The Easterlin Paradox Re-examined Using Propensity Score Matching on Philippine Social Weather Station Data

Gerardo L. Largoza[1]
[1]*School of Economics, De La Salle University*
gerardo.largoza@dlsu.edu.ph

Abstract: We subject a portion of Easterlin's controversial yet correlational claim of a non-monotonic relationship between happiness and income (the so-called "Easterlin Paradox", 1974) to a stronger test of causality. We do this by using propensity score matching (PSM) techniques to create "matched pairs" that mimic treatment and control groups on a non-experimental dataset, the 2008 wave of the Philippine Social Weather Survey. By matching individuals based on similar propensity scores, we are able to make causal claims over a subset of the data, the "region of common support." In implementing PSM, we use two versions of the outcome variable subjective well-being: self-reported happiness and life satisfaction; we also use three income proxies (class of dwelling, self-rated poverty, and number of hard-up times per month) in the absence of direct measures within the survey. Our findings indicate that Easterlin's assumed positive relationship between income and happiness within countries is reproduced in only four of six possible runs, either with naïve regression or PSM. However, in only one case does using PSM increase the significance of the effect; in all others, using a causal effects methodology reveals how naïve regression overstates the happiness-income relationship. Finally, in all cases where the treatment effect is significant, the magnitude of the relationship nevertheless remains weak.

Key words: Easterlin Paradox; happiness and income; propensity score matching

1. INTRODUCTION

Exactly fifty years have passed since four lads from Liverpool first declared love a non-market good, to great acclaim[1] — and forty since one professor from Southern California

[1] Upon release, their two-minute proposition rocketed an unprecedented 27 places straight

shook the (social science) world by claiming the same thing about happiness, to continuing controversy.

First described in 1974, then buttressed in 1995 and reworked in 2010, the so-called "Easterlin Paradox" shows large-scale evidence that individuals, can find themselves "'flat of the curve', with additional income buying little if any extra happiness" (Clark et al, 2008). As a direct challenge to some of the strongest assumptions of economic theory and indeed Western thought, no one will be surprised to find an entire sub-discipline has emerged to explain the (non) relationship between income and happiness (Kahneman et al, 1999), work out its implications on public policy (Layard, 2011), and continually subject the data to closer scrutiny (Deaton, 2006; Inglehart et al, 2008).

From the start, examining the Paradox has not been simple. Even in the 1974 paper, Easterlin already describes the relationship between income and subjective well-being in three empirical propositions, rather than a single over-arching statement: (1) within a single country cross-section, individuals within higher incomes are on average happier than those with lower incomes; (2) across countries however, wealthy countries are not significantly happier than poorer countries; and (3) between 1946 and 1970 incomes rose dramatically in the USA but without a commensurate rise in self-reported happiness.

Over the years, others (most notably Stevenson and Wolfers in 2008) have presented critiques and contrary evidence, and Easterlin has responded by including more measures of subjective well-being (life satisfaction and financial satisfaction), and adding a further claim that (4) in the short run (within ten years), income does correlate with happiness, only for the relationship to disappear in the long run. But with its current mix of cross-sectional, time-series, short and long run propositions, Easterlin's paradox now looks more like the proverbial riddle wrapped in a mystery inside an enigma.

And beyond issues of measurement and data scope lies another often-ignored proviso: all these claims about happiness and income are correlational[2], not causal in the Rubin (1974) sense of measuring treatment effects from approximated counterfactual outcomes. Each of the relationships so far asserted is based on observational data, which fall short of the "gold standard for the estimation of causal effects": randomised experiments (Rubin et al, 2006). So while we may observe that in the short term, increases in income accompany (perhaps even predict) increases in happiness, we cannot really assert that higher incomes cause happiness to increase – at least not without data generated from an income-happiness experiment or quasi-experiment.

In this paper, we present the results of a small econometric exercise meant to estimate the causal effect of income on happiness from an observational dataset. It is "small" because it works only on the first and least controversial of the Easterlin claims: that within a country, individuals with higher incomes are happier than individuals with lower incomes. We subject the 2008 wave[3] of data from the Philippine Social Weather Survey (SWS) to a technique first proposed by Rosenbaum and Rubin (1983) called Propensity Score Matching (PSM). It works by constructing "a statistical comparison group that is based on a model of the probability of participating in the treatment, using observed characteristics" (Khandker et al, 2009). The method allows us to more closely approach the

to Number One, ironically prompting lead author McCartney to muse many years later that, "It should have been *Can* Buy Me Love" (Barry, 1997)

[2] Easterlin actually addresses issues of causality in the 1974 paper, but does so only constructively by pointing out that theory is silent on the impact emotions have on income (thus ruling out simultaneity) and citing the inclusion of the "hereditary upper class" in his cross-country samples (for whom emotions would play no role in influencing income).

[3] The most recent survey of self-reported happiness available for the Philippines, the next most recent being the 2001 wave of the World Values Survey.

Rubin counterfactual ideal by creating the equivalent of treatment and control groups from non-experimental data, and, under favourable statistical conditions, calculate causal effects — often labeled the "average treatment effect on the treated".

In this paper, we consider two alternative definitions of subjective well-being available: self-reported happiness and life satisfaction. And since income is not directly measured in the SWS dataset, we use three proxies: type of dwelling, self-rated poverty, and the number of times respondents experienced being hard-up that month.

We first try to replicate the Easterlin findings (that within countries, wealthy individuals are on average happier than poor individuals) using a binomial definition of subjective well-being. We then implement PSM using a number of available covariates to predict participation in the treatment (i.e., predict the likelihood of receiving high income in a hypothetical experiment), choose a matching procedure for the statistical pairs generated, and evaluate the treatment effect of "income" on well-being. A more detailed explanation follows.

2. METHODOLOGY

Dataset. The Fourth Quarter 2008 Social Weather Survey covered the entire Philippines, balancing respondents across the National Capital Region (NCR), Luzon, Visayas, and Mindanao. Sample size was 1500, with the minimum age of respondents at 18.

Variables: Subjective well-being.

Happiness: response to the question "If you were to consider your life in general these days, how happy or unhappy would you say you are on the whole?" Responses are coded as 1=very happy, 2=fairly happy, 3=not very happy, 4=not at all happy (question and responses translated into English). We consolidate the categories into a dummy in which 1 (happy)=responses 1 and 2, and 0 (unhappy)=responses 3 and 4.

Life satisfaction: response to the question "On the whole, are you:" 1=very satisfied, 2=fairly satisfied, 3=not very satisfied, 4=not at all satisfied. In the Filipino version of the question, the "summing-up of entire life" aspect is clearer, and distinguishes it from the more current evaluation of happiness above. As with happiness, we create strong and weak versions for life satisfaction.

Variables: Income.

Class of dwelling. The most direct proxy for income, as explained by Mangahas (2010): "Commercial survey interviewers are trained to assign their sample households into the following groups based mainly on the quality of dwelling: AB, upper class…C, middle class…D, lower class…E, extreme lower class." (note equivalence of dwelling with socio-economic class in definition). Partly to simplify, but mostly to deal with the extreme inequality of income within the dataset (households classified as A=3), we create a dwelling dummy: 1=ABC, 0=otherwise.

Self-rated poverty. Response to the question: "Where would you place your family?" 1=not poor, 2=on the line, 3=poor. Again, to counterbalance poor income distribution within the dataset, we create the dummy 1 (not poor)=1 and 2, and 0 (poor)=3.

Hard-up times. Response to the question: "How often in a month do the hard-up times come?" Available responses are 1=not hard up, 2=once, 3=twice, 4=three to five times, 5=six times to every other day, 7=everyday — for which we generate the dummy 1 (not hard up)=1, 0 (hard up)=2 to 7.

Variable: Treatment covariates. We list here predictors of treatment required to implement PSM. In most cases, we consolidated variables with multiple categories into dummies, to simplify the process of predicting likelihood of treatment.

Locale of household. 1=urban, 0=rural

Sex of the household head. 1=male, 0=female.

Age group of the household head. Twelve categories, min=18 max=88, which although defined as a range, we treat here as a continuous variable.

Educational attainment of the household head. Ten categories ranging from 1=no formal education to 5=finished high school to 10=postgraduate; transformed into dummy where 1 (secondary and below)=1 to 5, 0 (beyond secondary)=6 to 10.

Marriage status of the household head. Over ten categories describing various relationship states from never married to previously separated but with current partner; we combine categories to create a dummy 1 (no partner)=11 to 13, 0 (with partner)=all other categories.

Work status of household head. 1=working, 2=not working but worked before, 3=has never worked; consolidated into dummy 1 (working)=1, 0 (not working)=2 and 3.

OFW in household. 1=OFW member, 0=no OFW member.

Steps: Propensity score matching. Caliendo and Koepeinig (2005) lay out an eight-step procedure for implementing PSM, but for brevity, we focus on only three here. First, we regress a logit model that predicts the probability of treatment based on a set of observable covariates:

T_y = *f*(locale of household, sex of household head, age group of household head,…, OFW in household) (Eq.1)

where T_y is a binary variable representing available income proxies.

Assuming that the probability of being treated (i.e., of receiving high income in a hypothetical experiment) is sufficiently predicted by the covariates provided — and also that the covariates affect treatment without being affected by treatment —then composite matchable "twins" will be found for the observations that received treatment.

Second, these twins can then be matched using a variety of techniques: nearest neighbour, caliper, etc. (see Caliendo & Kopeinig, 2005 for a summary and practical discussion of each). Because of the extensive list of covariates provided (including one for locale), we allow the default unrestricted matching option used by Stata's *psmatch2* module.

Finally, *psmatch2* calculates the average treatment effect on the treated (ATT), which gives us the average causal effect of income on subjective well-being. But it also provides a useful comparison based on a "naïve" calculation of treatment effects from an "unmatched" sample: i.e., a straightforward regression of subjective well-being on income and all covariates. By comparing the estimates from the unmatched sample to the treatment effects generated by implementing PSM, we hope to replicate the Easterlin result (unmatched sample, correlational, should show income highly significant in predicting subjective well-being) and subject it to the stronger test of causality (ATT from propensity score matching).

3. RESULTS AND DISCUSSION

Below we present two tables consolidating the various average treatment effects on the treated (ATT) generated. They are organised per outcome variable (subjective well-being): Table 1 summarises the results for happiness while Table 2 summarises those for life satisfaction.

The rows pair estimates and t-stats of the (naïve) unmatched treatment effects (Easterlin-style correlations generated by logit regression over the whole sample) to the ATTs generated by propensity score matching, which use only a subset of observations that form the treatment and control "twins". These ATTs are found under the *Difference* column in Stata's *psmatch2* module; they are obtained by subtracting the mean outcomes of control from the mean outcomes of treated, both of whose values we do not reproduce here). We repeat this pairing of unmatched versus PSM results for all three available measures of the treatment variable income: class of dwelling, self-rated poverty, and hard-up times.

As we consider them an intermediate step en route to the main findings, we do not reproduce the logit regression results meant to predict treatment from observable covariates; they are, however, available upon request. We can also confirm that in all six PSM runs (two versions of subjective well being x three income proxies), locale and age of household head were significant 100% of the time, followed by educational attainment of household head, and presence of OFW (significant in four out of six runs).

If the first of Easterlin's propositions is supported in the Philippine sample, we would see it in significant and strongly positive effects, even within the unmatched sample. Note also that since all estimates are obtained from a logit regression, the coefficients are interpreted as odds ratios.

Table 1. Average treatment effects on the treated for outcome variable happiness; with number of treated (T) per income proxy; t-stats in parentheses

Happiness (1=happy, 0=unhappy)		
Income proxy	Unmatched Naïve	Matched PSM
Class of dwelling (T=103)	0.518 (1.20)	0.029 (0.47)
Self-rated poverty (T=369)	0.025 (0.97)	0.081 (2.15)*
Hard-up times (T=460)	0.113 (4.81)*	0.109 (3.00)*

Source: Author's computations

Table 2. Average treatment effects on the treated (unmatched [UM] and PSM) for outcome variable life satisfaction; with number of treated (T) per income proxy; t-stats in parentheses

Life satisfaction (1=satisfied, 0=unsatisfied)		
Income proxy	Unmatched Naïve	Matched PSM
Class of dwelling (T=103)	0.137 (2.88)*	0.116 (1.77)
Self-rated poverty (T=369)	0.114 (4.10)*	0.100 (2.50)*
Hard-up times (T=460)	0.132 (5.09)*	0.104 (2.63)*

Source: Author's computations

Can we replicate Easterlin's within-country findings predicting higher levels of subjective well-being associated with higher levels of income? The results are mixed. When subjective well-being is defined as (current) happiness (Table 1), we find Easterlin's findings replicated in only the hard-up version of income in the unmatched regression (as a weakly positive but significant odds ratio of high incomes raising happiness). When class of dwelling and self-rated poverty proxy for income, even the unmatched sample yields no positive relationship.

The PSM results (which cover only those observations in the dataset paired with the "treated" [high-income] individuals) do not differ substantially from the naïve regression. Class of dwelling has no causal impact on happiness, self-rated poverty has a significant but very weak causal effect, as does hard-up times.

Taken together, however, results show that the non-correlation between income (proxies) and happiness operates even in the within-country sample for the Philippines. While it is uncertain what this does to the Paradox (in a rejoinder to Deaton [2006], Easterlin constructs the Paradox thus: within countries, higher incomes raise happiness; but across countries, they do not), our findings strengthen the claim of a non-relationship between income and happiness by exhibiting it even in a within-country sample.

When (overall) life satisfaction is used, as in Table 2, we obtain a significant but weaker correlation between class of dwelling in the unmatched sample, which turns out to be insignificant in the alternative PSM estimation. For self-rated poverty and hard-up times, the other income proxies, the effect of using PSM is to reduce the significance levels of the original unmatched estimates, even as all the magnitudes suggest a weakly positive odds ratio between income and life satisfaction.

Reproducing the initial Easterlin findings is best done with self-rated poverty and hard-up times as income proxies, but even with asymptotically significant t-stats, the odds ratio is only weakly positive.

Most interestingly, the correlational claim arising from class of dwelling on income (significant and highest on life satisfaction) does not survive the test of causal effects.

4. CONCLUSIONS

This exercise shows that even the most uncontroversial of Easterlin's claims — that within countries, wealthier individuals are happier than poor individuals — cannot be completely reproduced, even with naïve regression. Indeed, in the four of six cases where a statistically significant relationship can be shown, the odds-ratios remain uniformly weak, never rising above 0.14. This suggests that the ability of high incomes to produce greater happiness is, at least within the Philippine sample, over-stated. It lends additional support to the Easterlin Paradox by broadening the scope of the claim: no longer does it operate only across countries and over the long term — it operates even within countries. Furthermore, using causal effects methodologies like PSM shows how naïve regression can over-estimate the significance and magnitude of treatment effects, albeit over a sample dramatically limited by the search for matched pairs.

In a later version of this paper, we shall examine the effects of various matching schemes (nearest neighbour, caliper, etc.) to see how sensitive results are to these options.

Natural extensions to the work include implementing causal effects methodologies on the remaining Easterlin propositions (i.e, cross-country and long-term happiness-income relationships), using a similar approach of first trying to replicate the findings, then comparing the naïve regression results to those obtained by stricter causality tests. For these, it may be practical to use difference-in-difference methods, or even regression discontinuity.

Finally, analysis of causal effects will remain incomplete until one looks into the mechanisms that show just *how* income affects subjective well-being. With more extensive surveys on the horizon (a new World Values Survey wave will be released in April), we may be able to better proxy some of the leading causal mechanisms proposed in the literature: hedonic adaptation, status anxiety, the erosion of relational goods (see Bruni & Porta, 2005). By combining more direct instruments to embody causal mechanisms with methods that allow us to construct reasonable counterfactuals, we can help disentangle one of the most complex, controversial, yet important relationships in the social sciences.

5. ACKNOWLEDGMENTS

An earlier version of the problem was submitted to the DLSU School of Economics as the undergraduate thesis of Jose Miguel C. Reyes, Feifei Shih, Dorothy D. Tan and Cherica Y. Vicente (author supervisor, April 2013). I thank Krista Danielle S. Yu for assistance with Stata.

6. REFERENCES

Barry, M. (1997). Paul McCartney: Many years from now. New York: Henry Holt and Company

Bruni, L., & Porta, P. L. (2005) Economics and happiness: Framing the analysis. Oxford University Press

Caliendo, M., & Kopeinig, S. (2005). Some practical guidance for the implementation of propensity score matching. IZA Discussion Paper No. 1588, Institute for the Study of Labour, Bonn

Clark, A. E., Frijters, P., & Shields, M. A. (2008). Relative income, happiness, and utility: An explanation for the Easterlin Paradox and other puzzles. Journal of Economic Literature, **46** (1), 95-144

Deaton, A. (2008). Income, health, and well-being around the world: Evidence from the Gallup World Poll. Journal of Economic Perspectives, **22**, 53-72

Easterlin, R. A. (1974). Does economic growth improve the human lot? Some empirical evidence. In Paul A. David and Melvin W. Reder, eds., Nations and Households in

Economic Growth: Essays in Honor of Moses Abramovitz, New York: Academic Press, Inc.

Easterlin, R. A. (1995). Will raising the incomes of all increase the happiness of all? Journal of Economic Behavior and Organization, **27**, 35-74

Easterlin, R. A., McVey, L. A., Switek, M., Sawangfa, O., & Zweig, J. S. (2010). The happiness-income paradox revisited [Electronic version]. Proceedings of the National Academy of Science, **107** (52), 22463-22468

Inglehart, R., Foa, R., Peterson, C., & Wetzel, C. (2008). Development, freedom, and rising happiness: A global perspective (1981-2007). Perspectives in Psychological Science, **3**, 264-285

Kahneman, D., Diener, E., & Schwarz, N. (eds.). (1999). Well-being: The foundations of hedonic psychology. New York: Russell Sage Foundation

Khandker, S. R., Koolwal, G. B., & Samad, H. A. (2009). Handbook on impact evaluation: Quantitative methods and practices. The World Bank

Layard, R. (2011). Happiness: Lessons from a new science, 2nd revised edition. Penguin

Leonhardt, D. (16 April 2008). Maybe money does buy happiness after all. The New York Times

Mangahas, M. (2001). Subjective poverty and affluence in the Philippines. Philippine Review of Economics, **38** (2), 122-134

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. Biometrika, **70**, 41-55

Rubin, D. (1974). Estimating causal effects of treatment in randomized and non-randomized studies. Journal of Educational Psychology, **66** (5), 688-701

Rubin, D. B., & Waterman, R. P. (2006). Estimating the causal effects of marketing interventions using propensity score methodology [Electronic version]. Statistical Science, **21** (2), 206-222

Stevenson, B., & Wolfers, J. (2008). Economic growth and subjective well-being: Reassessing the Easterlin Paradox. Brookings Papers on Economic Activity, Spring, 1-8