Title: The Causal Effect of Age on Subjective Well-Being

Extended Abstract for PAA Meeting in San Diego, 2015

Authors: Clemens Noelke[1], Marcia P. Jimenez[2], Daniel Corsi [3]

Affiliations:

1. Center for Population and Development Studies, Harvard School of Public Health, Cambridge, MA
2. Brown University School of Public Health, Department of Epidemiology, Providence, RI
3. Clinical Epidemiology Program, Ottawa Hospital Research Institute, Ottawa, Canada

Word count: 2,851

Address for correspondence:

Clemens Noelke
Center for Population and Development Studies
9 Bow Street
Cambridge, MA 02138
Email: cnoelke@hsph.harvard.edu

Abstract:

Empirical research has thus far not clarified whether and how subjective well-being changes with age. We replicate, clarify and reconcile evidence from three widely-cited studies that use the same US dataset to arrive at very different conclusions about how well-being changes with age. We explore three sources of bias to explain the discrepant findings: bias due to period or cohort effects, collider bias, and functional form misspecification. Our results indicate that estimated age effects are highly sensitive to the inclusion of collider variables, such as marital status or health, which lie on the causal path from age to happiness. We suspect that inappropriate conditioning on collider variables when estimating age effects is more widespread, which is why our results illustrate a potentially important form of misspecification in research on aging.

Do we get happier as we age? Numerous theories have been advanced about the relationship between subjective well-being and age over the life course (see review in Ulloa, Moller and Sousa-Poza 2013). Because growing older tends to be associated with poorer health, less money, and fewer social contacts, one might expect that subjective well-being declines in older ages. However, ageing may be associated with increased happiness as unrealistic ambitions are given up and individuals growing awareness of their own mortality causes them to shift focus on activities that increase present rather than future well-being. Set-point theory presumes that individuals have a fixed level of well-being. While major live events cause temporary departures, well-being levels eventually return to this fixed level, suggesting no major change of well-being with age.

Theoretical predictions diverge, and empirical studies using datasets from across Western countries have been unable to clarify the association between age and happiness (see review in Ulloa, Moller and Sousa-Poza 2013). Even studies using the same US dataset and similar model specifications come to very different conclusions about how happiness changes with age: Blanchflower and Oswald (2008, 581 citations on Google Scholar) find that happiness declines until mid-adulthood and then increases ("U-shape"). Easterlin (2006, 278 citations on Google Scholar) finds that happiness increases until mid-adulthood and then declines ("inverted U-shape"). Yang (2008, 243 citations on Google Scholar) finds both evidence of a U-shape effect and increasing happiness throughout the life course.

In this study, we replicate, clarify and reconcile Blanchflower and Oswald's (2008), Easterlin's (2006) and Yang's (2008), discrepant findings. We conceptually illustrate and empirically explore three sources of bias to explain why the results of these studies differ and offer some general lessons about the challenges involved when estimating the causal effect of age on happiness and other outcomes.

First, each study opts for a different strategy to identify age effects net of period and cohort effects, but none of them tests whether the results are robust to alternate identification strategies. We illustrate, using Directed Acyclic Graphs (DAGs), the assumptions underlying the common approaches to the age-period-cohort problem (APC) that the studies adopt to estimate age effects. We then test empirically how robust the results are to using different identifying approaches. Second, using DAGs we explore the role of collider bias or endogenous selection bias (Elwert 2013). Both Yang (2008) and Blanchflower and Oswald (2008) use condition on marital status and self-rated health, variables that lie on the causal path from age to happiness. Using DAGs, it can be shown that if these variables have an unobserved cause that also predicts happiness, the resulting age effect estimate will be biased. We explore empirically whether results are sensitive to the inclusion of these colliders. Third, while all three studies opt for a quadratic functional form, we argue that this may be restrictive and consider additional moments in the happiness age response function.

We use the same data and variables as the three studies cited, i.e., repeated cross-sections for nearly forty years from the General Social Survey (1972-2010). Our measure of happiness is the response to the question "Taken all together, how would you say things are these days: Would you say that you are very happy, pretty happy, or not too happy?". The GSS provides the longest-running individual data source on happiness, capturing a large number of cohorts over long periods of time and is therefore ideally suited for a methodological exploration of age-period-cohort models.

Concerning the identification of age effects net of cohort and period effects, we emphasize that the effect of ageing on happiness (or any other outcome) is fundamentally unidentifiable because of the exact linear dependency between age, period and cohort. Age effects are only identified by imposing restrictions on how cohort and period effects are adjusted for. Using DAGs, we aim to provide a crisp illustration of how demanding the assumptions necessary to identify causal age effects are, focusing on the three currently dominant approaches: Models using a continuous age variable with coarsened cohort fixed effects and unrestricted period fixed effects (used by Blanchflower and Oswald [2008] and Easterlin [2006]), models using a continuous age variable and cross-classified age-period random effects (used by Yang[2008]), and Winship and Harding's (2008) approach.[1] The cross-classified period-cohort random effects model proposed by Yang (2008), for example, assumes that age is conditionally independent from birth cohort random effects without providing evidence that this assumption is plausible.[2] However, the approach used by Blanchflower and Oswald (2008) and Easterlin (2006) similarly requires demanding assumptions that go unchecked, leaving readers wondering whether results are robust to different ways of solving the APC problem.

Conditioning on an endogenous mediator variable introduces bias in an otherwise well-identified treatment effect (Elwert 2013, Angrist and Pischke 2009). Figure 1 visualizes the associations between age, marital status and happiness. Marital status lies on the causal path from age to happiness, because changes in marital status increase in probability with age. Marital status is likely to be selective in terms of unobserved characteristics, such as personal maturity (U). If U also predicts the outcome variable, conditioning on marital status introduces collider bias or endogenous selection bias by inducing a non-causal association between age and happiness via the variable U (Elwert 2013). Even if age was randomized, the age effect would be biased if marital status is adjusted for. Blanchflower and Oswald (2008) and Yang (2008) condition on marital status in their final models as well as numerous other potential colliders such as labor market status, having children, income, religious attendance and self-rated health, each of which is likely endogenous, i.e. predicted by another unobserved variable that also predicts happiness.

[1] Applied to the causal analysis of age effects, the latter approach omits, for example, cohort effects from the estimation model but uses individual covariates to fully block any non-causal association between age and happiness that would run through cohort differences.

[2] The assumption is prima facie implausible since in most individual dataset age and birth year are highly correlated.

Using DAGs, it can be clearly shown that the inclusion of these variables contributes nothing to identification of the age effect but rather introduces bias.

Figure 1. Directed Acyclic Graph Illustrating Collider Bias

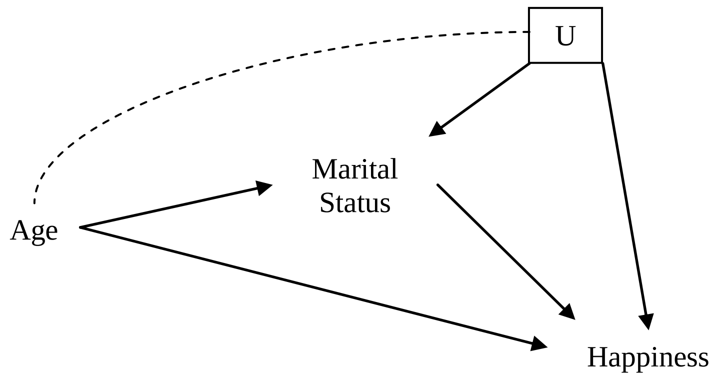

We find that the estimated age effects are highly sensitive to the inclusion of collider variables. Glenn (2009) already noted that the effect of happiness on age is sensitive to the inclusion of marital status in models using coarsened cohort and unrestricted period fixed effects. Beyond this work, we show that none of the approaches to model APC effects referred to above is robust to the inclusion of collider variables: Using the three different approaches without additional controls and an age quadratic, we show that the effect of age on happiness follows an inverted U-shape ($p<0.001$), increasing until mid-adulthood and then decreasing. In each of the three modeling frameworks, adding a single collider variable, marital status, causes the sign of both quadratic terms to flip, yielding an inverted U-shape ($p<0.001$). Furthermore, we illustrate that the age effect is similarly sensitive to the inclusion of other colliders such as health and income.

The conditioning on collider variables is our main explanation of the discrepant findings from GSS data reported above. The increase of happiness after mid-adulthood that Blanchflower and Oswald (2008) as well as Yang (2008) report is entirely driven by the inappropriate adjustment for a collider variable. We furthermore show that the estimated happiness-age function is similarly sensitive to the inclusion of other collider variables that Blanchflower and Oswald (2008) and Yang (2008) use. We suspect that the use of collider variables in research on age effects is more widespread, which is why our results illustrate a potentially important form of misspecification in research on ageing.

Furthermore, we show that the quadratic functional form for the age effect is too restrictive. The inverted U-shape is essentially driven by small increases in well-being in early life and small decreases in well-being in late life. Using Deviance Information Criterion (DIC) statistics to compare model fit, our preliminary results indicate that a fifth order polynomial provides the best description of the variation in happiness with age. Using this model, any increase of well-being in mid-life disappears, and we observe roughly trendless fluctuation over the life course, a finding that we interpret to be most consistent with set-point theory. This pattern is robust to different approaches to identify age effects. In this particular application, the inappropriate adjustment for collider variable as well as an inappropriately specified functional form contributed to biased results in published research.

References

Angrist, Joshua D. and Pischke, Jörn-Steffen (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton: Princeton University Press.

Blanchflower, David and Andrew Oswald (2008). "Is Well-Being U-shaped over the Life Cycle?" *Social Science & Medicine* 66: 1733-1749.

Blanchflower, David and Andrew Oswald (2009). "The U-shape without Controls: A Response to Glenn". *Social Science & Medicine* 69: 486-488.

Elwert, Felix. "Graphical Causal Models." Pp. 245-273 in S. Morgan (ed.), Handbook of Causal Analysis for Social Research. Dodrecht: Springer

Glenn, Norval (2009). "Is the Apparent U-shape of Well-Being over the Life Course a Result of Inappropriate use of Control Variables? A Commentary on Blanchflower and Oswald". *Social Science & Medicine* 69: 481-485.

López Ulloa, Beatriz Fabiola, Valerie Moller and Alfonso Sousa-Poza (2013). "How Does Subjective Well-Being Evolve with Age? A Literature Review." *IZA Discussion Paper* No. 7328. Bonn: Institute for the Study of Labor (IZA).

Yang, Yang (2008). "Social Inequalities in Happiness in the United States, 1972 to 2004: An Age-Period-Cohort Analysis." *American Sociological Review* 73:204-236.