Unhappy Cities

Edward L. Glaeser, *Harvard University and National Bureau of Economic Research*

Joshua D. Gottlieb, *University of British Columbia and National Bureau of Economic Research*

Oren Ziv, *Michigan State University and Dartmouth College*

There are persistent differences in self-reported subjective well-being across US metropolitan areas, and residents of declining cities appear less happy than others. Yet some people continue to move to these areas, and newer residents appear to be as unhappy as longer-term residents. While historical data on happiness are limited, the available facts suggest that cities that are now declining were also unhappy in their more prosperous past. These facts support the view that individuals do not maximize happiness alone but include it in the utility function along with other arguments. People may trade off happiness against other competing objectives.

According to the Behavioral Risk Factor Surveillance System (BRFSS), only 35.9% of the residents of the Gary, Indiana, metropolitan area report themselves as very satisfied with their lives, as opposed to 45.7% across the United States as a whole. Self-reported unhappiness is high in other

We are grateful to David Albouy, Ori Heffetz, Naomi Hausman, Jordi Jofre-Monseny, Jacob Leshno, Carolin Pflueger, Kelly Shue, Doug Staiger, Dave Zhao, and especially Erzo F. P. Luttmer for valuable comments and assistance. We thank Kristina Tobio and Ian Zhu for outstanding research assistance. We acknowledge support from the Taubman Center for State and Local Government. Gottlieb acknowledges support from the Social Sciences and Humanities Research Council of Canada, the Institute for Humane Studies, and the National Institutes on

[*Journal of Labor Economics,* 2016, vol. 34, no. 2, pt. 2]
 0734-306X/2016/34S2-0016$10.00
Submitted April 8, 2014; Accepted October 17, 2014; Electronically published February 11, 2016

declining cities, and this tendency persists even when we control for income, race, and other personal characteristics. Why are the residents of some cities persistently less happy? Given that they are, why do people choose to live in unhappy places?

The presence of significant differences in self-reported well-being across places within the United States poses something of a challenge for the reigning paradigm of urban economics—the concept of a spatial equilibrium. This central idea—proposed by Alonso (1964), Muth (1969), Rosen (1979), and Roback (1982)—assumes that wages and prices adjust so that in equilibrium there are no arbitrage opportunities across space. In equilibrium, individuals cannot improve their overall utility levels by migrating within the United States.

There are two ways to reconcile differences in self-reported well-being with the notion of a spatial equilibrium. First, subjective well-being (SWB) may not be equivalent to the economist's concept of utility. Under this view, agents make decisions in order to jointly maximize expected future happiness and other objectives. Compensating differences in other dimensions offset persistent spatial differences in happiness. Second, the observed differences in subjective well-being may not reflect the permanent life-long well-being for otherwise identical people. The unhappiness might be transitory or explained by unobserved individual heterogeneity, especially if some areas attract people who are disproportionately prone to be more or less happy.

In Section I of this paper, we follow Oswald and Wu (2011) and use the BRFSS to measure subjective well-being across the United States. We extend their work by calculating SWB at finer geographic levels, adjusting for observable individual differences, and correcting for sampling error. We find significant, but not huge, differences across metropolitan areas both with and without controlling for state fixed effects. After correcting for sampling noise, we find that the cross-city standard deviation of happiness is about 6% of a standard deviation of individual happiness. This is approximately the difference in subjective well-being between the sexes or that between high school graduates and those with some college. This difference is roughly the order of magnitude caused by a one standard deviation decline in neighborhood poverty (Ludwig et al. 2012). We also find that this variation persists when we control for a rich battery of individual controls, including employment status and income.

Aging via grant number T32-AG000186-23 to the National Bureau of Economic Research. Ziv acknowledges support from the NE-UTC and the NSF via grant nos. DGE0644491 and DGE1144152. Contact the corresponding author, Edward Glaeser, at eglaeser@harvard.edu. Information concerning access to the data used in this article is available as supplementary material online.

One primary concern is whether these differences are caused by unobserved heterogeneity, either in human capital or in propensity toward happiness. We address this using the National Survey of Families and Households (NSFH). This is a panel survey, with which we can estimate area-level happiness by looking at individuals who move across metropolitan areas between the survey's first wave (1987–88) and second wave (1992–94). Differences in happiness persist, even when we control for individual fixed effects. The correlation between area-level estimates with and without individual fixed effects is 0.69. This leads us to believe that much of the difference in happiness across space reflects more than the selection of unhappy people into unhappy places.

We next document that area-level happiness is essentially uncorrelated with many area attributes. For example, metropolitan area population and housing values are orthogonal to subjective well-being in the BRFSS. Like Florida, Mellander, and Rentfrow (2013), we find that area-level education is positively associated with subjective well-being, but we find that this effect vanishes when we control for individual-level education. If more educated individuals only became educated because of the education level of the area, then it can be fairly said that these places have made them happier. But if they would have been educated regardless of place, then the happiness of more educated areas should be interpreted as differential selection.

In Section II, we document the one robust fact that emerges clearly from multiple data sets: places with lower levels of population and income growth are less happy (Glaeser and Redlick 2009). Lucas (2013) also finds higher rates of migration to counties with higher subjective well-being in the BRFSS, arguing that the migration patterns are consistent with a spatial equilibrium with happiness as a measure of utility. We find that the relationship persists for quite long periods (from 1950 to 2000). Moreover, we find the strongest effect at the left tail of SWB. It is not that high-growth places are particularly happy but rather that very low-growth areas are particularly unhappy. It is possible that people flee areas that produce unhappiness, but the long time periods involved make it hard to believe that these differences are transitory.

We show that the connection between low well-being and decline persists when we control for a bevy of individual controls, including education and income, and even when we control for state fixed effects. This fact appears in three independent surveys. In the NSFH, the effect does not persist in the general individual fixed effects estimation, but it reemerges when we limit our sample to cities with more than 250 respondents across both waves. None of these results speak to whether unhappiness is causing decline or whether decline is causing unhappiness.

Section II also notes three other facts about urban decline and unhappiness. First, while Oswald and Wu (2010) document the relationship between state-level happiness and amenities, we find that the connection between unhappiness and decline in the BRFSS does not reflect the role of urban disamenities associated with decline, such as crime, coldness, and inequality. Second, we find that the connection between urban decline and low SWB is just as strong among recent migrants as among longer-term residents. This latter fact leans against the interpretation that happiness was ex ante identical across areas but that some areas experienced negative shocks, people were stuck in those areas, and their happiness fell accordingly.

Third, we ask whether the unhappiness of declining cities is a new phenomenon, perhaps caused by decline, or represents a more historic tendency. The General Social Survey (GSS) enables us to look back as far as the early 1970s, and these data suggest that the connection between decline and unhappiness was stronger in the past than it is today. These facts lead us to suspect that the connection with unhappiness and urban decline more likely reflects long-standing attributes of these cities rather than a causal effect of the decline itself.

In Section III, we propose a framework that incorporates spatial differences in SWB into the spatial equilibrium framework. Following writers as diverse as Epictetus, de Mandeville, Irving Fisher, and Gary Becker, we assume that happiness is desirable—but not equivalent to utility. We have life objectives other than being satisfied, and we may knowingly make choices that reduce happiness, such as exposing ourselves to a more competitive environment, if those choices further other aims (Luttmer 2005; Benjamin et al. 2011). According to the spatial equilibrium logic, a city's unhappiness must be offset by some other amenity, such as higher real income.

In our model, happiness is generated through experiences, which can be improved by spending money, and happiness is but one ingredient in the utility function. Individuals have other objectives, which we refer to as achievements, such as raising a family. These are also produced with a combination of money and time. The model suggests that the connection between money and happiness may significantly understate the connection between money and utility because a higher opportunity cost of time causes individuals to engage in less happiness-generating leisure. In a spatial equilibrium, higher wages are compensated shifts, typically offset by higher real estate prices, so higher area wages could easily be associated with lower happiness levels even if utility levels are equalized across space.

In Section IV, we examine whether individuals in declining or otherwise unhappy places are being compensated for their unhappiness. In the 1940 US Census, residents of declining cities were receiving significantly higher incomes. A one standard deviation drop in population growth post-1950 was associated with $222 more in income ($3,655 in current dollars), which is more than 10% of average income. Presumably, high labor costs were one reason why businesses left these areas. One interpretation of these results is that the industrial cities were less happy in 1940 but their residents were being compensated with earnings that could achieve other ends, such as nurturing a family.

The data also show that housing prices in 1940 were higher in areas that subsequently declined, yet there are essentially no housing quality controls in that early data. As such, while it is possible that some of the high earnings in declining cities were eaten away by higher housing rents, it is also possible that these rents were actually compensation for better housing quality.

When we turn to 2000 US Census data, we find that the unhappy declining cities are no longer receiving higher wages. Wages are essentially uncorrelated with our growth variable in the more modern data. But decline is correlated with house prices and rents. In 1940, the residents of unhappy declining places seem to have been compensated with higher incomes. In 2000, the residents of those same cities seem to have been compensated with lower housing costs.

We have also examined the direct correlation between our area-level happiness measure and area-level rents and incomes, as in Oswald and Wu (2011). We do find some evidence that residents of happy cities pay higher rents, suggesting some form of offset for the added level of happiness. The results are certainly compatible with the view that individuals trade other objectives against happiness when they are choosing where to live. Section V concludes.

I. Unhappiness across Cities

To begin, we briefly document five stylized facts about urban happiness, primarily in the United States but also abroad. We discuss the connection between unhappiness and decline in Section III.

Throughout this paper, we follow the literature in measuring happiness using self-reported survey data on subjective well-being. Our primary data source is a large national survey, the Behavioral Risk Factor Surveillance System, conducted by the Centers for Disease Control and Prevention (CDC), which asks individuals to report on their own life satisfaction using a discrete response scale.

Since 2005, the CDC has asked all respondents, “In general, how satisfied are you with your life?” Respondents were given four possible categories: “very satisfied,” “satisfied,” “dissatisfied,” and “very dissatisfied.” In each year between 2005 and 2010, around 300,000 subjects answered this question, along with all of the demographic variables listed below.[1] This ques-

[1] We discuss this survey in more detail in the appendix. We also explore the issues that arise from the discrete nature of the answers and explain why we do not think they are a problem, as well as other details of our estimation.

tion has been the focus of much of the previous literature on the economics of happiness, and we show the distribution of answers in appendix table A1. We recognize that satisfaction may strictly differ from happiness, but we will use the terms interchangeably.

In all of the work that follows, we recode these answers so that 4 indicates "very satisfied" and 1 indicates "very dissatisfied." We then rescale the answers linearly so that they have a mean of 0 and a standard deviation of 1. Because the BRFSS reports the county in which the respondent lives, we are able to link respondents to metropolitan areas.

This measure has several problems, even before considering whether it corresponds to the economic concept of "utility." First, respondents may have different interpretations of the response scale, or, equivalently, different reference points for life satisfaction. A situation that one person may consider very satisfactory may be merely satisfactory to another. If this leads to systematic differences across individuals, it could confound the variable's interpretation.

To address this, we estimate metropolitan statistical area j happiness as MSA j fixed effect in the following model:

$$y_{ijt} = \alpha + X_{ijt}\beta + \gamma_t + u_j + \epsilon_{ij}. \tag{1}$$

We estimate equation (1) at the individual level, so i indexes individual respondents, j indexes areas, and t indexes the survey wave. In this regression, y_{ijt} represents individual subjective well-being (SWB), X_{ijt} is a matrix of individual controls, u_j is a metropolitan area fixed effect, γ_t is a year fixed effect, and ϵ_{ij} is an uncorrelated error term. The individual controls include survey month, sex, a polynomial in age, eight race dummies, six marital status dummies, four educational attainment dummies, and variables representing various information about the children in the household. See the appendix for more details on these controls.

Second, respondents undoubtedly have a large degree of variability in their happiness at the moment they answer the survey. Because we only have responses from a small fraction of residents in each area (around 0.1%), this variability is likely to cause noisy estimates of area-level SWB. To account for this, we next measure area-level happiness using random effects instead of fixed effects. We estimate the following model, in which coefficients in bold type are considered to be fixed, while the others are random effects:[2]

$$y_{ijt} = \boldsymbol{\alpha} + X_{ijt}\boldsymbol{\beta} + \boldsymbol{\gamma}_t + u_j + \epsilon_{ij}. \tag{2}$$

[2] So one might prefer to call this a "mixed effects" model as opposed to a pure "random effects" model.

We consider the demographic characteristics to have a fixed relationship with individual happiness, and we allow for random metropolitan area effects as well as an individual error term.

This model enables us to compute a number of useful quantities. It allows us to calculate an estimate of the underlying variance of metropolitan area effects (σ_u^2). For each area, we can also determine the best estimate $\hat{u}_j$ of that area's u_j. We refer to these estimates as the metropolitan area's *adjusted life satisfaction*. We use them extensively in subsequent analysis as our estimate of the area's contribution to individual happiness.[3]

Finally, since the BRFSS has only asked about life satisfaction since 2005, we have a limited ability to address time-series variation in happiness. We will thus augment it with other data sources introduced below. We first turn to five sets of facts about life satisfaction across space.

A. Are There Significant Differences in Life Satisfaction across Space?

We first address whether there is a meaningful difference in happiness levels across geographic areas, both before controlling for individual demographic characteristics and after including these controls. We answer this question in multiple ways. First, we run the fixed effects model (1) and perform an *F*-test of the joint significance of the metropolitan area fixed effects. Second, we determine whether the estimated variance of metropolitan area random effects in equation (2), σ_u^2, is significantly different from zero. Third, we perform a likelihood ratio test of the random effects model (2) against a constrained model in which the random effects are removed (we force $u_j = 0$ for all j).

We run each of these tests on a model with no demographic controls and also with the full set of demographic controls shown in appendix table A2. In both cases, all three tests strongly reject the null hypothesis that metropolitan area effects are irrelevant, all with $p < .0001$.

Our next task is to quantify the differences across regions. We do so using two different measures from the random effects estimates in (2). First, σ_u^2 provides an estimate of the variance across the full population of metropolitan and nonmetropolitan areas. Second, the empirical variance of the adjusted life satisfaction values, $\mathrm{Var}(\hat{u}_j)$, quantifies the dispersion of estimates in the sample of areas where we are able to compute happiness.

In the unadjusted random effects model (where X_{ijt} is empty so we have no demographic estimates $\boldsymbol{\beta}$), we find $\sigma_u = 0.063 \pm 0.004$ and $\mathrm{sd}(\hat{u}_j) = 0.058$.

[3] Our calculation of these adjusted life satisfaction measures $\hat{u}_j$ recognizes the problem of potentially large sampling variation when measuring SWB in a survey. We therefore calculate the best linear unbiased predictor (BLUP) based on our MSA-level random effects from (2), following the method of Bates and Pinheiro (1998) as implemented in Stata 11.1.

Since all of our analyses use measures of SWB rescaled to have zero mean and unit variance across individuals, the variation across geographic regions is around 6% of the individual-level variation in happiness. These numbers shrink by about one-quarter, to $\sigma_u = 0.047 \pm 0.003$ and $\mathrm{sd}(\hat{u}_j) = 0.042$ when we include the demographic controls in model (2). Figure 1 of Glaeser, Gottlieb, and Ziv (2014), the working paper version of this paper, shows the distribution of these adjusted life satisfaction estimates.

To get a better sense of what this means quantitatively, we can compare it to the estimates of the impact of other characteristics on individual SWB. For example, moving across one standard deviation in geographic areas has an impact one-third as large as the difference between being a high school graduate or not graduating or 1.8 times the estimated male-female gap.

The values of our local happiness estimates themselves are shown visually in figure 1. This map shows adjusted life satisfaction estimated at the MSA- and rural-area level after controlling for individual demographics.[4] The map shows a band of less happy areas in parts of the Midwest and the Appalachian states, stretching from Missouri in the West and Alabama in the South well into Pennsylvania and even New Jersey in the East. New York City, Detroit, and much of California also have lower SWBs, while the happiest areas are concentrated in the West, the upper Midwest, and the rural South. Appendix table A3 shows specific values for a handful of metropolitan and nonmetropolitan regions, including the highest and lowest values that we estimate.

A third potential problem with these results is that they may reflect differences in the ways in which states implement the BRFSS. Unlike many surveys, the BRFSS is not centrally administered. Instead, individual state agencies perform the surveys. We cannot be sure of what biases may be created through this decentralized implementation, but it is at least possible that state-level implementation has caused some of the variance that we see in the data.

To address this possibility we reestimate model (2) controlling for state fixed effects. Since there are a relatively few number of metropolitan areas in many states, we will not use these state-corrected area fixed effects in general. Still, it is important to note the reduction in variance that occurs when we look only at the within-state variance. The standard deviation of $\hat{u}_j$ falls from 0.043 to 0.017 when we control for state fixed effects as well as demographic controls. The variance is significantly reduced, but these effects remain statistically distinct from zero. As such, we conclude that

[4] Because income depends on numerous individual choices, including where to live and possibly including happiness, we do not include it among our demographic controls. For the interested reader, our working paper version of this paper (Glaeser et al. 2014) shows a version of this map after adjusting for individual income.

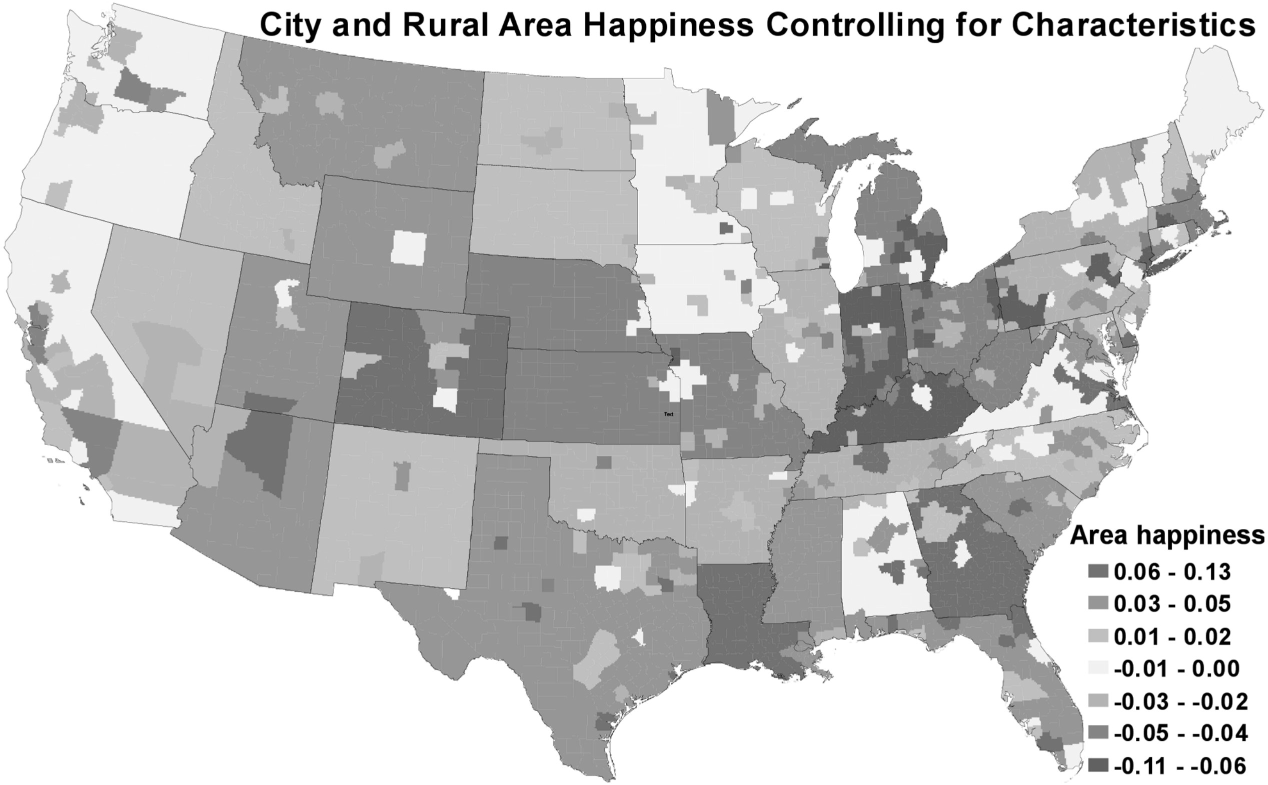

FIG. 1.—Estimated metropolitan- and rural-area-adjusted happiness. This figure shows each metropolitan and rural area's adjusted life satisfaction, after controlling for demographic covariates in a mixed effects model. Data are from Centers for Disease Control and Prevention (2005–9).

metropolitan differences would persist even if all the state-level variation reflected only state-level differences in implementing the BRFSS.

This evidence does not rule out the possibility that these differences reflect unobserved individual characteristics. One approach to unobserved heterogeneity is to estimate metropolitan area fixed effects controlling for individual fixed effects. This requires us to use a panel, rather than a repeated cross section, which forces us to move from the very large BRFSS to the much smaller National Survey of Families and Households (NSFH). The NSFH is a longitudinal study, from which we use the first two waves (completed in 1987–88 and 1992–94; Sweet, Bumpass, and Call 1988; Sweet and Bumpass 1996; Trull and Famularo 1996).

In both waves, the data contain information on family and personal characteristics of individuals and on individual subjective well-being. In particular, the NSFH asks: "First taking things all together, how would you say things are these days?" Appendix table A1 shows the responses.

We will later use this measure to examine whether the link between area attributes and well-being is stronger for recent migrants or long-term residents. Here we restrict our attention to the heterogeneity in subjective well-being across space. We first estimate adjusted life satisfaction for the merged sample of NSFH waves 1 and 2. The variance of these estimates is 0.0007, roughly in line with the estimates from the BRFSS. The raw variation of metropolitan area fixed effects is larger in the NSFH, but the variance correction is also much larger because the sample size is so much smaller.

We then estimate a Primary Metropolitan Statistical Area (PMSA) fixed effect variable using the two waves including individual level fixed effects. The correlation between these estimates and the estimates without the individual fixed effects is 0.69. The variance of the PMSA fixed effect with individual fixed effects is 0.64. We conclude from these results that there appears to be significant variation in subjective well-being across space even when we control for unobservable individual-level heterogeneity by using individual fixed effect estimates.

B. Do Metropolitan Area Differences in Subjective Well-Being Persist?

Having established the existence of spatial differences in happiness and estimated their magnitude, we now want to see how they evolve over time. Hypotheses about the temporal pattern of spatial SWB could range from a completely permanent local characteristic (e.g., Honolulu has gorgeous weather and is on the beach, which always makes its residents happy), to a long-term shock common to area-level residents (e.g., the economy in Detroit was poor and declining during our sample period, making its residents unsatisfied), to an extremely transitory common shock caused by the weather or local sports outcomes.

We first test the stability of area effects in two ways. First, we run versions of equation (2) separately for each year, so without year fixed effects,

$$y_{ij} = \boldsymbol{\alpha}^{(t)} + X_{ij}\boldsymbol{\beta}^{(t)} + u_j^{(t)} + \epsilon_{ij}. \tag{3}$$

We then compare the adjusted life satisfaction estimates across different years ($\hat{u}_j^{(t)}$ vs. $\hat{u}_j^{(t')}$ for $t' \neq t$). Our second method is to augment equation (2) by adding an area-year random effect, v_{jt}, to the random effects regression:

$$y_{ijt} = \boldsymbol{\alpha} + X_{ijt}\boldsymbol{\beta} + \boldsymbol{\gamma}_t + u_j + \nu_{jt} + \epsilon_{ij}. \tag{4}$$

This model estimates the time-invariant area effect, u_j, and the time-varying area effect, v_{jt}, simultaneously. We can test the statistical impact of each of these effects separately and quantify the importance of permanent and transitory area effects. For this analysis, we use the sample of respondents in the 177 MSAs with at least 200 respondents in all years of our sample.

These tests reveal very clearly that the permanent effects are far more important than the transitory components. When estimating equation (4) without demographic controls, we find $\sigma_u = 0.064 \pm 0.004$, while $\sigma_v = 0.018 \pm 0.002$. Thus, there is a statistically significant transitory component, but it varies by 70% less than the permanent area component, and its standard deviation is around 2% of the individual-level standard deviation.

Another way to see this variation is to relate adjusted life satisfaction from equation (3) in one year to that in another year. Using the measures adjusted for demographic controls from 2005 ($\hat{u}_j^{(2005)}$) and from 2009 ($\hat{u}_j^{(2009)}$), we find an extremely strong positive relationship, with a correlation of 0.48. Thus one-quarter of the variation in adjusted life satisfaction is driven by permanent metropolitan area–level shocks and the rest by transitory shocks and estimation error. Although we have adjusted for the effect of sampling error in computing adjusted life satisfaction, we should expect to see a correlation less than one if our correction is imperfect. Hence, the random effects results discussed in the previous paragraph give the more accurate assessment of the relative importance of permanent and transitory components to well-being.

C. Is Urbanization Associated with Happiness or Unhappiness?

One natural question is whether happiness increases or diminishes in large cities. Cities have often been seen as entities that create financial wealth but diminish other types of well-being. We first test this hypothesis by examining the correlation between adjusted life satisfaction and the logarithm of metropolitan area population. If we use the 2010 population, we find a weak positive correlation of 0.07. As metropolitan area population increases by one log point, SWB increases by 0.003 standard deviations, and the effect is quite imprecise. We will later show that in an individual-level regression, there is also no significant relationship between area-level population and self-reported well-being, holding individual-level characteristics constant.

Using past population levels, instead of current levels, we find a positive correlation with population in 2000 and 1990 and a negative correlation with population levels before that point. The relationship between recent levels of SWB and metropolitan area population before 1960 becomes significantly negative. While larger cities today do not evince significantly lower levels of unhappiness, residents of cities that were large in the past do seem to be less happy. We return to this topic later, when we discuss the connection between SWB and population growth.

Following Stevenson and Wolfers (2008), we now briefly turn to worldwide data. Using the World Values Survey, we estimate subjected well-being in rural and urban areas in 39 countries throughout the world. Across the entire sample, we find that urban happiness is on average higher than rural happiness. This effect is, however, driven primarily by poorer countries.

Figure 2 shows the correlation between the logarithm of per capita gross domestic product (GDP) in the country in 2007 and the rural-urban gap in subjective well-being. The coefficient is significantly negative, and the R^2 is

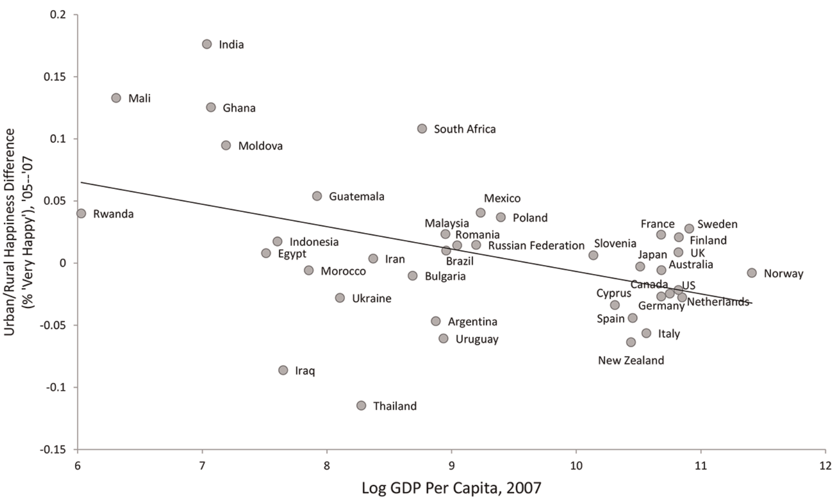

FIG. 2.—Urban-rural happiness gradient versus gross domestic product across countries. Regression line: $y = -.0181\log\text{GDP} + 0.174$, $N = 39$, $R^2 = 0.20$. The source is World Values Survey and World Bank.

0.2. In poorer countries, which often have cities that seem particularly hellish, the residents of cities say that they are significantly happier than the residents of rural areas. It is perhaps unsurprising that the developing world is urbanizing so rapidly as urban residents appear to be both far better paid and happier.

D. Unhappiness and Urban Characteristics

We now turn to area-level correlates of self-reported well-being. Most area-level attributes are relatively uncorrelated with subjective well-being, at least conditional on individual education.

Table 1 presents these facts using area characteristics as of the year 2000. We use the year 2000 both because it predates our well-being data and because it is the last year with a comprehensive US Census. Our core specification includes a bevy of individual attributes that have been found to correlate with happiness, including education, age, race, and family status. We do not include income or employment controls, as these represent outcomes that may be caused by an area's economic success. Education and marital status may themselves be determined by the urban environment, and we include regressions both with and without those controls. All regressions cluster the standard errors at the area-year level and include year and month fixed effects.

Column 1 shows the relationship between the population size of the metropolitan area and self-reported well-being. When we do not control for education and marital status, the statistical relationship is small and statistically indistinct from zero. When we include these more endogenous controls, the relationship becomes more negative and statistically significant.

In column 3, we control for the share of the adult population in the area with a college degree. Using the fixed effects estimated without controlling for individual-level education, this variable is strongly positive. Using the fixed effects estimated conditional on these controls, the variable's estimated effect drops by two-thirds, and it becomes statistically indistinct from zero. Columns 5 and 6 examine the share of the adult population with a college degree. The picture is much the same as with the other education variable. The coefficient is large and statistically significant when we control only for area-level attributes but not when we control for area-level education. These regressions can be interpreted as suggesting that area-level education boosts self-reported well-being by increasing individual educational attainment, or that area-level education has no independent effect.

Columns 7 and 8 examine racial segregation, as measured by a standard dissimilarity index. In this case, we also interact segregation with a dummy variable that takes on a value of one if the individual is black. Both with and without individual controls, segregation is negatively associated with

Table 1
Happiness Levels across Space, BRFSS

	Dependent Variable = Self-Reported Well-Being									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log population, 2000	−.0066	−.0085**							−.00325	.000921
	(.00577)	(.00371)							(.00463)	(.00334)
% BA graduate, 2000			.330***	.0977					.394***	.185***
			(.108)	(.0768)					(.114)	(.0624)
% High school graduate, 2000					.421***	.0716			−.276**	.0893
					(.151)	(.116)			(.125)	(.0833)
Segregation index, 2000							−.160***	−.130***	−.130***	−.0323
							(.0379)	(.0250)	(.0342)	(.0326)
Segregation × black							−.263***	−.144***	−.113**	−.0573
							(.0507)	(.0482)	(.0440)	(.0387)
Segregation × Asian									.0506	.0872
									(.0882)	(.0582)
Segregation × HPI									−.0370	−.0488
									(.151)	(.154)
Segregation × other									.217**	.230**

									(.101)	(.101)
Segregation × AIAN									−.0695	−.0576
									(.0832)	(.0822)
Segregation × multiracial									−.160***	−.140***
									(.0618)	(.0527)
Segregation × Hispanic									−.0549	−.252
									(.247)	(.205)
Log median house value									−.0398	−.0447***
									(.0258)	(.0114)
Additional controls	No	Yes	No	Yes	No	Yes	No	Yes	Yes	Yes
State fixed effects	No	No	No	No	No	No	No	No	No	Yes
No. of observations (1,000s)	1,185	1,185	1,185	1,185	1,185	1,185	1,134	1,134	1,134	1,134
R^2	.007	.076	.008	.076	.008	.076	.008	.076	.076	.078

SOURCES.—Authors' regressions on microdata from the Behavioral Risk Factor Surveillance System Survey (CDC), the US Census (Ruggles et al. 2010), and Glaeser and Vigdor (2001).

NOTE.—All regressions control for year fixed effects, month fixed effects, age, race, and sex. "Additional controls" include education, marital status, and family size. Standard errors in parentheses are clustered at the metropolitan statistical area level. HPI and AIAN are dummy variables for self-reported race of Hawaiian and Pacific Islander and American Indian and Alaskan Native, respectively.

** $p < .05$.

*** $p < .01$.

well-being and this effect is approximately twice as large for African Americans as for whites.

Column 9 controls for all of the metropolitan area variables and has the full set of individual-level controls. We also interact segregation with all of the race categories. The results here remain similar to our previous specifications. The population share with a college degree has a positive effect on self-reported happiness, although the share with a high school degree has a negative relationship. Segregation continues to have a negative connection to self-reported well-being, and this effect is much stronger for African Americans. In this specification, housing value has a somewhat surprisingly negative but insignificant effect on subjective well-being.

Column 10 adds state fixed effects and is our most complete specification. As many states have only one metropolitan area, this reduces our effective sample and eliminates any variation that represents larger regional trends. In this specification, the positive effect of college education remains and the effect of high school graduates becomes positive. Housing values now have a negative and significant relationship with happiness, while segregation is insignificant.

Putting together these results, we draw two tentative conclusions. There is some possibility that individuals report higher levels of well-being in more educated areas, although this is true only when we include a full range of area controls or when we fail to control for individual-level education. Segregation is associated with lower levels of subjective well-being but only when we do not control for state fixed effects. Overall, these results do not suggest a robust series of correlations between urban attributes and SWB.

II. Unhappiness and Urban Decline

We now turn to a particularly striking correlation between urban unhappiness and decline (Glaeser and Redlick 2009; Lucas 2013). We first examine linear specifications and then allow the impact of population growth on subjective well-being to have a piecewise linear shape. We will focus on changes in the logarithms of population and median household income between 1950 and 2000. We first focus on the BRFSS and then turn to the NSFH and the GSS, which enable us to look at movers and estimate equations with individual fixed effects.

A. Linear Effects of Population Growth and Income

Table 2 presents our first set of results on the correlation between SWB and urban change. Columns 1–3 show results for population change. Column 4 shows results for income change. Columns 5 and 6 show results for both variables together and include other area-level controls.

Table 2
Happiness and Urban Change, BRFSS

	Dependent Variable = Self-Reported Well-Being					
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Change in log population, 1950–2000	.0635***	.0412***	.0174***		.0270***	.00312
	(.0146)	(.00925)	(.00649)		(.0104)	(.00710)
Change in log income, 1950–2000				.0586***	.0597**	.0301
				(.0156)	(.0298)	(.0216)
% BA graduate, 2000					.124	−.0260
					(.113)	(.0406)
% High school graduate, 2000					−.170	.185***
					(.140)	(.0706)
Segregation index, 2000					−.0930***	−.0484***
					(.0230)	(.0180)
Additional controls	No	Yes	Yes	Yes	Yes	Yes
State fixed effects	No	No	Yes	Yes	No	Yes
No. of observations	1,182,563	1,182,563	1,182,563	1,166,056	1,114,898	1,114,898
R^2	.008	.076	.078	.078	.077	.078

Sources.—Authors' regressions on microdata from the Behavioral Risk Factor Surveillance System Survey (CDC), the US Census (Ruggles et al. 2010), and Glaeser and Vigdor (2001).

Note.—All regressions control for year fixed effects, month fixed effects, age, race, and sex. "Additional controls" include education, marital status, and family size. Standard errors in parentheses are clustered at the metropolitan statistical level.

** $p < .05$.
*** $p < .01$.

Column 1 shows the relationship between population change and self-reported well-being, controlling for individual attributes. The coefficient of 0.0635 implies that a one log point increase in population growth is associated with about a one-sixteenth of a standard deviation increase in self-reported well-being. Column 2 controls for the more endogenous individual characteristics. The coefficient on population change remains statistically significant, but it falls in magnitude by about one-third. Column 3 controls for state fixed effects. In this case, the coefficient falls to about one-fourth of its value in the first regression, although it retains statistical significance.

Column 4 looks at income change instead of population change, including all of the same controls as column 3. In a sense, this is the local version of the classic Easterlin (1974) work on income change and happiness. It shows a strong positive relationship between income growth and self-reported well-being. The coefficient is somewhat larger than that on population growth, but since the variation of income growth is smaller, the impact of a one standard deviation change in income growth is actually smaller than the impact of a one standard deviation change in population growth.

Column 5 includes both change variables and other area-level controls, and the change variables both remain statistically significant. The coefficients are modest but continue to suggest that growth is associated with positive levels of well-being. The only other control that is statistically distinct from zero is segregation, which remains negative.

As the BRFSS is administered at the state level, there could be cross-state differences in the survey's implementation. To adjust for any such differences, we add state fixed effects in column 6. Note that the fixed effects may be overcontrolling in important ways as they eliminate all regional variation from our estimates. Our map of adjusted happiness (fig. 1) shows clear regional patterns, and we now eliminate that variation and more. With this caveat, column 6 shows that these fixed effects eliminate the otherwise robust relationship between urban growth and subjective well-being. But unfortunately they do not tell us whether this reflects variability in survey implementation or is simply because the bulk of geographic differences in subjective well-being are regional in nature.

As we will see in the next subsection, the growth-happiness relationship is driven by the lower end of the city growth distribution. This part of the relationship remains robust to state fixed effects, reducing the importance of distinguishing between columns 5 and 6 of table 2.

B. The Nonlinear Relationship between Unhappiness and Population Growth

Figure 3 shows the correlation between population growth and adjusted life satisfaction. As the figure makes clear, the effect is much stronger at the lower end of the population change distribution. Low levels of happiness are particularly common in areas with declining population, but higher levels of happiness are not especially prevalent in areas where population is growing rapidly.

There are several hypotheses that could explain this nonlinearity. For example, if decline is actually causing unhappiness—rather than merely being correlated with it—it might be that decline itself creates urban stresses, relative to stasis, but that urban growth does not particularly alleviate those stresses. Declining cities, such as Detroit, often find it difficult to cover the costs of their historic footprint and infrastructure. Decline may be associated with crumbling social or physical infrastructure. It could also be that unhappiness is caused by other attributes that cause decline but that, among growing cities, the differences come mainly from differences in housing supply and economic productivity, which perhaps have little impact on happiness.

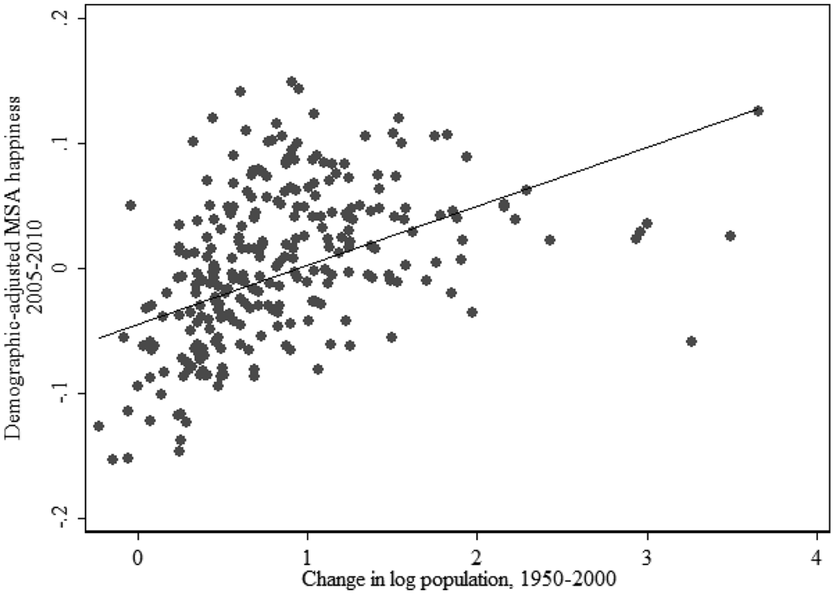

FIG. 3.—Population change and adjusted happiness. This figure shows each metropolitan and rural area's adjusted life satisfaction, after controlling for demographic covariates in a mixed effects model, against metropolitan statistical level population change from 1950 to 2000. Data are from CDC (2005–10).

Whatever the cause, the nonlinear relationship is obvious in the data. Table 3 shows the connection between SWB and urban growth in the BRFSS, where we have allowed the break in the slope to occur at a value of 0.75, the median for our sample of metropolitan areas. Column 1 shows that, controlling for exogenous demographic controls, the coefficient on growth when growth is below the median is 0.214, meaning that a 0.5 change in log population growth is associated with a 0.1 standard deviation increase in SWB. The result is extremely significant and remains so in the second column, where we include the endogenous demographic controls. The coefficient here drops to 0.134, meaning that a 0.5 change in log population growth is associated with a 0.065 standard deviation increase in SWB. This is roughly equivalent to one standard deviation of the metropolitan area fixed effects, and it is roughly equivalent to the difference in SWB between high school graduates and individuals who have some college education.

In column 3, we control for income and employment status. While we recognize the endogeneity of these outcomes with respect to local labor market conditions, we still think it is worthwhile knowing whether the connection between urban decline and unhappiness disappears when

Table 3
Happiness and Urban Population Growth Differences, BRFSS

	Dependent Variable = Self-Reported Well-Being							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in log population (below median), 1950–2000	.214***	.134***	.101***	.0972***	.0781***	.156***	.0863***	.0830*
	(.0186)	(.0146)	(.0174)	(.0180)	(.0127)	(.0344)	(.0251)	(.0446)
Change in log population (above median), 1950–2000	.00409	.00503	.00929	.00771	−.00443	.0231	.0134	.0164
	(.0127)	(.00795)	(.00711)	(.00642)	(.00564)	(.0331)	(.0224)	(.0240)
Average January temperature						.00153**	.00177***	−.00281**
						(.000780)	(.000442)	(.00134)
Precipitation						.000117	−.000146	.00190**
						(.000682)	(.000480)	(.000840)
Log of crime						.00457	.00258	.0115
						(.0115)	(.00703)	(.00707)
Pollution						.000171	.000487	−3.81e-05
						(.00140)	(.000955)	(.00107)
Gini coefficient, 2000						−.0538	.325	.863***
						(.643)	(.420)	(.237)
Additional controls	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Employment and income controls	No	No	Yes	Yes	No	No	No	No
Health controls	No	No	No	Yes	No	No	No	No
State fixed effects	No	No	No	No	Yes	No	No	Yes
No. of observations	1,182,563	1,182,563	1,182,563	1,164,203	1,182,563	261,987	261,987	261,987
R^2	.009	.077	.125	.185	.078	.010	.078	.079

Sources.—Authors' regressions on microdata from the Behavioral Risk Factor Surveillance System Survey (CDC) and the US Census (Ruggles et al. 2010).

Note.—All regressions control for year fixed effects, month fixed effects, age, race, and sex. "Additional controls" include education, marital status, and family size. Standard errors (in parentheses) are two-way clustered (Cameron, Gelbach, and Miller 2011) at both the metropolitan statistical area and year levels.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

we control for them. In this case, the coefficient falls to 0.101. In column 4, we control for health status (including both a general question about overall health status and a question about days spent ill over the past year). The coefficient falls to 0.097 and remains quite statistically significant. Column 5 adds state fixed effects but excludes the health and income questions. Here we estimate a coefficient of 0.078. Appendix table A4 demonstrates that the relationship in this column is robust to numerous other functional forms for the nonlinearity.

C. Can Urban Disamenities Explain the Correlation between Unhappiness and Urban Decline?

We next consider whether observable urban disamenities can explain the correlation between unhappiness and urban decline. We add a variety of correlates of decline to our previous specifications and ask whether these variables reduce the coefficient on urban decline. While none of these estimates can be taken as being causal, they represent a rough pass at judging whether the correlation between decline and happiness merely represents the correlation between decline and some other more important variable.

Our first control is January temperature. The correlation between warm weather and metropolitan growth is well known (Glaeser and Tobio 2008), and it is certainly possible that tough winters are depressing. While Oswald and Wu (2010) find that climate has a significant relationship with self-reported happiness, we find no connection in our specification once we have controlled for population growth nonlinearly. Moreover, this control does little to the estimated coefficient on population decline.[5] Note that we have a smaller sample in these regressions because we lack crime data for many of our metropolitan areas.

Our second climate variable is precipitation, measured in annual inches of rain. This variable again has little correlation with SWB in our data after controlling for population decline. The third variable we test is the log of the number of serious crimes per capita, and it also has no detectable relationship. While being victimized may certainly make someone unhappy, it seems quite possible that crime is sufficiently concentrated in certain population subgroups that it has little impact on average happiness.

The fourth variable captures pollution, which might well be higher in America's erstwhile industrial heartland. We have tried many different measures of local pollution levels, but none of them correlate well with happiness. Here we include total particulates (mean of 10 micron particulate matter, from 2000), and it has little correlation with happiness.

[5] In our working paper (Glaeser et al. 2014), we also show regressions that include these controls one at a time. That approach has no impact on the conclusions we draw here.

Our fifth variable is the Gini coefficient, which measures income inequality as of the year 2000. While Alesina, Di Tella, and MacCulloch (2004) find that happiness decreases with inequality, especially in Europe, we find a slight positive relationship between happiness and inequality across US metropolitan areas. Moreover, controlling for inequality does little to change the estimated impact of population decline on unhappiness. The weak connection between inequality and SWB in the BRFSS is somewhat odd, because it is quite strong in the GSS (Glaeser, Resseger, and Tobio 2009).

Column 6 shows that including all of these variables reduces the coefficient on decline from 0.21 to 0.156. Adding the endogenous demographic controls causes the coefficient to drop further to 0.08. This should be compared with column 2's coefficient of 0.134, which is the effect of decline on happiness without these other amenity controls but with endogenous demographics. Finally, in the last regression, we include state fixed effects, which change the coefficients on some of the area amenity controls but have little impact on population growth coefficients.

Throughout the specifications in this table, bolstered by the robustness checks in appendix table A4, the coefficient on urban decline is statistically robust and the magnitudes remain quite similar. While it is certainly true that controlling for education and family status significantly reduces the estimated coefficient, other individual controls change the coefficient only slightly. We believe that this suggests that this effect is less likely to reflect unobserved heterogeneity, but to address that issue, we now turn to the NSFH.

D. Urban Decline and Unhappiness with Movers and Stayers

We now ask whether the unhappiness of declining cities appears to be limited to longer-term residents or whether they are similar for recent migrants. We explore this question in the first two columns of table 4 using the NSFH. In its first two waves, the NSFH is a clean panel that can, in principle, enable us to look at SWB for people who move between areas. Two significant challenges with the NSFH are that the samples are small and the time between the first and second waves is small (under 5 years), so the number of movers is smaller still.

Column 1 shows the effect of the population growth spline with exogenous demographic controls. The coefficient is 0.14, somewhat smaller than in the BRFSS, but the question is different, and the controls are not identical. This result confirms our baseline finding in a separate data set.

We next explore whether people who move to a new metropolitan area experience the happiness level of the area where they are newly arrived. To do this, we look at observations in the second wave of the NSFH, where we can distinguish between recent movers and previous residents. In column 2, we estimate whether the coefficient on decline is different for individuals who moved into the metropolitan area between the first wave and the second wave. The coefficient on decline for stayers is very similar to that in the previous regression. The interaction between the decline measure and being a mover is negative, meaning that decline is less strongly associated with unhappiness for the movers. However, while the interaction is not small, it is not distinct from zero. It is difficult to conclude much from this regression.[6]

One possible explanation for the relationship between decline and unhappiness is selective migration. Individuals who leave declining cities may be happier than their neighbors or growing cities may attract individuals who are happier than the population as a whole. The panel nature of the NSFH allows us to test this hypothesis.

Columns 3 and 4 use the entire NSFH sample in order to test whether individual and PMSA characteristics in wave 1 predict whether an individual moves between waves 1 and 2. Both columns control for subjective well-being in wave 1, the population growth spline of an individual's wave 1 PMSA, our standard set of individual controls, and income in wave 1. In column 3, the upper part of the population growth spline positively predicts mover status, reflecting the high degree of population churn in the upper tail of growing cities. Subjective well-being in wave 1 is not predictive of whether an individual will move between waves 1 and 2.

Column 4 adds an interaction between individual wave 1 subjective well-being and the population growth spline. Subjective well-being is now marginally positive and significant. Critically, the interaction between the lower spline and subjective well-being is negative and significant. Happier people are less likely to leave declining cities, relative to rising cities. Put another way, we can reject the hypothesis that the happiest individuals are selectively moving out of declining areas.

In the remaining columns, we focus on the subsample of the NSFH who moved MSAs between waves 1 and 2. In columns 5 and 6, we test for the hypothesis that happier migrants select growing cities. We focus on the 935 movers in our sample for which we have data on PMSA population for waves 1 and 2. In column 5, we use our set of exogenous controls, wave 1 individual subjective well-being, and wave 1 PMSA population growth spline. In column 6, we add controls for wave 1 endogenous individual characteristics and income. Although we cannot reject a positive relationship between wave 1 subjective well-being and wave 2 PMSA population growth, we find no evidence for selection of individuals with higher subjective well-being into growing cities.

[6] In table 5 of the working paper version (Glaeser et al. 2014), we run additional robustness checks on these regressions, including additional controls, fixed effects, and sample restrictions. These have little impact on our conclusion here.

Table 4
Happiness, Urban Decline, and Mobility, NSFH

	Dependent Variable							
	Self-Reported Well-Being		Mover Status		Wave 2 PMSA Growth		Wave 2 PMSA BLUP	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in log population (below median), 1950–2000	.141***	.142**	.00942	.00804	.580***	.619***		
	(.0362)	(.0702)	(.0331)	(.0333)	(.176)	(.178)		
Change in log population (above median), 1950–2000	−.0574***	−.0387	.0422***	.0431***	.0139	−.00488		
	(.0160)	(.0356)	(.0151)	(.0147)	(.0988)	(.0997)		
Mover		.0869						
		(.124)						
Mover × change in log population (below median), 1950–2000		−.0851						
		(.211)						
Mover × change in log population (above median), 1950–2000		.0217						
		(.0681)						
SWB × change in log population (below median), 1950–2000				−.0327**				
				(.0160)				

SWB × change in log population (above median) 1950–2000				.0152				
				(.0119)				
Wave 1 subjective well-being (SWB)			−.00137	.0128*	.0227	.0218	.00227	.00165
			(.00352)	(.00723)	(.0262)	(.0279)	(.00150)	(.00149)
Wave 2 subjective well-being (SWB)	−.0131***							
	(.00403)							
Wave 1 PMSA BLUP (below median)							.170***	.152**
							(.0598)	(.0583)
Wave 1 PMSA BLUP (above median)							.118	.108
							(.112)	(.117)
Log Household income			−.0123**	−.0124**		−.00476		−.00353**
			(.00537)	(.00537)		(.0332)		(.00171)
Additional controls	No	Yes	Yes	Yes	No	Yes	No	Yes
No. of observations	17,019	8,491	8,528	8,528	935	935	1,513	1,513
R^2	.010	.046	.066	.066	.093	.113	.039	.062

SOURCES.—Authors' regressions on microdata from the National Survey of Families and Households (Sweet and Bumpass 1996) and the US Census (Ruggles et al. 2010). SWB = subjective well-being.

NOTE.—All regressions control for month fixed effects, age, race, and sex. "Additional controls" include education, marital status, and family size. Standard errors in parentheses are robust to heteroskedasticity.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

Finally, columns 7 and 8 analogously assess whether happier migrants select happier cities, using data from all 1,513 movers in the NSFH. In this specification, we also find little connection between wave 1 "happiness" for movers and choosing, conditional upon moving, to relocate to a happier locale. In column 8, we add controls for endogenous characteristics and income in wave 1. The positive relationship between wave 1 and wave 2 subjective well-being decreases in size and continues to be insignificant. The data do not support the hypothesis that unhappy migrants choose declining or unhappy cities, but the results are not strong enough to reject the possibility of selective migration.

Table 5 now turns to a different data set, the GSS. The public version of the GSS contains state name and city-level population. These two variables enable us to predict the population decline in the area with a fairly high degree of accuracy for the overwhelming majority of data points.

In column 1, we again estimate the spline controlling for exogenous individual attributes, this time in the GSS. We continue to find a strong positive relationship between growth and happiness for the areas where growth rates are below the sample median. The regression of column 2 interacts these variables with an indicator variable denoting whether the individual has moved among metropolitan areas since age 16. The interaction between this variable and population decline is negative but very close to zero. In this larger sample, we see little evidence suggesting that the unhappiness associated with urban decline is limited to longer-term residents.

E. Is the Unhappiness of Declining Cities New or Old?

Historical data can help us assess whether the relationship between unhappiness and decline reflects the impact of decline itself or whether these now declining cities were historically defined more by productivity than by pleasure. According to the first view, Detroit was once a place of happiness as well as prosperity, but as the prosperity declined and the social problems increased, unhappiness spread. According to the second view, Detroit was unhappy even during its heyday, but historically its residents were well compensated for their joylessness. Capital and labor located in the city historically because it had natural advantages, such as access to waterways, that made up for the loss in happiness.

The era of comprehensive urban happiness measures really only began 10 years ago with the BRFSS. The NSFH goes back 20 years, but even that is a relatively short historical window. To investigate the more distant past, we turn to the GSS. The GSS has relatively comprehensive personal controls, but it still dates back only to the early 1970s.

Our approach is to estimate the impact of area-level population change and then to examine how this effect changes over the decades. We do this

Table 5
Happiness Regressions Using the General Social Survey

	Dependendent Variable = Self-Reported Happiness				
	(1)	(2)	(3)	(4)	(5)
Change in log population (below median), 1950–2000	.214***	.222***	.521***	.485***	.459***
	(.0527)	(.0818)	(.104)	(.113)	(.113)
Change in log population (above median), 1950–2000	−.0295	−.0382	−.0961	−.0819	−.0752
	(.0438)	(.0658)	(.0768)	(.0853)	(.0803)
Change in log population (below median), 1950–2000 × individual moved		−.0355			
		(.118)			
Change in log population (above median), 1950–2000 × individual moved		.0543			
		(.0780)			
Moved		.0378			
		(.0555)			
Change in log population (below median), 1950–2000 × 1980 decade dummy			−.287**	−.296**	−.277**
			(.134)	(.123)	(.122)
Change in log population (below median), 1950–2000 × 1990 decade dummy			−.372***	−.451***	−.421***
			(.121)	(.0992)	(.0952)
Change in log population (below median), 1950–2000 × 2000 decade dummy			−.556***	−.446**	−.440***
			(.144)	(.169)	(.159)
Change in log population (below median), 1950–2000 × 2010 decade dummy			−.510***	−.344	−.298
			(.172)	(.225)	(.214)
Change in log population (above median), 1950–2000 × 1980 decade dummy			.130	.134	.128
			(.103)	(.103)	(.0969)
Change in log population (above median), 1950–2000 × 1990 decade dummy			.0398	.113	.111
			(.0978)	(.128)	(.118)
Change in log population (above median), 1950–2000 × 2000 decade dummy			.0107	.0116	.00589
			(.0795)	(.110)	(.106)
Change in log population (above median), 1950–2000 × 2010 decade dummy			.188**	.161	.180

Table 5 (*Continued*)

	Dependendent Variable = Self-Reported Happiness				
	(1)	(2)	(3)	(4)	(5)
Additional controls	No	Yes	No	Yes	Yes
Income and employment controls	No	Yes	No	No	Yes
No. of observations	9,995	7,541	9,995	7,541	7,541
R^2	.021	.051	.024	.040	.054

SOURCES.—Authors' regressions on microdata from the General Social Survey and the US Census (Ruggles et al. 2010).

NOTE.—All regressions control for a year trend, age, race, and sex. "Additional controls" include education, marital status, and family size. Standard errors in parentheses are clustered at the metropolitan statistical area level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

using the GSS in the last three columns of table 5. We again estimate a spline for population growth, but we interact the coefficients on that spline with indicator variables that represent each decade. The population growth is defined over the entire 1950–2000 period, but the interactions allow the connection between decline and happiness to differ across the decades.

In column 3, we control for standard demographic variables and a year trend variable. As the regression shows, the interaction is strongly negative after the 1970s, meaning that the correlation between unhappiness and decline has decreased over time. Indeed, by the 2000s, the connection has disappeared entirely, which is, of course, not what we observed in the BRFSS. This shows that the cities that are declining over the entire period were unhappier in the 1970s, relative to other areas, than they were after 2000. These results are compatible with the view that unhappiness caused the decline or that declining cities have long-standing attributes associated with unhappiness, but they do not seem compatible with the view that unhappiness has grown following decades of decline.

Column 4 of table 5 includes controls for the endogenous demographics. While the overall negative relationship weakens, the time pattern is unchanged. The final regression includes controls for income and unemployment. Again, the basic time pattern remains clear. Declining cities were even unhappier in the past than they are today.

Our working paper presents results from even farther in the past. Using Gallup surveys from the 1940s, we show a significant negative connection between unhappiness and city population during those years, although that connection is not stronger in states that experienced more subsequent decline (Glaeser et al. 2014). Nonetheless, eight of the 10 largest US cities in 1950 lost at least one-fourth of its population over the next 50 years. So these results support the view that the large cities of the 1940s, which typically experienced subsequent decline, were also places marked by somewhat lower happiness levels during their heyday.

These results are hardly definitive, but taken together they suggest that urban unhappiness is not exclusively recent. The GSS shows larger results in the past than in the present. This corresponds to results we can see in the BRFSS estimates. The correlation between log metropolitan area population in 2010 and adjusted life satisfaction is 0.03. The correlation between that same happiness outcome and log area population in 1950 is −0.28.

III. Why Does Happiness Differ across Space?

If self-reported happiness has any equivalence to the economist's concept of utility, then modestly enduring differences in self-reported life satisfaction seem to challenge the view that migration and the free operation of housing markets ensure that utility levels are equalized across space. Alternatively, if there are persistent differences in subjective well-being for identical people across space and a spatial equilibrium does hold, then this would imply that subjective well-being is just not equivalent to the economists' conception of utility.

Perhaps the differences that we measured above may not really represent differences in subjective well-being among otherwise identical human beings. The residents of declining cities may have less marketable skills, of various forms, than residents of growing cities, and as such, they would naturally earn less and have lower levels of life satisfaction or utility in any metropolitan area. Yet our results control for a bevy of individual characteristics, and controlling for added metropolitan area–level variables, including the percent with college degrees or the share of the population that is white, only modestly reduces the relationship between decline and self-reported life satisfaction. The estimated relationship actually increases in magnitude if we restrict our samples to metropolitan areas with relatively similar levels of college graduates.[7] Finally, while the individual fixed effects results on urban decline were inconclusive outside of the larger cities, there are still significant differences in SWB across cities when we control for individual fixed effects.

It is also possible that individuals on the margin of moving across areas receive the same welfare but that inframarginal individuals differ in their average level of well-being across space. Yet, for this view to be correct, we would need an explanation of why the average inframarginal welfare in

[7] The coefficient when our happiness variable is regressed on the change in population between 1950 and 2000 is 0.023. When we control for share of the population with college degrees and percent white, the coefficient drops to 0.02. When we restrict our sample to metropolitan areas in which 20%–30% of adults have college degrees, then the coefficient rises to 0.031.

declining areas is significantly lower than in growing areas even if the marginal happiness levels are the same.

Another interpretation is that, when equivalent individuals made location decisions, their expected happiness was equal across space but that ex post some migrants have fared worse than others, either because they were bad at projecting the happiness that different places bring or because some areas have received particularly adverse shocks. According to this view, ex post welfare differs across space, even though ex ante welfare does not. But if this view is correct, then the connection between urban decline and unhappiness should exist primarily for longer-term residents of the area, such as people who are unlikely enough to be born in metropolitan areas in decline, not recent migrants. Yet, as we have discussed, the connection between unhappiness and urban decline is stronger for individuals who chose to come to the area as adults than for individuals who were born into the area. Moreover, given the well-advertised urban problems of many declining cities, such as Detroit, it is hard to imagine that migrants are all that surprised—although it is certainly true that there are general problems in forecasting happiness (Gilbert 2006; Kahneman and Krueger 2006).

A. Is Happiness a Measure of Utility?

While we acknowledge that the preceding facts are open to multiple interpretations, we focus on one fact in particular: individuals maximize neither happiness nor life satisfaction and will sacrifice both for the right reward. In this view, individuals in less happy areas forgo well-being in exchange for some other advantage. While we do not suggest that SWB and happiness are meaningless concepts, this view posits that their meaning is distinct from economists' conception of utility—which is merely a representation of choice. In standard microeconomic theory, an outcome yields higher utility if and only if it is preferred.

The debate over whether individuals either do or should maximize happiness is ancient and intellectually rich. Bentham (1789) famously wrote: "It is for [*pain* and *pleasure*] alone to point out what we ought to do, as well as to determine what we shall do." Bentham's claim that happiness is both the positive and normative determinant of behavior greatly influenced nineteenth-century economists, such as John Stuart Mill.

These economists reflect an ancient philosophical tradition, dating back to the Cyrenaic School. The Cyrenaics emphasized pleasure as life's central goal and influenced the later Epicureans. Giants of medieval philosophy, including Augustine and Aquinas, accepted that human beings pursued happiness above all, but they taught that true happiness is to be found by following God's will.

Some modern researchers on happiness, though certainly not all, have also conflated happiness with utility (Alesina et al. 2004), or at least social welfare (Easterlin 1995).[8] Yet it is quite possible to believe that happiness is interesting and important without accepting the equivalence. There is also an equally ancient and distinguished philosophical tradition rejecting the notion that individuals either do or should maximize happiness.

While the Epicureans believed in maximizing pleasure and minimizing pain, the Stoics did not. About 1900 years ago, Epictetus wrote: "What is our nature? To be free, noble, self-respecting. . . . We must subordinate pleasure to these principles" (Epictetus 1916, 89). Epictetus is making the normative claim that other goals—freedom, nobility, and self-respect—trump happiness.

Economists from Fisher (1892) through Stigler (1950) and Becker and Rayo (2008) have followed this approach. Becker and Rayo (2008, 89) wrote of "an alternative interpretation of the happiness data, namely, that happiness is a commodity in the utility function in the same way that owning a car and being healthy are." Perhaps the best known evidence supporting this interpretation is that parents of small children typically report lower happiness or life satisfaction (Baumeister 1991).[9] If happiness measured utility, then presumably this relationship should impede the survival of the species. Yet in Becker and Rayo's formulation, this negative relationship is no puzzle at all—parents receive ample compensation, in the form of progeny, for their suffering.

Leaving aside survey evidence on SWB, we can look for reasonably direct evidence of low utility levels by measuring suicide (Becker and Posner 2004). If suicide reflects low utility, then the relationship between suicide and self-reported happiness provides another way to consider

[8] Alesina et al. (2004) explicitly state that they "measure 'utility' in terms of survey answers about 'happiness'." They elaborate in footnote 7 that they, and in their view much of the literature on economics of happiness, aim to measure "*experienced* utility, a concept that emphasizes the pleasures derived from consumption." They view these survey responses, in certain circumstances, as "reasonable substitutes to observing individual choices." Many other prominent papers in this literature implicitly posit such an equivalence, such as Easterlin (1995, 36). Easterlin writes, "Formally, this model corresponds to a model of interdependent preferences in which each individual's utility or subjective well-being varies directly with his or her own income and inversely with the average income of others."

Of course the literature has also considered many subtle points about the appropriate conception of subjective well-being. For example, Kahneman and Thaler (1991) and Kahneman and Krueger (2006) distinguish between decision utility and experienced utility. We certainly do not claim to introduce a novel distinction here. Our contribution, in part, is to use the decision-utility maximization embodied in spatial equilibrium to put more structure on the theoretical and empirical relationships between choices and subjective well-being.

[9] But see Deaton and Stone (2014) on the importance of sorting into parenthood.

whether low SWB captures low utility. In the working paper version of the current paper (Glaeser et al. 2014), we show that suicide rates are uncorrelated with subjective well-being across metropolitan areas, corroborating Daly et al.'s (2011) cross-state evidence.

Some suggest that, in its very wording, life satisfaction should capture all the elements of utility. While it seems implausible to hope that maximizing utility should automatically mean maximizing joy or happiness, might individuals answer the question about life satisfaction in such a way that actually ranks their preferred outcomes, as does a utility function? If so, an individual who has received a preferred outcome will report a higher level of life satisfaction. Hence, utility and subjective well-being converge because well-being acts as a barometer to measure how well people have achieved their goals.

Yet this view seems barely more tenable than the view that happiness should miraculously map onto human preferences. Someone may choose a more competitive environment with more opportunity to shape the world and yet know that this environment will—by opening up opportunities and inviting comparisons with high achievers—lead to less satisfaction. A rational person could select a PhD program, or a city, despite recognizing that it will lead to less satisfaction.

Among members of the classical economic tradition, Bernard de Mandeville may be the most powerful proponent of the view that human beings should not maximize happiness, especially not in location choice. In *The Fable of the Bees* (1714), he writes, "To be happy is to be pleas'd, and the less Notion a Man has of a better way of Living, the more content he'll be with his own . . . the greater a Man's Knowledge and Experience is in the World, the more exquisite the Delicacy of his Taste, and the more consummate Judge he is of things in general, certainly the more difficult it will be to please him. . . . But when a Man enjoys himself, Laughs and Sings, and in his Gesture and Behaviour shews me all the tokens of Content and Satisfaction, I pronounce him happy, and have nothing to do with his Wit or Capacity." Clearly, de Mandeville thinks little of happiness. When he writes "ask'd where I thought it was most probable that Men might enjoy true Happiness, I would prefer a small peaceable Society, in which Men, neither envy'd nor esteem'd by Neighbours, should be contented to live upon the Natural Product of the Spot they inhabit, to a vast Multitude abounding in Wealth and Power," he is not espousing such places, but arguing that it is perfectly sensible to choose busier, but less happy, locales.

We do not dispute the desirability of happiness, and for that reason, the spatial equilibrium logic of Rosen (1979) and Roback (1982) implies that there must be some compensation offsetting the unhappiness of declining cities. Residents must receive some other benefit, such as higher real wages, that offsets the costs of lower life satisfaction. Otherwise, it would be hard to understand why they remain in unhappy cities. We formalize these issues in the model that follows.

B. Happiness and Utility

To formalize this discussion, we begin with a general framework meant to capture the difference between happiness and utility. We then adapt our structure to deal with cities and urban decline, which requires considerably more assumptions about structure and ultimately even functional forms. This latter section puts forward the model that will be taken to the data in Section IV.

In Becker (1965), individuals maximize a function $U(\cdot)$ defined over a vector of objectives $\tilde{Z}$, where each element in that vector Z_i is a function of time (t_i) and spending (s_i). One possible approach is to assume that life satisfaction is defined over an alternative function $H(\cdot)$ of those same objectives, but that approach provides little guidance for modeling or testable implications.

We assume that subjective well-being represents an alternative function $W(\cdot)$ over the same set of objectives. It may be that welfare is a function of well-being and other objectives, or that well-being is simply a slightly different function of exactly the same inputs that guide utility. In the first case, utility can be described as $U(W(\tilde{Z}), \tilde{Z}_{\mathrm{NH}})$, where $\tilde{Z}_{\mathrm{NH}}$ refers to objectives that enter into utility directly, such as childrearing, as well as possibly also having an impact on well-being.

We will approach well-being and utility as reflecting a combination of experiences and achievements. Well-being or happiness will be conceived as experience-based utility, following Bentham (1789/1996) and Kahneman and Krueger (2006). Individuals care about experienced utility, but they also care about achievements, which can also be produced with time and money. We lose little generality by assuming at this point that there is a single achievement, which is produced with achievement-specific time denoted t_A and achievement-specific spending s_A. Individual earnings are the product of wages and time spent working wt_w and unearned income y_0, which includes the fixed cost of housing.

Time spent working and time pursuing the alternative achievement convey experienced utility of t_w and t_A per unit time. The remainder of hedonic time generates well-being equal to $h(s_h)t_h$, where s_h reflects the total amount of spending on these activities. The term $h(s_h)t_h$ is meant to aggregate all other time, including sleeping. So the individual's problem is to maximize

$$U(h_w t_w + h_a t_a + h(s_h)t_h,\ Z(s_a,\ t_a)), \tag{5}$$

subject to the time budget constraint $t_w + t_a + t_h = 1$ (where total time available is normalized to one) and the cash budget constraint $wt_w + y_0 = s_h + s_a$.

The two budget constraints can be combined to create a single total budget constraint of $w + y_0 = w(t_a + t_h) + s_a + s_h$.

In this model, as in almost all economic models, more income is preferred to less and translates into higher levels of utility. Yet the link between happiness and wages is less clear. If, for example, $Z(\cdot)$ is produced entirely with earnings, then as long as the uncompensated wage elasticity is positive, happiness diminishes with wages even though utility increases. If $h(s_h) = h_0$ is independent of income, then the derivative of happiness with respect to the wage equals $(h_0 - h_w)$ times the derivative of t_h with respect to the wage, which equals

$$\frac{\partial t_h}{\partial w} = \frac{-wZ'(s_a)U_Z + (1 - t_h)Z'(s_a)[(h_0 - h_w)U_{HZ} - wZ'(s_a)U_{ZZ}]}{-\left[U_{HH}(h_0 - h_w)^2 - 2(h_0 - h_w)wZ'(s_a)U_{HZ} + (wZ'(s_a))^2U_{ZZ}\right]}. \tag{6}$$

Across space, the impact of income on happiness may be even more negative. Suppose that amenities are constant across space and that utility levels are unchanged with changes in wages; $\partial y_0/\partial w = -(1 - It_0)$: the change in housing costs exactly offsets the change in earnings. If this is the case, then in the case discussed above, where spending does not impact the hedonic flow of time,

$$\frac{\partial t_h}{\partial w} = \frac{-wZ'(s_a)U_Z}{-\left[U_{HH}(h_0 - h_w)^2 - 2(h_0 - h_w)wZ'(s_a)U_{HZ} + (wZ'(s_a))^2U_{ZZ}\right]} < 0,$$

so happiness is always lower in higher-wage cities. Since the impact of area-level wages is compensated, rather than an uncompensated change in wages, it will invariably cause an increase in hours worked and a decrease in time spent in household production.

We now turn to the spatial equilibrium, where we assume that $h_w = h_a = 0$. We also assume a Cobb-Douglas utility function, with a weight of α on happiness, power functions for producing the other goods, and that $Z(s_a, t_a) = z_0(s_a)^z(t_a)^{1-z}$, where z_0 is a city-specific production shifter. We assume that time spent at work is fixed at $\hat{t}_w$ but that time can still be allocated between leisure and the other achievement. Further, $h(w\hat{t}_w + y_0) = h_0(w\hat{t}_w + y_0)^\gamma$, where h_0 is a city-specific amenity. Given these assumptions, then indirect utility is proportional to $(h_0)^\alpha(z_0)^{1-\alpha}(w\hat{t}_w + y_0)^{\alpha\gamma+(1-\alpha)z}$ and happiness is proportional to $h_0(w\hat{t}_w + y_0)^\gamma$.

The Cobb-Douglas welfare function generates a happiness-income tradeoff of log (Happiness)$/d$log(Income) $= -(1 - \alpha)/\alpha$. This tradeoff is a distinct concept from the derivative of happiness with respect to the wage (assuming unearned income is negligible), which equals γ.

We have two options here, choosing fixed or flexible working time, but the simpler functional forms come with fixed hours. In that case, the spatial equilibrium condition can be written as:

$$w\hat{t}_w + y_0 = k_0 h_0^{-[\alpha/[\alpha\gamma+(1-\alpha)z]} z_0^{-(1-\alpha)/[\alpha\gamma+(1-\alpha)z]}. \tag{7}$$

The values of h_0 and z_0 are determined both by natural amenities, such as climate, and amenities tied to public services, such as safety. Declining areas could well have lower levels of quality of life both because they are in relatively cold areas of the United States and because a reduced level of spending leads to lower levels of public amenities.

An urban equilibrium involves three separate equations. The first is the spatial equilibrium curve for consumers in which welfare—but not happiness—must equal a constant reservation utility across space. The second condition is that firm profits are equalized across space. The third condition is that the cost of housing equals the cost of supplying homes.

We assume a linear housing supply curve, so that the flow cost of housing in a city, denoted r, is $r = c_0 + c_1\log(N_t) + c_2\log(N_t/N_{t-1})$, where N_t reflects the population in the place, and we assume that $y_0 = -r$. This can be generated by an assumption that houses are created with a Cobb-Douglas utility function using traded and nontraded capital, where nontraded capital is in fixed supply. In principle, c_0, c_1, and c_2 might all vary across areas.

Finally, we have linear labor demand, so that $w\hat{t}_w = A - B\log(N_t)$. This can be generated by assuming that there are a fixed number of firms with Cobb-Douglas production functions and two types of labor, one of which is traded and the other which is not (Glaeser 2007). Again, A and B might differ across metropolitan areas.

Using the housing supply curve and the labor demand curve and taking logs of equation (7) yields

$$\log(N_t) = \frac{1}{B + c_1 + c_2}\Big\{A - c_0 + c_2\log(N_{t-1}) - k_0 h_0^{-\alpha/[\alpha\gamma+(1-\alpha)z]} z_0^{-(1-\alpha)/[\alpha\gamma+(1-\alpha)z]}\Big\}, \tag{8}$$

$$w\hat{t}_w = \frac{1}{B + c_1 + c_2}\Big\{(c_1 + c_2)A + Bc_0 - \mathrm{B}c_2\log(N_{t-1}) + Bk_0 h_0^{-\alpha\gamma/[\alpha\gamma+(1-\alpha)z]} z_0^{-(1-\alpha)/[\alpha\gamma+(1-\alpha)z]}\Big\}, \tag{9}$$

$$r = \frac{1}{B + c_1 + c_2}\Big\{(c_1 + c_2)\big(A - k_0 h_0^{-\alpha/[\alpha\gamma+(1-\alpha)z]} z_0^{-(1-\alpha)/[\alpha\gamma+(1-\alpha)z]}\big) + B[c_0 - c_2\log(N_{t-1})]\Big\}, \tag{10}$$

$$\log(\text{Happiness}) = \log[k_0^{\gamma}(1-\hat{t}_w)] + \frac{(1-\alpha)z}{\alpha\gamma + (1-\alpha)z}\log(h_0) - \frac{(1-\alpha)\gamma}{\alpha\gamma+(1-\alpha)z}\log(z_0). \tag{11}$$

Population is increasing with productivity, decreasing with the cost of providing housing, and increasing with the two amenity variables. Income is rising with productivity and housing cost and falling with amenities. Housing rents increase with productivity, with the cost of housing and with amenities. Happiness is rising with the happiness-related amenity and declining with the non-happiness-related amenity. This becomes a four-equation system for empirical work, where the impacts of local variables can be traced through these four distinct outcomes.

In this formulation, happiness is a measure of local amenities—and local amenities only—because population and housing prices adjust to shifts in local demand and construction costs. The spatial equilibrium requires that gaps in real income end up being proportional to happiness, holding as such happiness should be declining in real income. The slope is predicted to equal $[(1-\alpha)z]/\alpha$ on real income, which equals $(1-\alpha)/\alpha$—the basic happiness-income tradeoff—times z—the elasticity of the non-happiness-related component of welfare with respect to earnings.

We can also use the spatial indifference condition and find that happiness is proportional to $(z_0)^{(\alpha-1)/\alpha}(w\hat{t}_w + y_0)^{[(\alpha-1)/\alpha]z}$. Holding z_0 constant, we expect to find that richer places are less happy, and holding income constant, we expect to find happier places deficient in some other desirable (non-happiness related) amenity. The unhappiness of declining cities, therefore, needs to be compensated either with higher real incomes or with some other asset.

IV. Are Individuals Compensated for Unhappiness?

In this section, we do not implement the four-equation empirical estimation section discussed above but rather restrict ourselves to a simpler empirical approach inspired by the theory. We will test whether individuals in declining or unhappy cities are being compensated for their misery by either lower housing costs or higher wages.

This approach begins with a simple view of America's changing urban system. We initially built cities in places where firms had a productive advantage because of proximity to waterways and coal mines. Moreover, we also built those cities in ways that favored productivity rather than pleasure. Over time, declining transport costs enabled capital and labor to flee low-amenity places (Glaeser and Kohlhase 2004) and move to "consumer cities" endowed with higher amenity levels (Glaeser, Kolko, and Saiz 2001). An increasingly wealthy population also built new cities that were more oriented toward consumer well-being.

Within the context of the model, this can be understood as a change in the covariance between productivity and the amenity parameters. In early twentieth-century America, productivity may have been higher in lower amenity places but in late twentieth-century America, that negative covariance disappeared. As a result, population growth was faster in places that had higher amenities initially and lower levels of productivity.

This argument provides a slightly different interpretation of Easterlin (1974), at least insofar as it applies to America's metropolitan areas. In the early twentieth century, a city needed to be unpleasant to be productive. In the late twentieth century, it did not. Since technological change favored pleasant, happier locales, it seemed as if happiness was tied to income growth, even if it was ultimately driven by the local environment.

We now turn to the question of compensation. Since no one would presumably have built an intrinsically unhappy city unless it was more productive, we first look at income in 1940. We test whether declining cities, which seem also to have been unhappy in the past, paid higher wages in 1940. Columns 1 and 2 of table 6 show these results (which we expand on in Glaeser et al. 2014). In both columns, we look at earnings for males aged 25–55. We include a full battery of controls for age, race, and education. Column 1 shows that as population growth increases by 0.1 log point, for cities with population growth below the median, wages drop by 0.014 log points. In 1940, the cities that would subsequently decline were very well paid. We do not mean to imply causality with this regression. The wage outcomes precede the decline and may have caused the decline. We only mean to suggest that residents of cities that declined after 1950, and that are unhappy today, were relatively well compensated in 1940.

The second column repeats this exercise controlling (coarsely) for city size today. In this column, we simply include an indicator variable that takes on a value of one if the population of the metropolitan area is greater than 5 million. The coefficient on population is quite large, and it causes the coefficient on population growth below the median to decline to −0.1. The strong positive coefficient on large population size is also quite compatible with the compensation hypothesis, for the Gallup data show that people who live in extremely large cities were dramatically less happy in the 1940s (Glaeser et al. 2014).

We next look at housing costs in that year to test whether either wages or rents compensate the residents in unhappy declining cities. Columns 3 and 4 show these results. With both measures, cities with subsequent decline had higher housing costs in 1940. These higher costs would mean lower real incomes in these areas, which should eat away some of the compensation received for living in less happy places. However, we are

Table 6
Relationships among Income, Housing Costs, Population Growth, and Happiness from 1940 and 2000 Censuses

	Dependent Variable									
	Regressions from 1940 Census				Regressions from 2000 Census					
	Log Income		Log House Value	Log Housing Rent	Log Income	Log House Value	Log Housing Rent	Log Income	Log House Value	Log Housing Rent
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Population growth, 1950–2000, below median	−.144**	−.104***	−.219*	−.294***	.0109	.302**	.208**			
	(.0577)	(.0324)	−.128	(.103)	(.0533)	(.132)	(.0807)			
Population growth, 1950–2000, above median	−.0442	−.0261	−.157	−.0619	.00640	.0953	.171***			
	(.0455)	(.0404)	(.124)	(.119)	(.0196)	(.0579)	(.0467)			
Indicator for MSA with population ≥ 5 million		.0671***	.296***	.376***	.0900***	.555***	.387***	.0950***	.489***	.338***
		(.0211)	(.0656)	(.0541)	(.0310)	(.123)	(.0514)	(.0284)	(.134)	(.0617)
MSA-adjusted life satisfaction								.154	.157	.659*
								(.197)	(.387)	(.335)
No. of observations (1,000s)	7,766	7,766	6,535	10,527	30,044	42,914	22,125	30,044	42,914	22,125
R^2	.211	.215	.035	.095	.217	.411	.108	.217	.400	.089

Source.—Authors' regressions on US Census microdata (Ruggles et al. 2010).

Note.—Rent variable is monthly contract rent. The 1940 Census data do not include housing quality information. The sample for income regressions is employed males aged 25–55 who work full-time and earn more than half the federal minimum wage for a full-time worker. All income regressions include age, race, and education controls. All regressions include household or person weights. Standard errors are clustered at the metropolitan statistical area (MSA) level.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

wary of putting much weight on these regressions, since we are not able to effectively control for housing quality in the 1940 US Census data. As such, the results could just be showing that the residents of industrial cities in 1940 had substantially better housing than the residents of the Sunbelt in those years.

Column 5 shows the result for income in the 2000 US Census microdata. Declining metropolitan areas are not particularly well or poorly paid relative to the United States overall. Columns 6 and 7 show results for housing costs. The house value regression estimates a coefficient of 0.3 on population growth below the median, implying that a 0.1 log decrease in population growth between 1950 and 2000 is associated with housing values that are 0.03 log points lower. The result is similar in column 7 when using rents. The residents of declining cities may be less happy, but they are being at least modestly compensated for lower levels of happiness with lower housing costs.

The last three columns in table 6 look at whether rents or incomes seem to be directly compensating for happiness in 2000. We now correlate the Census outcomes with area happiness itself. We have again controlled for individual attributes, but distinct issues with this specification remain. If the spatial equilibrium is imperfect, then a temporary shock to local income should make people happier—not unhappy. As such, we may not see the negative relationship between happiness and real incomes that is predicted by the model.

Column 8 shows that the happiness variable is not particularly correlated with area income.[10] One interpretation of this is that, while people need higher real incomes to put up with unhappy places, these higher incomes also make them happier. These two offsetting forces could create the near-zero coefficient. The remaining regressions examine housing costs, and the results are positive although estimated imprecisely. The coefficient relating rents to happiness in column 10 is 0.66. A 0.1 increase in area happiness is associated with a 0.07 log point increase in rents.

This relationship is shown in figure 4, which collapses the data to the metropolitan area level, but excludes California (which has much higher rents). The positive relationship is generally visible, but a small number of cities with higher rents and lower levels of self-reported happiness can be seen in the upper-left-hand corner of the graph. These are the large metropolitan areas of the east coast, such as New York and Boston. These places also tend to pay high wages, which is presumably how their residents are compensated for lower levels of happiness.

[10] Glaeser and Gottlieb (2008) show that happiness and income are also uncorrelated across cities in the GSS.

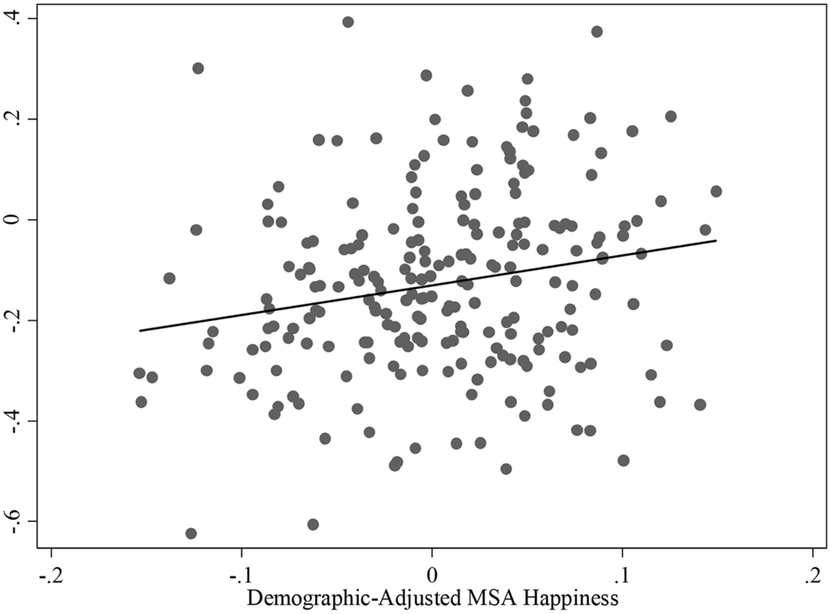

FIG. 4.—Metropolitan statistical area rent and adjusted happiness. This figure shows each metropolitan and rural area's adjusted life satisfaction, after controlling for demographic covariates in a mixed effects model, against adjusted housing rent from the Census (the median of each metropolitan statistical area's residuals from regressing rent on housing characteristics). Data are from Centers for Disease Control and Prevention (2005–10) and Ruggles et al. (2010).

In sum, this table suggests that higher wages compensated for the unhappiness of cities that were large and productive in the 1940s but that would subsequently decline. The population decline has not offset the unhappiness but is associated with lower housing costs that could partly compensate for the lower reported well-being in such places. This tradeoff is consistent with a model in which happiness is one argument in utility. It is harder to reconcile with views that emphasize happiness as equivalent to utility, or as individuals' ultimate objective.

V. Conclusion

In this paper, we have documented significant differences in self-reported well-being across American cities that persist when we control for individual demographics and even for individual fixed effects. These facts are not reliably correlated with many area level attributes, but they do seem to be connected with urban decline across at least three large data sets. We do not interpret this correlation as evidence that population decline causes unhappiness. Indeed, cities that have declined also seem to have been unhappy in the past. This suggests that a better interpretation might be that these areas were always unhappy—and that was one reason why they declined.

Differences in happiness or subjective well-being across space weakly support the view that these desires do not uniquely drive human ambitions. If we chose only that which maximized our happiness, then individuals would presumably move to happier places until the point where rising rents and congestion eliminated the joys of that locale. An alternative view is that humans are quite understandably willing to sacrifice happiness or life satisfaction if the price is right. This view rationalizes the well-known tendency of parents to report lower levels of happiness and life satisfaction. Indeed, the residents of unhappier metropolitan areas today receive higher real wages—presumably as compensation for their misery.

Declining cities seem also to have been unhappy during the past, but in 1940, the cities that were prone to future decline earned outsized incomes, both nominal and real. The industrial cities of the Midwest may have reported lower happiness levels, but their residents were getting richer as a result. As transportation cost declines freed industry from the Great Lakes and the coal mines, we should not be surprised that people left less pleasant locations. Today the residents of cities that declined are not receiving higher nominal wages, but they do seem to be paying lower rents. As such, the unhappiness of America's declining cities may have been compensated with higher incomes in the past and lower housing costs today.

Appendix

BRFSS

Throughout this paper, we follow the literature in measuring happiness using self-reported survey data on subjective well-being (SWB). We use a large national survey, the Behavioral Risk Factor Surveillance System (BRFSS), conducted by the Centers for Disease Control and Prevention (CDC), which asks individuals to report on their own life satisfaction using a discrete response scale.

The CDC (2005–10) has conducted BRFSS surveys annually since 1985, in order to study risk factors for various diseases. This is a large, nationally representative survey, involving more than 350,000 respondents in over 2,000 counties annually.

The BRFSS survey is administered by individual states via telephone interviews. The interviews are collected via computer-assisted phone calls to randomly selected landlines.[11] During our sample period of 2005–9, the

[11] CDC provides weights to adjust for differences in phone line density across areas, but we do not use these weights.

survey covers all 50 states and Washington, DC.[12] Individuals report their county to the interviewer, and we drop observations where the county is not reported.

Based on the self-reported county, respondents live in 367 metropolitan statistical areas (MSAs) and nonmetropolitan regions.[13] When we examine temporal patterns in the data, we restrict the sample to the 177 MSAs with at least 200 respondents in each year.

The life satisfaction question we use has been a part of the BRFSS "core" since 2005. Core questions are asked in every interview with minor exceptions. In 2009, the life satisfaction question was not asked in less than 5% of BRFSS surveys, which is approximately the same percent unasked of similar questions in the survey. This number is slightly lower in other years. Responses to LSATISFY of "refused" and "unsure" are treated as missing responses and dropped from the data set.

One might be concerned that individual SWB is reported on a discrete scale, with values whose interpretation is not obvious. When we summarize one area's happiness as a linear average of these discrete values, the resulting summary is undoubtedly a noisy and imperfect measure of area-level happiness. We cannot solve this problem, but Stevenson and Wolfers (2008) find that more sophisticated methods yield results that are extremely highly correlated across countries (correlations are regularly above 0.99) with results from this method.

We standardize each year's data separately, with respect to the overall mean and standard deviation for the survey year in question.

One wave of the BRFSS may actually be administered in two different years (e.g., the 2009 wave interviewed respondents from January 2009 through January 2010). The year fixed effects γ_t that we estimate represent the survey wave as opposed to the actual year of the interview.

The concern about systematic differences in individual SWB is not merely hypothetical. On the contrary, a large body of research has documented regular patterns based on age, sex, income, life events, and other demographic characteristics.[14] To the extent that people sort across areas

[12] Puerto Rico, Guam, and the US Virgin Islands are also included, but we drop the three territories.

[13] We use the county FIPS code to assign the respondent to a metropolitan area. We use the Office of Management and Budget's definitions of metropolitan areas from 1999 (which correspond to data from the 2000 Census). We use Primary Metropolitan Statistical Areas (PMSAs) rather than Consolidated Metropolitan Statistical Areas (CMSAs), where applicable. We classify respondents in New England according to their New England Consolidated Metropolitan Statistical Area (NECMA) rather than PMSA or CMSA. We classify all respondents not living in an MSA, PMSA, or NECMA as part of one "nonmetropolitan region" for their state (e.g., "nonmetropolitan Texas").

[14] See, e.g., Sacks, Stevenson, and Wolfers (2010).

based on these same characteristics, our estimates of area-level happiness will be biased.

A small percentage of survey respondents refuse to respond to one or more of the demographic questions asked. The total fraction refusing to answer, unsure of, or not being asked at least one demographic question of interest is about 2.3% in any year. We drop any observation with any such missing demographic information, as well as respondents over 85 years old.

The controls for children's characteristics deserve further elaboration. While the survey nearly always has information about the number of children in the household, more detailed information is available for only one randomly selected child. In most states during most years, the BRFSS asks about the age of one randomly selected child in the household, as well as the respondent's relationship to that child.[15] We therefore create indicator variables for four age ranges of the randomly selected child and six categories for the respondent's relationship.[16] The omitted group for these questions is respondents with no children. We add a separate dummy variable indicating respondents with children in state-years when no question was asked about a child's age.

Table A2 reports the coefficients on the controls in this regression when it was run on our full sample of 1,574,361 respondents across five waves of the BRFSS. For the most part, these coefficients are consistent with findings in the previous literature, and they are robust to the inclusion or exclusion of area fixed effects. In column 1, we include only the basic demographic controls discussed above. We find that age has an important influence on subjective well-being, as estimated by a fifth-order polynomial in age. On average, men are 0.036 standard deviations less happy than women. There are strongly significant differences across races, with whites reporting the highest average well-being.

The most significant correlates of happiness in column 1 are education level and marital status. Education has one of the largest impacts on individual responses, with a range of nearly half a standard deviation from high school dropouts to college graduates. But bear in mind that this regression does not control for individual-level income, which may mediate this relationship somewhat. Marital status is also extremely important, with married individuals half a standard deviation happier than single or divorced

[15] The survey is divided into core questions and modules, the latter of which each state individually elects whether to ask in their phone interviews. Individual states sometimes add additional questions on their own. None of the questions we focus on are module or state questions in any year, except for the age of one randomly selected child.

[16] In the 2006 survey, the age of the child is not recorded, but it is imputed from the reported birthdate. In 2007, the age is recorded in the BRFSS in months, and we round this down to an integer number of years.

respondents. Those reporting being separated are one-sixth of a standard deviation less happy than singles or divorcees.

Our estimates of the relationship between happiness and the presence of children in the household differ from previous findings. The existing literature has generally found a significant negative association between happiness and having children, especially young children.[17] In the BRFSS data, however, there seems to be a more complex relationship. This regression allows us to compute the connection between a respondent's subjective well-being and the presence of children with various characteristics in the household. To calculate the complete relationship, we need to add the coefficients for the appropriate number of children (one, two, three or more), the age of the randomly selected child (one of four categories, or unknown), and the respondent's relationship to the randomly selected child. For all of these characteristics, the coefficients presented in table A2 are expressed relative to the omitted group of respondents with no children in the household.

Parents in a one-child household are, on average, anywhere from 0.01 standard deviations less happy than similar respondents with no child to 0.07 standard deviations happier, depending on the child's age. Older children appear to be associated with less happiness, all else equal, with 11–17-year-olds having a coefficient 0.076 standard deviations below 0–1-year-olds. We find increasingly positive coefficients as the number of children increases, with a bump of 0.04 standard deviations for a second child and a further 0.01 standard deviation gain with a third child or beyond.

These benign or positive relationships between children and happiness disappear if the respondent is the child's guardian but not the biological parent. Grandparents, foster parents, and unspecified other relatives have very strong negative coefficients, which wipe out the (otherwise positive) associations with most categories of number and age of children. In other specifications (not reported), we interact the number or age of children with the respondent's marital status or relationship with the random child. These regressions tend to confirm that the positive correlation between children and respondents' well-being is concentrated among married couples and respondents who are the child's biological parents, while the other

[17] The negative relationship between children—especially young children—and parents' happiness is widely accepted in the literature. Di Tella, MacCulloch, and Oswald (2001) report increasingly negative coefficients on life satisfaction in the EuroBarometer as the number of children increases (table A1). This finding dates back at least to Glenn and Weaver (1979), which finds the negative coefficient to be largest for children under 5 years old in the General Social Survey (table 1). The closest finding to ours is Clark and Oswald (1994), which estimates a negative effect of having one child relative to no children and insignificant negative effects of having two or more children compared with none (tables 2 and 3). It does not report results controlling for children's ages.

groups tend to have negative associations between the presence of children and their own well-being.

Even without these interactions, our data suggest a more complex relationship than that previously found between subjective well-being and the presence and age of children. These correlations are sensitive to the relationship between the children present and the individual in question. Nevertheless, it is unlikely that the inclusion of controls for relationship with the child fully explains the difference between our results and the negative coefficients on children's presence reported in other papers. The cases of nonparental relationship status are probably not sufficiently prevalent to explain the aggregate negative associations found in other data sets.

Table A1
Distribution of Responses to Life Satisfaction Questions

Life Satisfaction Question	Number of Respondents
A. BRFSS life satisfaction question, 2005–9:	
Question: "In general, how satisfied are you with your life?"	
Answer:	
Very satisfied	717,779
Somewhat satisfied	766,374
Somewhat unsatisfied	72,258
Very unsatisfied	17,950
Total sample size:	1,574,361
B. NSFH life satisfaction question, wave 1 (1987–88):	
Question: "First taking things all together, how would you say things are these days?"	
Answer:	
1–Very unhappy	244
2	206
3	522
4	1,894
5	2,667
6	3,073
7–Very happy	2,723
Total sample size	11,329
C. NSFH life satisfaction question, wave 2 (1992–94):	
Question: "First taking things all together, how would you say things are these days?"	
Answer:	
1–Very unhappy	153
2	145
3	438
4	1,271

Table A1 (*Continued*)

Life Satisfaction Question	Number of Respondents
5	2,253
6	2,370
7–Very happy	1,874
Total sample size	8,504

SOURCES.—Panel A: Behavioral Risk Factor Surveillance System Survey (Centers for Disease Control and Prevention 2005–9); panel B: Sweet, Bumpass and Call (1988); panel C: Sweet and Bumpass (1996).

Table A2
Coefficients on Demographic Characteristics in Life Satisfaction Regression

	(1)		(2)		(3)	
Variable	Estimate	SE	Estimate	SE	Estimate	SE
Age/10	−2.061	.149	−.995	.185	.664	.183
Age^2/100	.761	.066	.266	.081	−.419	.080
Age^3/1,000	−.150	.014	−.038	.017	.092	.017
Age^4/10,000	.015	.001	.0034	.0017	−.008	.002
Age^5/100,000	−.0006	.0001	−.0002	.0001	.0002	.0001
Male	−.036	.002	−.037	.002	−.060	.002
Black	−.025	.003	.010	.004	.071	.004
Asian	−.124	.006	−.130	.008	−.093	.007
Pacific Islander	−.016	.017	.011	.020	.053	.020
Native American	−.079	.007	−.017	.008	.034	.008
Other race, non-Hispanic	−.119	.010	−.093	.012	−.050	.012
Multiple races	−.145	.006	−.103	.007	−.070	.007
Hispanic	−.014	.004	−.012	.004	.071	.004
Some high school	−.176	.003	−.101	.004	−.040	.004
Some college	.072	.002	.054	.002	.002	.002
College graduate	.273	.002	.229	.002	.096	.003
Married	.457	.003	.406	.003	.266	.003
Divorced	.003	.003	.005	.004	.005	.004
Separated	−.175	.006	−.141	.007	−.136	.007
In unmarried couple	.166	.005	.143	.006	.082	.006
One child < age 18 in household	.016	.006	−.002	.007	.004	.007
Two children < age 18 in household	.057	.006	.032	.007	.035	.007
Three or more children < age 18 in household	.067	.007	.041	.008	.054	.007
Random child < 2 years old	.052	.009	.069	.010	.074	.010
Random child 2–4 years old	−.020	.008	−.017	.010	−.019	.010
Random child 5–10 years old	−.021	.008	−.025	.009	−.031	.009
Random child 11–17 years old	−.024	.007	−.032	.009	−.041	.008
Random child's age not asked	−.028	.006	−.025	.007	−.028	.007
Respondent is random child's parent	−.001	.007	.001	.008	.004	.008
Respondent is random child's grandparent	−.158	.010	−.109	.013	−.101	.012
Respondent is random child's foster parent	−.057	.017	−.030	.017	−.034	.017

Table A2 (*Continued*)

Variable	(1) Estimate	(1) SE	(2) Estimate	(2) SE	(3) Estimate	(3) SE
Respondent is random child's sibling	.072	.013	.098	.017	.013	.017
Respondent is random child's other relative	−.077	.016	−.039	.019	−.049	.019
Self-employed			.036	.003	.046	.032
Unemployed for more than 1 year			−.574	.007	−.416	.007
Unemployed for less than 1 year			−.436	.006	−.323	.006
Homemaker			−.005	.004	.031	.004
Student			−.028	.007	.049	.007
Retired			−.019	.003	.040	.003
Unable to work			−.717	.004	−.543	.004
Income $10,000–$15,000					.047	.006
Income $15,000–$20,000					.117	.005
Income $20,000–$25,000					.169	.005
Income $25,000–$35,000					.247	.005
Income $35,000–$50,000					.346	.005
Income $50,000–$75,000					.454	.005
Income > $75,000					.615	.005
Metropolitan and nonmetropolitan area fixed effects	Yes		Yes		Yes	
R^2	.076		.11		.13	
Sample size	1,574,361		1,084,596		1,084,596	

NOTE.—The table presents the results of the linear regression of individual responses to life satisfaction question in the Behavioral Risk Factor Surveillance System Survey (Centers for Disease Control and Prevention 2005–10), against the variables shown, month dummies, BRFSS wave fixed effects and dummies for 367 metropolitan statistical areas and nonmetropolitan regions. The omitted category of respondent is a single white female with a high school education and no children in the household, and in cols. 2 and 3, employed in the marketplace with income less than $10,000 per year.

Subsequent regressions in table A2 add controls for the respondent's economic situation. In column 2, we add dummies for labor force status. With employed individuals as the omitted group, we find that self-employment is associated with a 0.036 standard deviations more well-being, while the unemployed are 0.44–0.57 standard deviations less happy than employed workers. Retirees are 0.02 standard deviations less happy than workers, controlling for age, and those unable to work are 0.7 standard deviations less happy than workers. Including labor force status controls has only a modest impact on the coefficients on other demographics, with the notable exception of the indicator for being black. This dummy reverses signs, from −0.025 in column 1 to 0.01 in column 2.

Column 3 adds controls for reported income categories in addition to of the previous characteristics. These dummies show that happiness increases monotonically in income, with a range of 0.6 from the omitted category (less than $10,000 per year) to the highest income category (above $75,000 per year). Because income is correlated with many of the other covariates, its inclusion dramatically shifts some of the coefficients on other variables, including education, unemployment, race, and marital status, relative to their levels in column 2.

Aggregate Data

Our aggregate data about the metropolitan and nonmetropolitan areas in the country come from various sources. These data mostly come from the National Historical Geographic Information System (Minnesota Population Center 2004), which compiles data from the US Census. We obtain these data at the county level and consolidate them using the same metropolitan area definitions from 1999 as we use for the BRFSS. We obtain a number of quality of life measurements from Albouy (2008) and geographic data from Rappaport and Sachs (2003).

Data on Movers from the National Survey of Families and Households

The NSFH is a probability sample survey of 13,017 respondents in 9,643 households plus an oversampling of minority and single-family households and households with stepchildren. It is a longitudinal study with three waves, the first between 1987 and 1988, the second between 1992 and 1994, and the third wave between 2001 and 2002 (Sweet, Bumpass, and Call 1988; Sweet and Bumpass 1996; Trull and Famularo 1996).

We use data from the first two waves of the NSFH. In both waves, the data contain information on family and personal characteristics of individuals and on individual subjective well-being. In particular, the NSFH asks: "First taking things all together, how would you say things are these days?" Respondents may choose to respond on a 1–7 scale, 1 being very unhappy and 7 being very happy. The summary statistics from this survey are shown in Panels B and C of table A1. For our regressions, we normalize the responses to have mean zero and unit variance.

To examine the relationship between changes in subjective well-being and changes in geographic location, we need to match the longitudinal NSFH data to geographic data. Because the geographic locations of survey respondents are considered confidential, we cannot link individual responses to the names of the counties or PMSAs in which those individuals reside. However, the NSFH provided us with a match between survey respondent case IDs and certain geographic characteristics ("geomerge").[18] For each wave, for each publicly available observation, the NSFH provided a corresponding data set with the observation case ID number and the char-

[18] We are extremely grateful to Larry Bumpass, Jack Solock, Charles Fiss, and the Center for Demography of Health and Aging at the University of Wisconsin–Madison for generously conducting this geomerge for us and providing us with the data. The use of these geographically merged but not individually identified data was approved by the Institutional Review Board at the National Bureau of Economic Research.

acteristics of the respondent's county and PMSA. While we cannot link individual respondents to named geographic locations, we can link individuals with the relevant characteristics of their counties and PMSAs in each wave. Included in our match are census data on county and PMSA population, education, and income, other geographic amenities like crime statistics and temperature, and the county and PMSA fixed effects on subjective well-being that we estimated previously using the BRFSS.

With the geographic characteristics from both NSFH waves, we are able to isolate the population of NSFH respondents who moved counties or PMSAs. In NSFH2, 2,395 respondents report moving cities since NSFH1. Using our matched data set, we find 1,939 respondents who both answered the question on subjective well-being and have different county characteristics for NSFH1 and NSFH2, denoting a change in the respondent's county of residence. Of that group, we similarly find 1,480 respondents to have moved to a new PMSA.

Our analysis focuses on the relationship between the changes in reported subjective well-being of this population and the changes in the respondents' county and PMSA characteristics. We run regressions of the form

$$\Delta y_i = \tau + \psi \Delta \hat{u}_i + \varphi \Delta X_i + \varepsilon_i \tag{5}$$

across individuals who move. The coefficient ψ identifies the relationship between changes in area-level happiness and changes in individual happiness, possibly controlling for changes in other covariates (at the area or individual level) between the two observations, captured in ΔX_i.

Table A3
Selection of Metropolitan Areas Ranked by Happiness

Rank	Metropolitan Area	Adjusted Happiness (Controlling for Demographics)	Adjusted Happiness (Controlling for Demographics and Income)	Unadjusted Happiness
1	Charlottesville, VA	.150	.080	.150
2	Rochester, MN	.144	.089	.156
3	Lafayette, LA	.141	.146	.117
4	Naples, FL	.126	.077	.144
6	Flagstaff, AZ	.121	.071	.112
7	Shreveport, LA	.120	.125	.089
15	Nonmetropolitan Hawaii	.103	.109	.050
16	Galveston, TX	.103	.067	.098
20	Norfolk, VA	.100	.069	.071
22	Honolulu, HI	.094	.079	.042
26	Colorado Springs, CO	.089	.049	.090
29	Washington, DC	.087	.044	.045

Table A3 (*Continued*)

Rank	Metropolitan Area	Adjusted Happiness (Controlling for Demographics)	Adjusted Happiness (Controlling for Demographics and Income)	Unadjusted Happiness
31	Raleigh-Durham, NC	.084	.040	.064
41	Tallahassee, FL	.076	.052	.041
43	Atlanta, GA	.074	.024	.043
52	Anchorage, AK	.069	.058	.056
56	Nashville, TN	.064	.084	.054
58	West Palm Beach, FL	.062	.034	.071
70	Minneapolis–St. Paul, MN	.053	.022	.063
77	Burlington, VT	.049	.025	.067
92	Baltimore, MD	.044	.016	.028
108	McAllen, TX	.039	.015	−.027
121	Nonmetropolitan Texas	.029	.031	.008
174	San Jose, CA	.004	−.045	−.006
179	Chicago, IL	.002	−.028	−.018
187	Seattle, WA	−.003	−.032	.005
242	San Francisco, CA	−.027	−.037	−.032
250	Nassau-Suffolk, NY	−.030	−.067	−.020
279	Vallejo-Fairfield-Napa, CA	−.042	−.050	−.056
284	Boston, MA	−.044	−.032	−.039
287	Los Angeles, CA	−.047	−.035	−.094
301	Las Vegas, NV-AZ	−.059	−.037	−.065
328	Detroit, MI	−.080	−.088	−.110
350	Non-metropolitan Indiana	−.104	−.061	−.080
353	Gary, IN	−.111	−.087	−.175
355	Pittsburgh, PA	−.115	−.071	−.095
359	New York, NY	−.123	−.120	−.159
364	South Bend, IN	−.138	−.104	−.126
365	Erie, PA	−.147	−.103	−.126
367	Scranton, PA	−.154	−.086	−.126

SOURCE.—Each metropolitan and rural area's adjusted life satisfaction is estimated after controlling for demographic covariates in a mixed effects model. Data are from Centers for Disease Control and Prevention (2005–10).

Table A4
Robustness of Happiness-Population Decline Relationship to Alternative Functional Forms (Based on Col. 5 of Table 3)

	Dependent Variable: Self-Reported Well-Being				
	(1)	(2)	(3)	(4)	(5)
Change in log population (below sample median) 1950–2000	.0781***				.0672***
	(.0127)				(.0161)
Change in log population (above sample median) 1950–2000	−.00443				−.00655
	(.00564)				(.00576)

Table A4 (*Continued*)

	Dependent Variable: Self-Reported Well-Being				
	(1)	(2)	(3)	(4)	(5)
Change in log population (below sample mean) 1950–2000		.0699***			
		(.0104)			
Change in log population (above sample median) 1950–2000		−.00745			
		(.00561)			
Change in log population, 1950–2000			.0797***		
			(.0128)		
Change in log population, 1950–2000, squared			−.0106***		
			(.00361)		
Change in log population, 1950–2000, is below sample median				−.0240***	−.00767
				(.00430)	(.00537)
Observations	1,182,563	1,182,563	1,182,563	1,182,563	1,182,563
R^2	.078	.078	.078	.078	.078

SOURCE.—Authors' regressions on microdata from the Behavioral Risk Factor Surveillance System Survey (Centers for Disease Control and Prevention 2005–10) and the US Census (Ruggles et al. 2010).

NOTE.—All regressions control for state fixed effects, year fixed effects, month fixed effects, age, race, sex, education, marital status, and family size. Standard errors (in parentheses) are clustered at the MSA-year level

*** $p < .01$.

References

Albouy, David. 2008. Are big cities really bad places to live? Improving quality-of-life estimates across cities. NBER Working Paper no. 14472, National Bureau of Economic Research, Cambridge, MA.

Alesina, Alberto, Rafael Di Tella, and Robert MacCulloch. 2004. Inequality and happiness: Are American and Europeans different? *Journal of Public Economics* 88, nos. 9–10:2009–42.

Alonso, William. 1964. *Location and land use*. Cambridge, MA: Harvard University Press.

Bates, Douglas M., and José C. Pinheiro. 1998. Computational methods for multilevel modelling. In *Technical memorandum BL0112140-980226-01TM.* Murray Hill, NJ: Bell Labs, Lucent Technologies. http://stat.bell-labs.com/NLME/CompMulti.pdf.

Baumeister, Roy F. 1991. *Meanings of life*. New York: Gilford.

Becker, Gary S. 1965. A theory of the allocation of time. *Economic Journal* 75, no. 299:493–517.

Becker, Gary S., and Richard A. Posner. 2004. Suicide: An economic approach. Unpublished manuscript, University of Chicago. https://web.archive.org/web/20060830175355/http://economics.uchicago.edu/download/Suicide_An_Economic_Approach_4.pdf.

Becker, Gary S., and Luis Rayo. 2008. Comments on "Economic growth and subjective well-being: Reassessing the Easterlin Paradox." *Brookings Papers on Economic Activity*, no. 1:88–95.

Benjamin, Daniel J., Ori Heffetz, Miles S. Kimball, and Alex Rees-Jones. 2011. What do you think would make you happier? What do you think you would choose? *American Economic Review* 102, no. 5:2083–2110.

Bentham, Jeremy. 1789 (1996). *An introduction to the principles of morals and legislation*. Oxford: Clarendon.

Cameron, A. Colin, Jonah Gelbach, and Douglas Miller. 2011. Robust inference with multi-way clustering. *Journal of Business and Economic Statistics* 29, no. 29:238–49.

CDC (Centers for Disease Control and Prevention). 2005–10. *Behavioral risk factor surveillance system survey data and questionnaire*. Atlanta: Department of Health and Human Services, CDC. http://www.cdc.gov/brfss/.

Clark, Andrew E., and Andrew J. Oswald. 1994. Unhappiness and unemployment. *Economic Journal* 104, no. 424:648–59.

Daly, Mary C., Andrew J. Oswald, Daniel Wilson, and Stephen Wu. 2011. Dark contrasts: The paradox of high rates of suicide in happy places. *Journal of Economic Behavior and Organization* 80, no. 3:435–42.

Deaton, Angus, and Arthur A. Stone. 2014. Evaluative and hedonic well-being among those with and without children at home. *Proceedings of the National Academy of Sciences* 111, no. 4:1328–33.

Di Tella, Rafael, Robert J. MacCulloch, and Andrew J. Oswald. 2001. Preferences over inflation and unemployment: Evidence from surveys of happiness. *American Economic Review* 91, no. 1:335–41.

Easterlin, Richard A. 1974. Does economic growth improve the human lot? In *Nations and households in economic growth: Essays in honor of Moses Abramovitz*, ed. Paul A. David and Melvin W. Reder, 89–125. New York: Academic Press.

———. 1995. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization* 27, no. 1:35–47.

Epictetus. 1916. Epicitetus: The discourses and manual. Together with fragments of his writings. Vol. 2. Translated with introduction and notes by P. E. Matheson. Oxford: Clarendon.

Fisher, Irving. 1892. *Mathematical investigations in the theory of value and prices, and appreciation and interest*. New Haven, CT: Yale University Press.

Florida, Richard, Charlotta Mellander, and Peter J. Rentfrow. 2013. The happiness of cities. *Regional Studies* 47, no. 4:613–27.

Gilbert, Daniel. 2006. *Stumbling on happiness*. New York: Knopf.
Glaeser, Edward L. 2007. *The economics approach to cities*. NBER Working Paper no. 13696, National Bureau of Economic Research, Cambridge, MA.
Glaeser, Edward L., and Joshua D. Gottlieb. 2008. The economics of place-making policies. *Brookings Papers on Economic Activity*, no. 1: 155–239.
Glaeser, Edward L., Joshua D. Gottlieb, and Oren Ziv. 2014. Unhappy cities. NBER Working Paper no. 20291, National Bureau of Economic Research, Cambridge, MA.
Glaeser, Edward L., and Janet E. Kohlhase. 2004. Cities, regions and the decline of transport costs. *Papers in Regional Science* 83, no. 1:197–228.
Glaeser, Edward L., Jed Kolko, and Albert Saiz. 2001. Consumer city. *Journal of Economic Geography* 1, no. 1:27–50.
Glaeser, Edward L., and Charles Redlick. 2009. Social capital and urban growth. *International Regional Science Review* 32, no. 3:264–99.
Glaeser, Edward L., Matt Resseger, and Kristina Tobio. 2009. Inequality in cities. *Journal of Regional Science* 49, no. 4:617–46.
Glaeser, Edward L., and Kristina Tobio. 2008. The rise of the sunbelt. *Southern Economic Journal* 74, no. 3:610–43.
Glaeser, Edward L., and Jacob L. Vigdor. 2001. Racial segregation in the 2000 Census: Promising news. Survey Series Report, Brookings Institution, Washington, DC.
Glenn, Norval D., and Charles N. Weaver. 1979. A note on family situation and global happiness. *Social Forces* 57:960–67.
Kahneman, Daniel, and Alan B. Krueger. 2006. Developments in the measurement of subjective well-being. *Journal of Economic Perspectives* 20, no. 1:3–24.
Kahneman, Daniel, and Richard Thaler. 1991. Economic analysis and the psychology of utility: Applications to compensation policy. *American Economic Review: Papers and Proceedings* 81, no. 2:341–46.
Lucas, Richard E. 2013. Life satisfaction of US counties predicts population growth. *Social Psychological and Personality Science* 5:383–88.
Ludwig, Jens, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu. 2012. Neighborhood effects on the long-term well-being of low-income adults. *Science* 337, no. 6101:1505–10.
Luttmer, Erzo F. P. 2005. Neighbors as negatives: Relative earnings and well-being. *Quarterly Journal of Economics* 120, no. 3:963–1002.
Mandeville, Bernard. 1714. *The fable of the bees: or, private vices, publick benefits*. London: Roberts.
Minnesota Population Center. 2004. *National historical geographic information system: Pre-release version 0.1*. Minneapolis: University of Minnesota. http://www.nhgis.org/.
Muth, Richard. 1969. *Cities and housing*. Chicago: University of Chicago Press.

Oswald, Andrew J., and Stephen Wu. 2010. Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science* 327, no. 5965:576–79.

———. 2011. Well-being across America. *Review of Economics and Statistics* 93, no. 4:1118–34.

Rappaport, Jordan, and Jeffrey Sachs. 2003. The United States as a coastal nation. *Journal of Economic Growth* 8:5–46.

Roback, Jennifer. 1982. Wages, rents, and the quality of life. *Journal of Political Economy* 90, no. 6:1257–78.

Rosen, Sherwin. 1979. Wage-based indexes of urban quality of life. In *Current issues in urban economics*, ed. Peter Mieszkowski and Mahlon Straszheim. Baltimore: Johns Hopkins University Press.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. 2010. *Integrated Public Use Microdata Series* (IPUMS): Version 5.0. Minneapolis: University of Minnesota.

Sacks, Daniel W., Betsey Stevenson, and Justin Wolfers. 2010. Subjective well-being, income, economic development, and growth. NBER Working Paper no. 16441, National Bureau of Economic Research, Cambridge, MA.

Stevenson, Betsey, and Justin Wolfers. 2008. Economic growth and subjective well-being: Reassessing the Easterlin paradox. *Brookings Papers on Economic Activity, Economic Studies Program* 39:1–102.

Stigler, George J. 1950. The development of utility theory. *Journal of Political Economy* 58, no. 4:307–27.

Sweet, James A., and Larry L. Bumpass. 1996. The National Survey of Families and Households—waves 1 and 2: Data description and documentation. Center for Demography and Ecology, University of Wisconsin–Madison. http://www.ssc.wisc.edu/nsfh/home.htm.

Sweet, James A., Larry L. Bumpass, and Vaughn Call. 1988. The design and content of the National Survey of Families and Households. National Survey of Families and Households Working Paper no. 1, Center for Demography and Ecology, University of Wisconsin–Madison.

Trull, Elaine, and Lisa Famularo. 1996. The National Survey of Families and Households: Field report. Temple University Institute for Survey Research, Philadelphia. ftp://elaine.ssc.wisc.edu/pub/nsfh/cmapp_n.001.