suspended particulates of \$226 (in 2008 dollars). A $14.4 \mu\text{g}/\text{m}^3$ increase would be worth \$3,254, which amortized at 5 percent is worth \$163 per year, or considerably less than \$1 per day. Chay and Greenstone (2005) use an instrumental variables approach to compare housing values in U.S. counties according to whether they are in compliance with National Ambient Air Quality Standards and find that housing values in non-compliance counties grew by an average of \$2,774 between 1970 and 1980 (in 2008 dollars) due to the Clean Air Act. Amortized at 5 percent, this amounts to \$137 per year, comparable to the Smith and Huang numbers. More recent work by Bahari et al. (2012) uses repeated sales of the same houses to adjust for time-varying unobserved attributes and finds a WTP for a one $\mu\text{g}/\text{m}^3$ improvement of \$94 to \$104, which would be about \$4 per day for a $14.4 \mu\text{g}/\text{m}^3$ improvement. Bayer et al. (2009) use householders' birth cities to control for aversion to moving and find WTP of \$320 to \$397 per $\mu\text{g}/\text{m}^3$, or \$13-\$16 per day for $14.4 \mu\text{g}/\text{m}^3$, closer to the values here.

Probably the most controversial methodology for valuing environmental quality is contingent valuation, which asks respondents directly to place monetary values on environmental changes. A seminal example of this approach is an EPA-sponsored evaluation of air quality in California (Loehman et al. 1985). They asked respondents whether they would vote to improve air quality by 30 percent, along with associated health and visibility, at various costs, and showed them photographs of the sky with clean and dirty air. While not directly comparable to the $14.4 \mu\text{g}/\text{m}^3$ improvements discussed above, the average annual WTP was \$980 in Los Angeles and \$251 in San Francisco (in 2008 dollars), again considerably less than the EPA's values or those in Table 1.

The estimates of willingness to pay for improvements in air quality derived in Table 1 may be slightly overstated if the coefficient on income is underestimated due to attenuation bias from mismeasuring income, endogeneity of income, or omitted variable bias, but the general magnitudes are not out of line with the EPA's valuations of the particulate reductions attributable to the Clean Air Act. The estimates in Table 1 are, however, larger than those from most hedonic

regressions of property values on air quality and other housing characteristics, and from contingent valuation surveys of people's directly stated willingness to pay. One possible explanation is that existing methods measure willingness to pay for long-run air quality differences rather than daily fluctuations. To the extent people become habituated to systematic differences across jurisdictions, we would expect the happiness approach using daily fluctuations to generate higher valuations.

Nonlinearities: Interactions with other demographics

One natural test of whether these results truly measure reactions to air pollution, and not some spurious covariate, is to check whether they vary sensibly with respondents' characteristics. A natural candidate is income. If environmental quality is a normal good, we would expect WTP to increase with income. To test this directly, I include an interaction between the income variable and the daily PM10 count. To ensure that the coefficient α_1 can be interpreted in the same way as previously, at the average income, I interact pollution with the difference between the respondent's log income and the mean log income in the sample. Bars above variables denote means.

$$H_{ijt} = \alpha_1 P_{jt} + \gamma \ln Y_i + \alpha_2 P_{jt} (\ln Y_i - \bar{\ln Y}) + X'_{ijt} \beta + \delta_j + \eta_t + \delta_j \times \text{year}_t + \epsilon_{ijt} \quad (3)$$

Results are reported in the first column of Table 5. The pollution coefficient is unchanged by the inclusion of the interaction, and although the interaction term's coefficient ($\hat{\alpha}_2$) is not statistically significant, the two terms together ($\hat{\alpha}_1$ and $\hat{\alpha}_2$) are jointly significant, and the interaction coefficient is negative, suggesting that higher-income individuals are willing to pay more for clean air.

The marginal rate of substitution between income and air quality in this case, for the average level of pollution and log income, is

$$\left. \frac{\partial Y}{\partial P} \right|_{dH \equiv 0} = -Y \frac{[\hat{\alpha}_1 + \hat{\alpha}_2 (\ln Y - \bar{\ln Y})]}{\hat{\gamma} + \hat{\alpha}_2 P} \quad (4)$$

As shown at the bottom of Table 5, the point estimates in column (1) are such that people in the 25th percentile of the GSS income distribution appear willing to pay \$28 for a change of one standard deviation in air quality, and people in the 75th percentile appear to be willing to pay \$45.

Another variable we might expect to be correlated with WTP for daily air quality is the local average air quality. This could go in one of two directions. People could become habituated to poor air quality, and a one $\mu\text{g}/\text{m}^3$ change could affect people less in polluted areas than in clean areas. Or, if marginal disutility from pollution increases, we could find the opposite. In column (2) of Table 5, I estimate a version of

$$H_{ijt} = \alpha_1 P_{jt} + \gamma \ln Y_i + \alpha_2 P_{jt} I_{ijt} + X'_{ijt} \beta + \delta_j + \eta_t + \delta_j \times \text{year}_t + \epsilon_{ijt} \quad (5)$$

where I_{ijt} represents the interacted variable, in this case local monthly pollution. The interaction is statistically insignificant, but the interaction and the pollution variables together are jointly significant. The point estimate of the interaction is positive, suggesting if anything, pollution affects happiness less in polluted areas. The marginal rate of substitution can be calculated as

$$\frac{\partial Y}{\partial P} \Big|_{dH \equiv 0} = -Y \frac{[\hat{\alpha}_1 + \hat{\alpha}_2 I]}{\hat{\gamma}} \quad (6)$$

where I is the interacted variable. WTP appears to fall from \$52 at the 25th percentile of the PM10 distribution to \$38 at the 75th percentile, suggesting that habituation may overcome rising marginal damages.

PM10 is especially harmful for people with asthma or other respiratory problems. The GSS does not have data on respiratory problems *per se* but does have self-reported health status. In column (3) of Table 5, I include an interaction between the PM10 count and the indicator for whether a respondent's health status is fair or worse. The interaction term is statistically insignificant and positive, suggesting that people in poor health are not made even worse during high PM10 days than people in good health. This may reflect the crude nature of the health variable. For example, it could be that people in excellent health are more likely to exercise outdoors and therefore be affected by PM10 than people in poor health who remain indoors regardless of pollution levels. The bottom of column (3) reports the point estimates of WTP for people in better and worse health, \$41 and \$17, respectively.

Finally, in column (4) I interact the PM10 count with the indicator for whether the respondent has a college degree. That interaction coefficient is also statistically insignificant, though again it is jointly significant with daily PM10. The point estimates suggest college graduates are willing to pay \$20 more per day than those without college degrees for improvements in air quality. In sum, the general pattern of the interaction terms reported in Table

5 do not irrefutably demonstrate the merit of this happiness approach to valuing public goods, nor do they undermine it. Although the interaction coefficients do make intuitive sense, such as the fact that higher-income, more educated, respondents value clean air more than others, those differences are not statistically significant.

5. Conclusions: Advantages and disadvantages of the happiness approach

Economists estimate the benefits of public goods using several approaches. Each has associated advantages and disadvantages. Travel cost models face difficulty valuing time spent en route and on site. Contingent valuation methods are vulnerable to biases due to framing of the question, the monetary starting points used, strategic responses, and the critique that if respondents do not know about an environmental problem until it is described by the surveyor, the very fact of conducting the survey creates the WTP. Hedonic approaches suffer from Tiebout sorting and omitted variable bias. And using healthcare costs alone to value environmental quality understates the amount people would be willing to pay to avoid being sick in the first place.

The "happiness" approach to valuing public goods has its own set of weaknesses. It makes stronger assumptions about preferences than economists typically make, in that it compares the stated happiness of different individuals. It translates changes in stated happiness in response to temporary changes in pollution into systematic WTP, while at the same time, stated happiness does not seem responsive to systematic differences in pollution. And it treats household income as exogenous. Nevertheless, this new approach has a number of notable advantages.

First, the drawbacks of this approach are different from the drawbacks of the typically used approaches. It is more direct than hedonic or travel cost models, in that it relies on surveys of people's well-being, yet it is not as direct as the contingent valuation approach, in that it does not ask about environmental quality *per se*, avoiding any strategic response bias. As a result, this new approach, if nothing else, serves as a complement to existing approaches. Second, the happiness approach comes from nationally representative surveys and so can be used to assess how WTP varies over time and by income, health, education, and the current level of pollution. Finally, economists are increasingly interested in using happiness to measure the value of public goods and bads, such as unemployment and inflation, terrorism, airport noise, inequality, and flood control. These all face the obstacle that such public goods do not vary across individuals in the

same location during the same year. It seems only natural, therefore, to use this happiness approach to evaluate the economic benefits of the environment, and to take advantage of the fact that air quality changes daily in any given location.

What have we learned? This exercise is unlikely to be generally useful as an everyday cost-benefit tool, if only because its data demands are so extensive. It has been feasible in this one special case – a well-monitored, easily observable air pollutant that varies daily. We are not going to be able to use this approach to assess the value of environmental externalities that are imperceptible, such as carcinogens, or that do not vary on a daily basis, such as clean water or accident risk. The exercise has, however, demonstrated several important points. First, the results add to the evidence that self-reported subjective well-being captures something meaningful about people's circumstances – in this case, the quality of their daily local environments. Second, the results demonstrate that pollution has a direct effect on people's welfare, at least as self-reported well-being, in addition to any measured effects through health, lost work days, and other observable outcomes. Finally, the results demonstrate evidence of a substantial trade-off between income and environmental quality – a compensating differential for pollution.

References

Alesina, Alberto, Rafael Di Tella, and Robert MacCulloch. 2004. "Inequality and Happiness: Are Europeans and Americans Different?" *Journal of Public Economics*, 88(9–10): 2009–2042.

Bahari Patrick, Jane Cooley Fruehwirth, Kyoo Il Kim, and Christopher Timmins. 2012. "A Rational Expectations Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution," *American Economic Review* forthcoming.

Bayer, Patrick, Nat Keohane and Christopher Timmins. 2009. "Migration and Hedonic Valuation: The Case of Air Quality," *Journal of Environmental Economics and Management*, 58(1): 1-14.

Barrington-Leigh, Christopher P. 2008. "Weather as a Transient Influence on Survey-reported Satisfaction with Life." University of British Columbia Working Paper.

Blanchflower, David and Andrew Oswald. 2004. "Well being over time in Britain and the USA" *Journal of Public Economics*, 88(7-8): 1359-1386.

Chay, Kenneth, Carlos Dobkin and Michael Greenstone. 2003. "The Clean Air Act of 1970 and Adult Mortality." *Journal of Risk and Uncertainty*, 27(3).

Chay, Kenneth, and Michael Greenstone. 2003. "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession." *Quarterly Journal of Economics*, 118(3).

Chay, Kenneth, and Michael Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy*, 115(2): 376–424.

Clark, Andrew E., and Yannis Georgellis. 2004. "Kahneman Meets the Quitters: Peak-End Behavior in the Labour Market." DELTA Mimeo.

Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang. 2007. "Projection Bias in Catalog Orders" *American Economic Review* 97(4): 1217-1249.

Currie, Janet and Matthew Neidell. 2005. "Air Pollution and Infant Health: What Can We Learn from California's Recent Experience?" *Quarterly Journal of Economics*, 120(3): 1003-1030.

Deaton, Angus. 2008. "Income, Aging, Health and Wellbeing Around the World: Evidence from the Gallup World Poll." *Journal of Economic Perspectives*, 22(2): 53–72.

Di Tella, Rafael, Robert MacCulloch, and Andrew Oswald. 2001. "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness." *American Economic Review*, 91(1): 335–341.

Di Tella, Rafael, and Robert MacCulloch. 2008. "Gross National Happiness as an Answer to the Easterlin Paradox?" *Journal of Development Economics*, 86(1): 22–42.

Easterlin, R. 1974. "Does Economic Growth Improve the Human Lot? Some Empirical Evidence." In *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz*, ed. P.A. David and M.W. Reder, 89–125. New York: Academic Press.

EPA (Environmental Protection Agency). 1999. *Benefits and Costs of the Clean Air Act: Final Report to Congress on Benefits and Costs of the Clean Air Act, 1990 to 2010*. EPA 410-R-99-001. Washington, DC: EPA.

EPA (Environmental Protection Agency). 2011. *The Benefits and Costs of the Clean Air Act From 1990 to 2020*. Washington, DC: EPA.

Ferrer-i-Carbonell, Ada, and Paul Frijters. 2004. "How Important is Methodology for the Estimates of the Determinants of Happiness?" *The Economic Journal*, 114(497): 641–659.

Finkelstein, Amy, Erzo F.P. Luttmer, and Matthew J. Notowidigdo. 2009. "Approaches to Estimating the Health State Dependence of the Utility Function" *American Economic Review Papers and Proceedings*, 99(2): 116-121.

Frijters, Paul, John Haisken-DeNew, and Michael Schiels. 2004. "Money Does Matter! Evidence from Increasing Real Income and Life Satisfaction in East Germany Following Reunification." *American Economic Review*, 94(3): 730–752.

Frey, Bruno, Simon Luechinger, and Alois Stutzer. 2009. "The Life Satisfaction Approach to Valuing Public Goods: The Case of Terrorism." *Public Choice*, 138(3–4): 317–345.

Gardner, Jonathan, and Andrew J. Oswald. 2007. "Money and Mental Wellbeing: A Longitudinal Study of Medium-sized Lottery Wins." *Journal of Health Economics*, 26(1): 49-60.

Graham, Carol. 2009. "Why Societies Stay Stuck in Bad Equilibrium: Insights from Happiness Studies amidst Prosperity and Adversity." Paper presented to the IZA Conference on Frontiers in Labor Economics: The Economics of Well-Being and Happiness, Washington, DC.

Gruber, Jonathan, and Sendhil Mullainathan. 2005. "Do Cigarette Taxes Make Smokers Happier?" *Advances in Economic Analysis and Policy*, 5(1): Article 4.

Helliwell, John F., Christopher P. Barrington-Leigh, Anthony Harris, and Haifang Huang. 2010. "International Evidence on the Social Context of Well-being." In *International Differences in Well-Being*, ed. E. Diener, D. Kahnmena, J. Helliwell. Oxford University Press.

Kahneman, Daniel. 2000. "Experienced Utility and Objective Happiness: A Moment-Based Approach." In *Choices, Values and Frames*, ed. D. Kahneman and A. Tversky, 673–692. New York: Cambridge University Press.

Layard, Richard. 2006. "Happiness and Public Policy: A Challenge to the Profession." *The Economic Journal*, 116(510): c24–c33.

Ligon, Ethan. 1989. *The Development and Use of a Consistent Income Measure for the General Social Survey*. GSS Methodological Report No. 64. Chicago: NORC.

Loehman, Edna, David Boldt, and Kathleen Chaikin. 1985. *Measuring the Benefits of Air Quality Improvements in the San Francisco Bay Area: Property Value and Contingent Valuation Studies, Volume IV of Methods Development for Environmental Control Benefits Assessment*. EPA-230-12-85-022. Washington, DC: EPA.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility" 118(4): *Quarterly Journal of Economics* 1209-1248.

Luechinger, Simon. 2009. "Valuing Air Quality Using the Life Satisfaction Approach." *Economic Journal*, 119(536): 482–515.

Luechinger, Simon, and Paul A. Raschky. 2009. "Valuing Flood Disasters Using the Life Satisfaction Approach." *Journal of Public Economics*, 93(3–4): 620–633.

Luttmer, Erzo F.P. 2005. "Neighbors as Negatives: Relative Earnings and Well-being." *Quarterly Journal of Economics*, 120(3): 963–1002.

Neidell, Matthew, and Joshua Graff Zivin. 2009. "Days of Haze: Environmental Information Disclosure and Intertemporal Avoidance Behavior." *Journal of Environmental Economics and Management*, 58(2): 119-128.

Ng, Yew-Kwang. 1997. "A Case for Happiness, Cardinalism, and Interpersonal Comparability." *The Economic Journal*, 107(445): 1848–1858.

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

Oswald, Andrew, and Nattavudh Powdthavee. 2008. "Does Happiness Adapt? A Longitudinal Study of Disability with Implications for Economists and Judges." *Journal of Public Economics*, 92(5–6): 1061–77.

Pischke Jorn-Steffen. 2011. "Money and Happiness: Evidence from the Industry Wage Structure" NBER Working Paper 17056.

Powdthavee, Nattavudh. 2009. "How Much Does Money Really Matter? Estimating the Causal Effects of Income on Happiness." *Empirical Economics*, 39(1): 77–92.

Rayo, Luis, and Gary Becker. 2007. "Evolutionary Efficiency and Happiness." *Journal of Political Economy*, 115(2): 302–337.

Rehdanz, Katrin, and David Maddison. 2005. "Climate and Happiness." *Ecological Economics*, 52(1): 111–125.

Rehdanz, Katrin, and David Maddison. 2008. "Local Environmental Quality and Life-Satisfaction in Germany." *Ecological Economics*, 64(4): 787–797.

Schwarz, Norbert, and Fritz Strack. 1991. "Evaluating One's Life: a Judgment Model of Subjective Well-Being." In *Subjective Well-Being: An Interdisciplinary Perspective*, ed. Fritz Strack, Michael Argyle, and Norbert Schwarz, 27–47. Oxford, UK: Pergamon Press.

Smith, V. Kerry. 2008. "Reflections on the Literature." *Review of Environmental Economics and Policy*, 2(2): 292–308.

Smith, V. Kerry, and Ju-Chin Huang. 1995. "Can Markets Value Air Quality? A Meta-Analysis of Hedonic Property Value Models." *Journal of Political Economy* 103(1): 209–227.

Stevenson, Betsey, and Justin Wolfers. 2008. "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox." *Brookings Papers on Economic Activity*, Spring: 1–87.

van Praag, Bernard M.S., and Barbara E. Baarsma. 2005. "Using Happiness Surveys to Value Intangibles: The Case of Airport Noise." *The Economic Journal*, 115(500): 224–246.

Welsch, Heinz. 2002. "Preferences over Prosperity and Pollution: Environmental Valuation based on Happiness Surveys." *Kyklos*, 55: 473–494.

Welsch, Heinz. 2006. "Environment and Happiness: Valuation of Air Pollution Using Life Satisfaction Data." *Ecological Economics*, 58: 801–813.

Welsch, Heinz. 2007. "Environmental Welfare Analysis: A Life Satisfaction Approach." *Ecological Economics* 62: 544–551.

Table 1: Happiness, Pollution, and Income: Linear Regressions and Particulates (PM10) Coefficients

Means	Pollution and income only	Add average pollution	Add time and county f.e.'s	Baseline specification	
	(1)	(2)	(3)	(4)	(5)
PM10 daily ($\mu\text{g}/\text{m}^3$) [a]	30.4 (14.4)	-0.00143* (0.00061)	-0.00169* (0.00062)	-0.00122 [†] (0.00064)	-0.00136* (0.00065)
log(real income (\$1,000 2008)) [y]	3.75 (0.97)	0.133* (0.012)	0.133* (0.012)	0.135* (0.008)	0.065* (0.010)
Average PM10 by county and year			0.00183 (0.00121)		
Age (÷10)	4.4 (1.7)				-0.112* (0.030)
Age (÷10) squared	22.0 (16.4)				0.014* (0.003)
Female	0.56				0.042* (0.016)
Married	0.51				0.251* (0.018)
Kids	0.70				-0.111* (0.020)
Employed	0.66				-0.031 (0.020)
Unemployed	0.023				-0.190* (0.054)
College graduate	0.24				0.036 [†] (0.019)
Health fair or worse	0.20				-0.250* (0.022)
Health poor	0.044				-0.203* (0.042)
Rain (indicator)	0.45				-0.0035 (0.0190)
Rain (0.01 inches)	9.48 (24.23)				0.0002 (0.0004)
Temperature mean (10° F)	4.37 (1.42)				0.064* (0.029)
Temperature squared	21.1 (12.4)				-0.0063 [†] (0.0034)
Temp. diff. (daily max–min)	2.00 (0.78)				0.0092 (0.0127)
Constant		1.72* (0.04)	1.67* (0.05)	-3.02 (7.21)	--
Year, month, county f.e.'s, county-trends		no	no	yes	yes
Day-of-week and holiday fixed effects		no	no	no	yes
R ²	0.044	0.044	0.054	0.129	
No. of obs. = 6,035					
Years: 1984–1996, skipping 1992, 1995					
WTP to pay for a one $\mu\text{g}/\text{m}^3$ reduction	\$459* (188)	\$541 (194)	\$386 [†] (205)	\$891* (446)	
WTP to pay for a one std. dev. reduction for one day	\$18	\$21	\$15	\$35	

* Statistically significant at 5 percent. [†] Statistically significant at 10 percent. Std. deviations in column (1). Standard errors in columns (2)–(5) adjusted for clustering by county. Standard errors of WTP use the delta method. The dependent variable "happiness" has mean 2.17, std. dev. 0.63.

Table 2: Happiness, Pollution, and Income: Alternative Functional Forms and PM10

	PM10 without interpolation	Linear in income	In(Income) In(PM10)	Ordered probit: In(income)
	(1)	(2)	(3)	(4)
PM10 daily ($\mu\text{g}/\text{m}^3$) [α]	-0.0017* (0.0008)	-0.0014* (0.0006)	-0.044* (0.021)	-0.0027* (0.0012)
Income [γ]	0.082* (0.015)	0.0013* (0.0002)	0.065* (0.010)	0.130* (0.020)
Other covariates and fixed effects as in column (5) of Table 1		yes	yes	yes
R^2	0.154	0.130	0.129	
No. obs.	2,567	6,035	6,035	6,035
Years: 1984–1996, skipping 1992, 1995				
WTP to pay for a one $\mu\text{g}/\text{m}^3$ reduction [$-\alpha/\gamma$]	\$838 [†] (443)	\$1,075* (516)	\$947* (483)	\$890* (441)
WTP to pay for a one std. dev. reduction for one day	\$42	\$42	\$37	\$35

* See the footnote to Table 1.

Table 3: Alternative Approaches

	Instrument for income using average income by state, occupation and industry					
	Lagged environment	First stage: dependent variable = log(income)	Second stage: dependent variable = happiness	Baseline specification with smaller sample	Health as dependent variable	Main specification without health
	(1)	(2)	(3)	(4)	(5)	(6)
PM10 daily ($\mu\text{g}/\text{m}^3$) [a]	-0.0018 [†] (0.0011)	-0.0014 (0.0013)	-0.0021* (0.0010)	-0.0022* (0.0011)	-0.0005 (0.0008)	-0.0014* (0.0007)
PM10 previous day	0.0009 (0.0011)				--	--
log(real income (\$1,000 2008)) [y]	0.071* (0.019)		0.126 [†] (0.072)	0.050* (0.016)	0.153* (0.013)	0.090* (0.010)
log(CPS real income by year, state, occupation, industry)		0.301* (0.028)				
log(CPS real income for spouse's occupation, industry)		0.020 [†] (0.012)				
Other covariates and fixed effects as in column (5) of Table 1.	yes	yes	yes	yes	yes	yes
R ²	0.146	0.43	0.152	0.120	0.164	0.096
No. of obs.	1,588	2,599	2,599	2,599	6,035	6,035
Years: 1984–1996, skipping 1992, 1995						
F(2,2441) test excluded insts.		59.0				
Sargan overid test p-value			0.43			
WTP to pay for a one $\mu\text{g}/\text{m}^3$ reduction	\$1,057 (697)		\$728 (566)	\$1,922 [†] (1,126)	--	\$679* (322)
WTP to pay for a one std. dev. reduction for one day	\$54		\$29	\$76	--	\$27

* See the footnote to Table 1. Column (1) includes only observations where pollution was monitored in a county on successive days. Column (1) also includes lagged

temperature and rainfall. In column (4), health is coded from 1 ("poor") to 4 ("excellent").

Table 4: Other Pollutants

Dependent variable: Happiness (1–3)	Ozone	Sulfur Dioxide (SO₂)	Carbon Monoxide (CO)	PM10 and Ozone	PM10 and SO₂	PM10 and CO
	(1)	(2)	(3)	(4)	(5)	(6)
Pollution (daily) [α]	-0.0003 (0.0008)	-0.00021 (0.00077)	-0.0079 (0.0071)	-0.0015 [†] (0.0009)	-0.0019* (0.0007)	-0.0011 (0.0007)
log(real income (\$1,000 2008)) [γ]	0.067* (0.009)	0.071* (0.008)	0.067* (0.008)	0.056* (0.013)	0.064* (0.011)	0.063* (0.010)
Second pollutant				0.0019 (0.0013)	0.0025 (0.0019)	-0.0195 (0.0153)
R ²	0.134	0.133	0.130	0.134	0.135	0.133
No. of obs.	8,140	9,860	10,081	3,855	4,916	5,439
Years	1975–1996	1975–1996	1975–1996	1984–1996	1984–1996	1985–1996
WTP to pay for a one µg/m ³ reduction ^a	\$143 (457)	\$126 (469)	\$4,950 (4,529)	\$1,180 (715)	\$1,302* (565)	\$751 (488)
WTP to pay for a one std. dev. reduction for one day ^a	\$8	\$4	\$18	\$46	\$49	\$29

* See the footnote to Table 1. All regressions contain the other covariates and fixed effects as in column (5) of Table 1.

^aThe WTP calculations in columns (4)–(6) are based only on the coefficients on PM10 and income (α and γ), not the second pollutant.

Table 5: Interactions

Dependent variable: Happiness (1-3)		Income	Local	Health fair	College
			monthly pollution		
PM10	$[\alpha_1]$	-0.0014*	-0.0019*	-0.0016*	0.0012 ^t
		(0.0006)	(0.0008)	(0.0007)	(0.0007)
Income (\$1000 1986)	$[\gamma]$	0.084*	0.065*	0.065*	0.035*
		(0.020)	(0.010)	(0.010)	(0.010)
Interaction	$[\alpha_2]$	-0.00061	0.00005	0.0009	-0.0007
		(0.00058)	(0.00006)	(0.0013)	(0.0013)
Interacted variable			-0.0003	-0.279*	0.058
			(0.0030)	(0.045)	(0.044)
N		6035	6035	6035	6035
R ²		0.129	0.129	0.129	0.129
F test that pollution and interaction = zero		2.79*	3.00*	2.49*	2.39
WTP to pay for a one std. dev. reduction for one day when interaction = 25th percentile		\$28	\$52		
WTP to pay for a one std. dev. reduction for one day when interaction = 75th percentile		\$45	\$38		
WTP to pay for a one std. dev. reduction for one day when interaction = 0				\$41	\$31
WTP to pay for a one std. dev. reduction for one day when interaction = 1				\$17	\$51

* See the footnote to Table 1. All regressions contain the other demographic and local variables, location, year and month fixed effects, as in column (5) of Table 1.

