Dynamic Analysis and the Economics of Happiness: Rationale, Results and Rules

Piper, Alan T.

Universität Flensburg, Staffordshire University, UK

June 2012

Dynamic Analysis and the Economics of Happiness: Rationale, Results and Rules

Alan T. Piper (Universität Flensburg)[1,2]
December 2012

This paper provides a sustained introduction for the use of dynamic panel methods when analysing life satisfaction. As well as being able to address the issue of serial correlation, dynamic panel analysis also has the advantage of being able to treat variables as exogenous or endogenous, important for happiness, and can generate both contemporaneous and long run estimates for independent variables. A key result found initially for young people, but which is robust to different age ranges and countries, is that happiness is largely contemporaneous although there is a small, persistent effect of the past on current happiness. Additionally, decision rules are provided for the analysis of happiness using dynamic panel analysis.

1 Address for correspondence: alan.piper@uni-flensburg.de

2 I am enormously grateful to Nick Adnett, Andrew Clark, and Geoff Pugh for thorough, useful, and highly beneficial comments, suggestions and discussions. This paper is far better than it would have been without their input. The usual disclaimer applies.

Dynamic Analysis and the Economics of Happiness: Rationale, Results and Rules

1. Introduction

The aim of this paper is to introduce a new method to the economic analysis of the concept of happiness. The ‘workhorse’ model within this field is fixed effects analysis taking advantage of the rich nature of nationally representative samples, like the British Household Panel Survey (BHPS) and the German Socio-Economic Panel (SOEP). Such an analysis has provided many insights for a scientific understanding of well-being. Useful reviews of these studies include Clark et al. 2008 and MacKerron (2012). Using a subset of the BHPS, initially, this study demonstrates that many of these studies, by using fixed effects analysis, are potentially misspecified. The presence or otherwise of serial correlation is not tested for in the literature, and the analysis here, using a well-known and well-utilised data set demonstrate that it is a substantial issue. The presence of serial correlation in the idiosyncratic error term means that there are omitted dynamics in the FE estimates. As King and Roberts (2012) forcefully argue, this should not be treated as a problem to be fixed by adjusting the standard error but instead treated as an opportunity to take advantage of this information and respecify the model.

The respecification here, which results from the strongly significant finding of serial correlation in the idiosyncratic error term, is to employ dynamic panel methods. This introduces a lagged dependent variable on the right-hand-side of the equation, and substantially changes the interpretation of the coefficients for the independent variables. Such an analysis also introduces more methodological considerations, including the ability to choose whether the independent variables are endogenous and exogenous. As we shall see, such a choice can substantially change the outcome of the results for well-being estimates. As the main focus of this paper is about introducing and discussing dynamic panel methods and well-being, the bulk of the work below is given over to explanations and consideratons associated with the use of the method and the interpretation of the results. The twenties age range focus can be mainly seen as the frame that the analysis hangs on. That said the results are robust to other age ranges, and datasets too. What is described here appears somewhat more universal than for British individuals in their 20s.[3]

The analysis for the twenties, using standard controls, (marital status, job status, income, health, education) is in line with the U-shape found in many studies over the whole lifecycle (Clark and Oswald (1994), Blanchflower (2001), Di Tella et al. (2003), Blanchflower and Oswald (2008) are just four of many examples). Happiness appears to decline over the twenties age range. A novel result that stems directly from the use of dynamic methods is that happiness is largely contemporary: the influence of the past on contemporaneous well-being is approximately ten percent. This is robust for other age ranges, the whole life-cycle, other countries, and alternative ways of measuring wellbeing.

This paper is organised as follows. Section 2 provides a brief rationale for an investigation into part of the lifecycle generally, and the twenties generally, the main basis of this research. Due to a lack of previous economic analysis in this area, some relevant psychological literature is also reviewed.[4] Section 3 presents the static panel analysis regarding the happiness of young people (defined here as individuals in their

3 This twenties age range focus reflects the initial focus of the research, which was somewhat overtaken by the interesting methodological considerations and outcomes discussed here.

4 In psychology this period is known variously as late adolescence, emerging adulthood or early adulthood.

twenties) and provides the rationale for the dynamic panel analysis, the results of which are found in section 4. Section 4 also discusses the issues surrounding a dynamic analysis of wellbeing, and demonstrates that the central result is robust to other age ranges, data sets and countries. Section 5 offers concluding remarks.

2. Age and Well-Being

> "Despite all the recent research regarding happiness and subjective well-being a fundamental research question remains poorly understood. What is the relationship between well-being and age?"
>
> Blanchflower and Oswald (2008 p.1733)

This section reviews and discusses some of the issues regarding an economic assessment of age and well-being, and offers reasons why it is potentially valuable to investigate different age groups separately. This section also provides a rationale for investigating the twenties part of the lifecycle, the initial focus of the empirical work. Relevant literature, from both economics and psychology, is discussed. This reliance on psychology will be particularly important when discussing young people, given the nature of the studies over the lifespan and its different parts.

Blanchflower and Oswald (2008) investigate what they call the <u>ceteris paribus</u> relationship between age and well-being, factoring out many other considerations to get to an underlying pattern. This approach is taken to achieve an idea of the relationship between age and well-being not considering the many different elements that are important at different life stages: “the paper’s concern is with the ceteris paribus correlation between well-being and age. Hence we later partial out some other confounding factors, such as income and marital-status, that alter over a typical person’s lifetime and have an effect upon well-being” (p.3). Their results suggest a U-shape relationship between life satisfaction and age, with the conditional regressions finding reported well-being being higher at both the start and at the end of the adult lifespan.

Most studies in economics using the ‘many controls’ regressions of Blanchflower and Oswald (2008) find a similar U-shape. This U-shape has proved quite robust, being found in Clark and Oswald (1994), with Blanchflower (2001) demonstrating this shape in 23 East European transition nations and Di Tella et al. (2003) getting the same result for 12 Western industrial nations. The study mentioned above by Blanchflower and Oswald (2008) replicates the U-shape in 72 countries, when income and demographic variables are held constant.

A potential problem with multivariate regressions that find a U-shape (by controlling for many other factors) is that the controls assume the same definitions and standards for everyone, aged twenty, fifty, or eighty. Good health, for example, is assumed to have the same meaning for everyone regardless of age; yet an 80 year old may have a different conception of good health than a twenty year old. The multivariate regressions will not pick this up, and this is the specific reason for Blanchflower and Oswald (2008) not including physical health as a control. Clark (2007) explains similarly: “in the context of well-being and age… it is contentious to include health as a right hand side variable, although this practice is widespread in the literature. Including health does imply that we are comparing individuals of different (working) ages, but with the same level of health." (p. 11). As health in these studies is often self-reported (as excellent, good etc.) the multivariate regressions are, controlling for an individual's appreciation (or otherwise) of their health, which is perhaps less contentious, because we are asking how an individual's belief about their own health contributes to an estimation of their own life satisfaction. Objective health may be different but a subjective satisfaction with health is what is asked about and the assumption is that this satisfaction has the same impact on happiness. What the individual considers to be excellent health at 80 is highly likely to be different to excellent health at 20 but the contribution to subjective well-being is held (controlled) as the same. If differences in health matter for well-being, and the stylised result in the happiness literature is that health matters greatly, how should we account for it in an investigation of the underlying relationship between age and well-being? One solution would be to look at smaller age ranges and piece the results together. Another solution would be to use interaction terms to capture any differing effect of heath over the lifecycle. A further comment in the literature on health and education is that including them as controls could possibly bias the coefficient on income downwards "since you will not be capturing its impact through health and education" (Bottan and Perez-Truglia 2010, p.6), however they also state that "controlling for as many covariates as possible is our strategy to minimise the potential biases" (Bottan and Perez-Truglia, 2010, p.6).[5] However, such a strategy will not capture indirect effects, i.e. those mediated by other variables. The twenties age range, the original focus here, is obviously more homogenous than the whole adult lifespan[6], and the conditionality placed on the age and well-being

5

This paper is discussed further when the dynamic estimations are performed in section 4.

6

relationship means that the concerns about the health controls are thus mitigated somewhat. In our analyses below it is therefore less contentious to include health as a right-hand side variable.

Thinking about the left hand-side of the standard equations (i.e. happiness) also provides a rationale for a focus on a particular part of the lifecycle. In short, is the left hand side variable the same for everyone in the sample? Some studies assume that it is (Layard 2005; Myers and Diener 1995), whereas some studies argue that it is highly subjective and so different for everyone (Gilbert 2005). However, perhaps there is a systematic way that happiness differs between ages meaning that it is useful to study isolated parts of the lifecycle? There is some evidence that this is so. For example, Kamvar et al. (2009) in an analysis of twelve million blogs find that younger people (the paper is not precise about what this means) refer to happiness as excitement whereas older people refer to it as feeling peaceful. In the blogs being happy was associated with high arousal words for young people, whereas the association was with low arousal words for older people (again, there is no clear definition of older people). They offer support for this finding in three experiments which demonstrate the same thing (or similar) in different ways. These experiments demonstrate a statistically significant difference between the two age groups, though a weakness is that some of the experiments are based on a very small sample size. They argue that this change in how happiness is viewed is driven by an increasing sense of connectedness (to others and the present moment). The difference in what happiness (and therefore self-reported happiness) means to different age groups is potentially important.

Recall the discussion of the changing meaning of happiness over the lifecycle in section 3

An update of the analysis in Kamvar et al. (2009) demonstrates the relative importance of excitement and peacefulness with respect to what individuals' regard as happiness in different age ranges (Mogliner et al. 2011). For the twenties age range the ratio of excited happiness to peaceful happiness is about 1.5 to 1, and, as seen on the table below, the change throughout the lifecycle is striking. Note well that the first line of the table and figures are as presented in the original study, whereas the second line is a slight rebasing of the figures for easier comparison.

Table 1 Ratio of happiness as excitement to happiness as peacefulness across the lifecycle

	Teens	20s	30s	40s	50s
Excited happiness/ peaceful happiness	1.85:1 1: 0.54	1.48:1 1: 0.68	1:1.19 1:1.19	1:3.42 1:3.42	1:8.00 1:8.00

(Source: Mogliner et al., 2011, table 3)

The discussion now turns to the literature, again from within economics and psychology, which investigates the happiness of young people specifically. Psychologists have studied this particular age range in more detail. Some have suggested that fundamental changes occur over the twenties age range. Bee(1997), for example recognises that in this age range personality changes towards more autonomy, and more striving for achievement as well as increased self-confidence and personal assertiveness. She also notes that individuals in this age range become physically independent of their family, but more psychologically independent too. This personality shift is clearly one that is not limited to one particular cohort, and lends itself to a lifecycle interpretation of changes in happiness over the twenties. Similarly, according to Loevinger (1976, 1984), who set out a framework in psychology for measuring non-observables like the quality of life in 1957, the underlying change in early adulthood is a shift from an external to an internal definition of oneself. Demographer Ronald Rindfuss (1991) refers to the years between 20 and 30 as ‘demographically dense’, because of the many events that take place in this period. A focus in psychology is made on the ‘transition to adulthood’ is in line with a ‘lifecycle’ approach: i.e. the changes are due to getting older. This transition argument suggests that over the twenties – “years of profound change and importance” (Arnett, 2000) – what is important for subjective well-being may change.

Thus, the above discussion demonstrates why happiness or life satisfaction may change over the twenties age range (with potential reasons) and emphasises again that the age range is worth assessing separately from other age ranges. The brief review of research from psychology, it appears that different periods of the lifecycle (and different cohorts) require separate studies to properly understand the causes and correlates of happiness in these different age ranges. This is similar to the argument made by Buss, the evolutionary psychologist who has undertaken investigations into happiness, who argues that sources of happiness differ profoundly for individuals of different ages (Buss, 2000). Carstensen’s theory, (cited in Blanchflower and Oswald 2008), argues that ageing is associated with increasing motivation to derive emotional meaning from life and decreasing motivation to expand one's horizons. In short, as one ages one feels more connection to the present moment. Again, within psychology, the field of gerontology is clearly predicated on an explicit recognition that at different ages, different things are important for well-being and the quality of life and, as this quote from a survey of gerontology argues, “it is not appropriate to measure well-being in old age by the same standards that apply to middle age, namely, standards based upon activity or social involvement" (Neugarten et al 1964, p.134). Recent work has started to explore this notion that, for different age groups, happiness might mean, and be derived from, different things. Lelkes (2008) uses evidence of the U-shape as a start, via European Social Survey (cross-section data), to argue that changes for older people are explained, partly, by changing preferences and, partly, by changing circumstances.

Within economics, only one study (to my knowledge) has specifically investigated the well-being of younger individuals. Blanchflower and Oswald (2000) assess the well-being of the young (meaning individuals under 30) in the United States, and in many European countries, via data from the U.S. General Social Surveys and the Eurobarometer for approximately twenty years from the early 1970s. Using ordered logit analysis, they find that the probability of reporting a high happiness score is lowest at the end of one's twenties, which may be initially taken as suggestive of a 'lifecycle' result. However, since both the U.S. General Social Surveys and the Eurobarometer are cross-section data, the result could just as easily be a cohort result. This apparent declining happiness over the twenties period is consistent with the U-shape, a result examined here with more recent British data in section 3. They also find a statistically significant, positive time trend, suggesting that 'young Americans become steadily happier since the 1970s' (p.7). For comparison, a time trend effect for the over 30s was found to be small and negative. For Europe, Blanchflower and Oswald (2000) find a slightly upward trend in young people reporting increased happiness over roughly the same period. Though the abstract acknowledges that "many commentators believe that life in the industrial nations is getting tougher for the young… [and that] the evidence in this paper paints a different picture. The paper documents a rising level of happiness among young people in western countries" (p.289), these results are not necessarily in contrast to the claims made regarding a 'quarter-life crisis'. This is in part because the years of the data used (1972-93) do not reflect the main period of concern for the advocates of this phenomenon. Also as noted above the data sets used are not panels, so the same individuals are not being followed over time and that therefore the results could be due to different people being interviewed each year. This also restricts the econometric methods available to investigate the data. In Europe they generally find for individuals under 30 a statistically significant upward time trend in 13 out of 15 countries. However, the two countries where it is not found are Great Britain and Ireland, and Blanchflower and Oswald comment that "why the British Isles misses out on this recent growth of well-being amongst the young is a puzzle" (2000, p.302). Thus they cannot reject the belief that life 'is getting tougher for the young' in the UK, home of the individuals investigated here.

The next sections make use of panel data to empirically examine the happiness of individuals in the twenties age range in Britain via multivariate regressions on this shorter life span. As noted earlier this removes some of the methodological concerns discussed above, and enables a better picture of the happiness of young people to emerge, as compared to studies of the whole life span, where happiness has variable and often different meanings for the individuals' investigated.

3 Happiness in the twenties: panel analysis

The benefits of panel analysis as compared to pooled cross section analysis are numerous. Arguably the most important benefit is that individual heterogeneity can be controlled for, and this helps us overcome Bentham's well-known apples and oranges concern. This is particularly so for fixed effects estimations, where the comparison is 'within' a person, which removes the need to compare between individuals. This estimation method effectively 'controls' for the time invariant characteristics of each individual, meaning that FE regressions allow (or control) for differences in personality and disposition that may be important determinants of life satisfaction. The choice of estimating method is discussed further below, and the subsequent analysis follows the approach, initially, of much published literature (some of which are discussed above). The variables included in the regression are those suggested by the 'economics of happiness' literature which, in turn, provide the framework for the investigation. These variables are discussed below.

The data come from the British Household Panel Survey (BHPS), with the dependent variable being life satisfaction, measured on an ordinal scale from 1 to 7, 'not at all satisfied' to 'completely satisfied'. The independent variables are income, job status, marital status, education, and health, wave and regional dummies are also included. The other variables in the regression account for the region, the year (wave dummy variables) and age bands.

Combined, these are the variables that make up the happiness functions that are estimated in sections 3 and 4. The next few paragraphs briefly summarise the reasons advanced by orthodox economic theory for the inclusion of these variables in happiness functions. The specification adopted here is typical of the estimations in the empirical economic literature referred to throughout the thesis. There are anticipated results that reflect both orthodox economic theory and other empirical work. Work on well-being is continuing in each of these areas and, as yet, there are no stylised ‘facts’, particularly because of issues related to endogeneity and causality.

Initial diagnostic tests (not reported here) establish that the workhorse model, FE, is the preferred static model, being more appropriate than RE and OLS. Results for these regressions are presented in table 2: the different FE estimations in the table contain wave and period dummies (columns 2 and 4) and robust standard errors to correct for heteroscedascticity (columns 3 and 4). A discussion of these results follows the table.

Table 2 Fixed effects life satisfaction regressions for individuals in their 20s

	(1)	(2)	(3)	(4)
	FE	FE	FE	FE
			Robust SE	Robust SE
VARIABLES	Life Satisfaction	Life Satisfaction	Life Satisfaction	Life Satisfaction
Income	0.04***	0.05***	0.04***	0.05***
	(0.014)	(0.014)	(0.014)	(0.015)
Self-employed	-0.04	-0.04	-0.04	-0.04
	(0.050)	(0.050)	(0.049)	(0.049)
Unemployed	-0.30***	-0.31***	-0.30***	-0.31***
	(0.037)	(0.037)	(0.047)	(0.047)
Retired	-0.32	-0.28	-0.32***	-0.28***
	(1.004)	(1.004)	(0.074)	(0.079)
Other Labour Force Status	0.02	0.01	0.02	0.01
	(0.026)	(0.026)	(0.028)	(0.029)
Married	0.06*	0.07**	0.06**	0.07**
	(0.032)	(0.033)	(0.031)	(0.032)
Separated	-0.11	-0.11	-0.11	-0.11
	(0.086)	(0.087)	(0.116)	(0.117)
Divorced	0.05	0.07	0.05	0.07
	(0.099)	(0.100)	(0.144)	(0.146)
Widowed	-0.48	-0.49	-0.48	-0.49
	(0.511)	(0.512)	(0.992)	(0.985)
Education: High	-0.04	-0.03	-0.04	-0.03
	(0.086)	(0.087)	(0.099)	(0.101)
Education: Medium	-0.01	-0.02	-0.01	-0.02
	(0.087)	(0.088)	(0.099)	(0.100)
Health: Excellent	0.44***	0.44***	0.44***	0.44***
	(0.026)	(0.026)	(0.029)	(0.029)
Health: Good	0.31***	0.31***	0.31***	0.31***
	(0.021)	(0.021)	(0.024)	(0.024)
age2324	-0.06***	-0.04	-0.06***	-0.04
	(0.022)	(0.029)	(0.023)	(0.030)
age2526	-0.07***	-0.00	-0.07**	-0.00
	(0.025)	(0.044)	(0.027)	(0.045)
age2729	-0.11***	-0.01	-0.11***	-0.01
	(0.028)	(0.063)	(0.031)	(0.064)
Wave Dummies	No	Yes	No	Yes
Region Dummies	No	Yes	No	Yes
Constant	4.95***	4.97***	4.95***	4.97***
	(0.078)	(0.112)	(0.089)	(0.119)
Observations	22,967	22,695	22,967	22,695
Number of pid	7,312	7,284	7,312	7,284
Adjusted R-squared	-0.430	-0.432	0.026	0.028

Standard errors in

parentheses
*** p<0.01; ** p<0.05;
* p<0.1

The results indicate that the following are positive and significant for life satisfaction: income (which is in thousands of pounds and deflated by the CPI); being married; and categorising health as good or excellent. Being unemployed is statistically significant and negative As can be seen, altering the estimation to include wave and regional dummies, and using robust standard errors does not qualitatively change these results by much. The main difference is that the few individuals who are retired are significantly less happy, on average, than those employed: a result perhaps consistent with the non-pecuniary arguments for the loss of happiness due to unemployment (the size of the coefficient is similar to that for the unemployed). When the wave and region dummies are included, the dummies for the age categories are insignificant, but when they are not included the age dummies are statistically significant and negative (when compared to being in the 20 to 22 age range category). Further checks demonstrate, perhaps unsurprisingly, that it is the wave dummies responsible for the loss of significance of the age dummies. Thus the results for age dummies may be picking up wave effects. Despite this possibility, it is interesting to ask whether a focus on the twenties age range can offer support for the U-shape found via age and age-squared over the whole lifecycle. Over the twenties does ceteris paribus life satisfaction fall, ceteris paribus? The age dummies in table 2 (where significant, i.e. in columns 1 and 3, the estimates without the time and regional dummies) give a larger negative coefficient for the older age range: -.0.06 for 23 to 24 year olds, -0.07 for 25 an 26 year olds, and - 0.11 for 27 to 29 year olds, in comparison to the base age category of 20 to 22 years old. A separate regression was run with individual age dummies and figure 3 shows that life satisfaction declines, on average, over this age range. The vertical axis represents the coefficient on the age dummy, with age being on the horizontal axis. All of the coefficients were significant. (The estimates are the equivalent of column 1.)

Figure 3 Coefficients on age, in a life satisfaction regression.

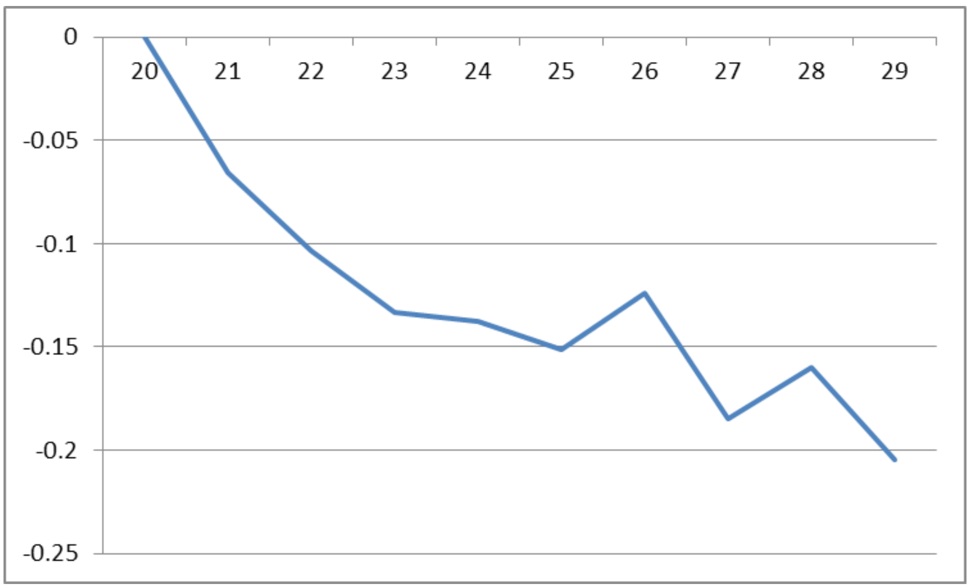

Wooldridge’s (2002) test for serial correlation, implemented in Stata by the user-written xtserial command (Drukker 2003), rejects the null of no first order autocorrelation. This is potentially useful information. One possibility is to recognise the clusters involved in the panel regression and to correct the standard errors accordingly. However this treats the omitted dynamics detected by the diagnostic test as a problem rather than an invitation to respecify the model to include the omitted dynamics in the estimated part of the model and thus to exploit this additional information in estimation. This argument has recently been strongly supported by King and Roberts (2012) in a study of robust standard errors:

> Robust standard errors now seem to be viewed as a way to inoculate oneself from criticism. We show, to the contrary, that their presence is a bright red flag, meaning "my model is misspecified"… it appears to be the case that a very large fraction of the articles published across fields is based on misspecified models. For every one of these articles, at least some quantity that could be estimated is biased (p. 2).[7]

Accordingly, a potentially more interesting solution is to estimate a dynamic panel model (which is undertaken in section 4). The outcome of this is interesting, and provides new insights for the empirical literature on the economics of happiness.

4 Happiness in the twenties: dynamic panel analysis

This section is informed by finding the presence of first order serial correlation in the idiosyncratic error term in the static estimations of section 3. Such a result can mean that the estimates generated by static panel analysis are inefficient and potentially misspecified. Adding dynamics to the model is usually undertaken by including a lag of the dependent variable as a right hand side variable. Hence we are estimating the following standard equation (with the independent variables excluded for clarity):

$$y_{it} = \beta y_{i,t-1} + (\alpha_i + \varepsilon_{it}) \qquad (1).$$

As it is a panel model each observation is indexed over *i* (= 1…N) cross-section groups (here individuals) and *t* (= 1…T) time periods (here, annual observations). Equation 1 is a first-order dynamic panel model, because the explanatory variables on the right-hand side

---

7 "We strongly echo what the best data analysts have been saying for decades: use all the standard diagnostic tests; be sure that your model actually fits the data; seek out as many observable implications as you can observe from your model. And use all these diagnostic evaluation procedures to respecify your model" King and Roberts (2012, p.8).

include the first lag of the dependent variable ($y_{i,\, t-1}$). The composed error term in parentheses combines a group-specific random effect to control for all unobservable effects on the dependent variable that are unique to the individual and do not vary over time (α_i), which captures specific ignorance about individual i, and an error that varies over both individuals and time (ε_{it}), which captures our general ignorance of the determinates of y_{it}. However, this cannot be estimated accurately by OLS or by fixed effects estimation. An OLS estimator of β in equation 1 is inconsistent, because the explanatory variable $\beta y_{i,t-1}$ is positively correlated with the error term due to the presence of individual effects. A fixed effects estimation does not have this inconsistency because the equation is transformed to remove the individual effect, as in equation 2.

$$y_{it} - y_{i,t-1} = \beta\left(y_{i,t-1} - y_{i,t-2}\right) + \left(\varepsilon_{it} - \varepsilon_{i,t-1}\right) \qquad (2).$$

However, this has the different problem of correlation between the transformed lagged dependent variable and transformed error term. The overall impact of the correlations is negative, and is the well-known Nickell (1981) bias. Bond (2002), in a paper that was very informative for the work of this section, states that these biases provide an informal test for an estimator of the lagged dependent variable. He suggests that the estimated coefficient should be bounded below by the outcome from OLS (which gives the maximum upwards bias) but above by the fixed effects estimate (which gives the maximum downwards bias).

Due to these problems, the standard approach is to find a suitable instrument that is correlated with the potentially endogenous variable (the more highly the better), but uncorrelated with ε_{it}. Because instrumentation is not confined to one instrument per parameter to be estimated, the possibility exists of defining more than one moment condition per parameter to be estimated. It is this possibility that is exploited in the General Method of Moments (GMM) estimation of dynamic panel models, first proposed by Holtz-Eakin et al. (1988). The two models popularly implemented are the "difference" GMM estimator (Arellano and Bond, 1991) and the "system" GMM estimator (Arellano and Bover 1995). Greene (2002, p.308) explains that suitable instruments fulfilling the criteria mentioned above come from within the dataset: the lagged difference ($y_{it-2} - y_{it-3}$) and the lagged level y_{it-2}. Both of these should satisfy the two conditions for valid instruments, since they are likely to be highly correlated with ($y_{i,t-1} - y_{i,t-2}$) but not be correlated with $(\varepsilon_{it} - \varepsilon_{i,t-1})$. It is this easy availability of such "internal" instruments (i.e., from within the dataset) that the GMM estimators exploit. The "difference" GMM estimator follows the Arellano and Bond (1991) data transformation, where *differences are instrumented by levels*. The "system" GMM estimator adds to this one extra layer of instrumentation where the original *levels are instrumented with differences*.

These estimators, unlike OLS and conventional FE and RE estimation, do not require distributional assumptions, like normality, and can allow for heteroscedasticity of unknown form (Verbeek, 2000, pp. 143 and 331; Greene, 2002, pp.201, 525 and 523). A more extensive discussion of these methods is beyond the scope of this thesis, but the references provided above and papers by Roodman (e.g. 2006, 2007, and 2009) are very informative.[8] A further advantage of GMM estimation and the use of 'internal'

8

instruments is that applied researchers can select which variables are endogenous and which are exogenous. Here we assume that, regarding life satisfaction, everything is potentially endogenous with the exception of the age, region and wave dummies. This is discussed below.

As a consequence of the lagged dependent variable being estimated (and the internal instruments used), the number of observations will shrink somewhat because two consecutive years are needed. In 2001, i.e. wave 11 of the BHPS, the life satisfaction question was not asked which will mean two years of no data because of the missing lags. Given the necessity of (a minimum of) two consecutive years, the number of observations from dynamic estimates will be smaller than the observations used in the static analysis of section 3

Before estimating any dynamic panel model there are two important considerations. Firstly, are the regressors potentially endogenous or strictly exogenous? And secondly, how many instruments to use? With happiness equations many of the regressors are potentially endogenous – does marriage, for example, make someone happy or are happy people more likely to get married (or are both determined by underlying but omitted variables) – and with this in mind, all of the controls apart from age and the wave and region dummies are treated as potentially endogenous. There is no perfect number of instruments to use, and a general rule is the more the better unless they start becoming too weak. However, there are standard diagnostic tests which can offer guidance with

---

The Roodman papers are particularly useful for applied researchers because they explains how to use the Stata software programme, xtabond2, that he created to implement the GMM dynamic estimators.

respect to the choice of instruments, and the discussion now turns to this aspect of dynamic estimation.

The "system" GMM estimation was undertaken twice, with the difference between the two being only the number of internal instruments the system uses (based on lag length choice), the minimum first and then the default maximum. In both cases the diagnostics indicate that first order autocorrelation is present (p=0.000 for the null of no first-order serial correlation, for both estimates), but second order is not (p=0.925 and p=0.623, for the minimum and default choice of instrumentation). This is expected (and necessary): the difference of lags and the difference of levels are correlated (first order), However, the second differences are not and thus valid for instrumentation. Also the various tests of the instruments indicate that they are suitable in each case – the null hypothesis of exogenous instruments (what we want) cannot be rejected. The Hansen (1982) test *J* statistic[9] of all overidentifying restrictions, with a p-value of approximately 0.37 (for both estimates), does not reject the null of instrument validity indicating an almost forty percent chance of a type one error if the null is rejected. This is higher than Roodman's recommended threshold of a p-value of 0.25 where he (2007, p.10) warns that researchers

> should not view a value above a conventional significance level of 0.05 of 0.10 with complacency. Even leaving aside the potential weakness of the test, those thresholds are conservative when trying to decide on the significance of a coefficient estimate, but they are liberal when trying to rule out correlation between instruments and the error term. A *p* value as high as, say, 0.25 should be viewed with concern. Taken at face value, it means that if the specification is valid, the odds are less than 1 in 4 that one would observe a *J* statistic so large.

9

This has the advantage over the Sargan *J* test because it works in the presence of heteroscedasticity. Indeed, if the errors are believed to be homoscedastic then the Hansen test is the same as the Sargan test.

Thus, for both estimations, we have the necessary second-order serial correlation and the instrumentation is valid. Satisfactory diagnostics means we can now turn to the actual results, though the interpretation of the coefficients is not the same as that for standard FE analysis:

> Adding dynamics to a model … creates a major change in the interpretation of the equation. Without the lagged variable, the "independent variables" represent the full set of information that produce observed outcome y_{it}. With the lagged variable, we now have in the equation the entire history of the right-hand-side variables, *so that any measured influence is conditional on this history*; in this case, any impact of (the independent variables) $\boldsymbol{x}_{it}$ represents the effect of *new* information. (Greene, 2008, p.468)

This quote is instructive for the discussion which follows table 3 below, and also regarding the decision as to whether to model the omitted dynamics more generally. A note on the internal instruments helps to explain the difference between the columns in the table. The problem of 'too many instruments', as identified by Roodman, is particularly problematic in small samples, but is less likely to be a problem here given that we are dealing with more than 13,000 observations. However, there is not yet any known way of choosing the instrument set beyond trial and error. Here we present the results from "system" GMM estimation with the minimum number of instruments (column 1), and the maximum default instrumentation (column 2), and this is the only difference between the estimations generating the results in the two columns.

Table 3 life satisfaction of young British people, assessed via dynamic analysis.

	(1) Life Satisfaction	(2) Life Satisfaction
Number of instruments	229	747
Lagged Life Satisfaction	0.10***	0.08***
	(0.021)	(0.019)
Income	-0.07	0.01
	(0.081)	(0.060)
Self-employed	-0.15	-0.10
	(0.210)	(0.185)
Unemployed	0.09	-0.22
	(0.356)	(0.216)
Other Labour Force Status	0.08	0.05
	(0.159)	(0.128)
Married	0.48***	0.46***
	(0.070)	(0.063)
Separated	-0.33	-0.25
	(0.393)	(0.287)
Divorced	0.19	0.25
	(0.245)	(0.207)
Widowed	-0.40	-2.02***
	(2.853)	(0.626)
Education: High	-0.02	-0.12
	(0.272)	(0.219)
Education: Medium	-0.20	-0.22
	(0.337)	(0.263)
Health: Excellent	1.07***	0.99***
	(0.248)	(0.157)
Health: Good	0.41*	0.50***
	(0.210)	(0.113)
age2324	-0.01	-0.01
	(0.034)	(0.032)
age2526	-0.05	-0.06
	(0.041)	(0.038)
age2729	-0.12***	-0.14***
	(0.047)	(0.044)
Wave Dummies	Yes	Yes
Region Dummies	Yes	Yes
Constant	4.30***	4.41***
	(0.346)	(0.279)
Observations	13,688	13,688
Number of pid	4,823	4,823

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A substantial discussion will follow regarding the lagged dependent variable, its result, and the interpretation of that result, but we note here that it is small and positive (approximately 0.1) and statistically significant. For the estimation with minimum instrumentation, the other statistically significant influences on life satisfaction are marriage and health that is either self-reported as excellent or good (with good having a p-value of 0.052). For the choice of default instrumentation, the results are the same (though good health is now statistically significant at or below the 1% level) with the addition of a negative impact from being a widow. As Greene reminds us, the interpretation of these coefficients is that they represent the effect of new information, whereas the lagged dependent variable reflects the influence of the past. Two changes from the static results of section 3 stand out. Both income and unemployment, when estimated in a dynamic model and treated as endogenous, are not significantly different from zero in their association with life satisfaction. Using GMM methods we can discover whether the reason for this is the dynamic estimation or treating the variables as endogenous.[10] Further estimations reveal that the reason for this loss of significance is not due to the move to dynamic analysis, but because we have chosen to treat both income and unemployment as endogenous. The results from a dynamic estimate treating these two variables as exogenous are presented below in table 4, and a discussion follows.

10 An advantage of GMM estimation when compared to FE is that applied researchers can, even in a static context, instrument for potentially endogenous variables.

Table 4 life satisfaction of young British people, assessed via dynamic analysis (income and unemployment treated as exogenous)

	(1)	(2)
	Life Satisfaction	Life Satisfaction
Number of instruments	229	747
Lagged Life Satisfaction	0.10***	0.08***
	(0.021)	(0.019)
Income	0.08**	0.08***
	(0.033)	(0.027)
Self-employed	-0.18	-0.08
	(0.219)	(0.171)
Unemployed	-0.32***	-0.36***
	(0.092)	(0.083)
Other Labour Force Status	0.09	0.07
	(0.166)	(0.129)
Married	0.43***	0.45***
	(0.069)	(0.064)
Separated	-0.20	-0.32
	(0.373)	(0.282)
Divorced	0.17	0.26
	(0.251)	(0.203)
Widowed	-1.00	-1.94***
	(2.697)	(0.743)
Education: High	0.06	-0.11
	(0.274)	(0.226)
Education: Medium	-0.04	-0.22
	(0.340)	(0.277)
Health: Excellent	1.21***	1.06***
	(0.257)	(0.107)
Health: Good	0.47**	0.54***
	(0.229)	(0.109)
age2324	-0.04	-0.03
	(0.032)	(0.031)
age2526	-0.09***	-0.08**
	(0.036)	(0.035)
age2729	-0.19***	-0.17***
	(0.041)	(0.040)
Wave Dummies	Yes	Yes
Region Dummies	Yes	Yes
Constant	3.99***	4.30***

	(0.330)	(0.273)
Observations	13,688	13,688
Number of pid	4,823	4,823

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

These results show that, when treated as exogenous, the effect of the income of British individuals in their twenties on self-reported life satisfaction is positive and significant, whilst that of unemployment is negative and significant. One question, therefore, is what our assumptions should be regarding the exogeneity or endogeneity of these variables. The economic literature has started to explore the notion that income may promote happiness, but happiness may also lead to increases in income or, as the title of one article has it, happiness may 'pay' (Powdthavee 2009; De Neue and Oswald 2012) but has not really addressed whether unemployment and life satisfaction are endogenous.[11] If they are connected by an omitted but time invariant factor, then FE analysis (though static) can be said to account for it. One possibility for an omitted variable might be depression, but an argument could be that this is itself indicated by low life satisfaction scores. If the endogeneity is caused by simultaneity rather than an omitted variable then static FE methods cannot account for this. It is clear that happiness work must carefully consider the potential relationships between the independent variables and life satisfaction. GMM methods can allow for both dynamic estimation and also potentially

11

It is worth noting that other possibilities were investigated. Interacting, variously, age, income, and gender with unemployment changes neither the result for unemployment, not the results for the other coefficients. Gender is not reported here because the aim of these regressions is to be as close to the estimation to the static FE as possible (where gender cannot be accounted for). However, running these regressions separately for each gender does not change the outcome nor does the inclusion of a gender dummy variable. Further checks were made with the 50s age range, and these also support the results of this paper.

endogenous relationships, and should be more routinely considered in the happiness empirical literature.

The coefficient on lagged happiness in these dynamic estimations is itself interesting and as Greene informs us (see the quote just above table 3), this coefficient represents the 'entire history of the model' i.e. the past history of the process that generates current levels of happiness. This coefficient is positive, suggesting a persistence or inertia effect from previous happiness: lagged happiness being positively associated with current happiness. That the coefficient is small (around 0.1) means that this influence of the past is minor, which demonstrates that what is (largely) important for the determination of current happiness are current circumstances and events. In all cases, dynamic GMM estimation for life satisfaction also passes Bond's (2002) informal test for a good estimator (mentioned above): the coefficient of 0.1 is lower than that obtained by OLS (which is biased upwards) and higher than that obtained by fixed effects (which is biased downwards). We return to this discussion of the approximate 0.1 coefficient on lagged life satisfaction, after discussing a different conjecture in the literature.

Another possible hypothesis has been put forward for the coefficient on lagged happiness. Bottan and Perez-Truglia (2010) assert that past happiness should be negatively correlated with current happiness supporting their notion of what they call general habituation (in contrast to specific habituation which is getting used to a single event like, for example, marriage or a pay rise). Because specific habituation occurs (see section 2 above), they argue that general habituation should occur: we get to adapt to marriage and divorce (see section 2) so perhaps we adapt to happiness overall (an argument that ignores events like unemployment which people do not appear to adapt to (Clark et al. 2008)). In “the very first paper to address [this] empirically” Bottan and Perez-Truglia (2010, p.2) argue for general adaptation in their opening paragraph with this rationale (the oft-found adaptation in the literature to specific incidents means that there is a general adaptation effect too). In a later version of this paper, in support of general adaptation they refer to “what we might recognise as the evolutionary origins of hedonic adaptation. The basic intuition is that positive and negative hedonic states are costly from a fitness perspective: e.g. generating feelings is a waste of energy for the brain. In order to minimise those fitness costs, humans have adapted with hedonic states that quickly return to ‘normal levels’” (Bottan and Perez-Truglia 2011, p.225). A little later on the same page they expand on this idea and suggest that “the reward centers in the brain may work as a spring: i.e. as soon as an individual excites some area in the brain, the corresponding reward system will automatically start pushing in the opposite direction”. Their empirical results, based on panel data from Britain (the BHPS), Germany (the GSOEP), Japan (the Japanese Panel Survey of Consumers) and Switzerland (the Swiss Household Panel) suggest that persistence (or inertia) exists – the sign on the coefficient for lagged happiness is positive and significant and not negative (as they expect) – and that this, they suggest, presents a ‘puzzle’. Another, and more likely, interpretation may be that their initial speculation is wrong, and that general adaptation is not found. They try many different specifications and consistently find a small, but positive impact of lagged happiness on current happiness. The empirical results above, tables 3 and 4, confirm their findings of no general habituation, but instead indicate that there is some happiness 'carry-over' from the previous period.[12]

Returning to our results, we investigate further this small persistence effect. Robustness checks follow supporting this finding of a low level of persistence from past happiness, but before then, and for clarity, it should be emphasised that using the term persistence does not necessarily indicate a fixed effect. The lagged dependent variable aggregates all previous influences, both observed and unobserved. Because these change, so does the value of the lagged dependent variable. Though we have a fairly persistent coefficient on the lagged dependent variable, it is not a fixed effect and hence not something that would be captured by fixed effects estimation. A little algebra demonstrates that this is the case, which just expands the lagged dependent variable to demonstrate that it contains more than just the fixed effect. In equation 3 LS_{it} is the life satisfaction of individual *i* in year *t*, βX_{it} is an independent variable and ϵ_{it} is the usual error term. Starting with our simplified specification in equation 5.3, we repeatedly substitute for the lagged dependent variable.

Substitute for LS_{it-1} in (5.3):

$$LS_{it} = \hat{\alpha}\, LS_{it-1} + \hat{\beta}\, x_{it} + \epsilon_{it} \quad (3)$$

$$LS_{it-1} = \hat{\alpha}\, LS_{it-2} + \hat{\beta}\, x_{it-1} + \epsilon_{it-1} \quad (4)$$

Substitute (5.4) into (3)

$$LS_{it} = \hat{\alpha}\left(\hat{\alpha}\, LS_{it-2} + \hat{\beta}\, x_{it-1} + \epsilon_{it-1}\right) + \hat{\beta}\, x_{it} + \epsilon_{it} \quad (5)$$

12 This paper is worrying for other reasons than the authors' refusal to accept what the empirical results tell them. They do not seem to understand the dynamic models they employ, neither appearing to understand the Sargan test (reporting when they reject the null of instrument validity only, without further comment) nor the meaning of the lagged dependent variable.

Substitute for LS_{it-2} in (3):

$$LS_{it-2} = \hat{\alpha}\, LS_{it-3} + \hat{\beta}\, x_{it-2} + \epsilon_{it-2} \tag{6}$$

Substitute (6) into (5)

$$LS_{it} = \hat{\alpha}\left(\hat{\alpha}\left[\hat{\alpha}\, LS_{it-3} + \hat{\beta}\, x_{it-2} + \epsilon_{it-2}\right] + \hat{\beta}\, x_{it-1} + \epsilon_{it-1}\right) + \hat{\beta}\, x_{it} + \epsilon_{it} (7)$$

Gather terms

$$LS_{it} = \hat{\alpha}^3\, LS_{it-3} + \hat{\alpha}^2\, \hat{\beta}\, x_{it-2} + \hat{\alpha}\, \hat{\beta}\, x_{it-1} + \hat{\beta}\, x_{it}$$

$$+ \hat{\alpha}^2\, \epsilon_{it-2} + \hat{\alpha}\, \epsilon_{it-1} + \epsilon_{it} (7')$$

Going back further than four lags introduces more past values and more idiosyncratic error terms too. By repeated substitution, we demonstrate that through the lagged dependent variable dynamic specifications contain the entire history of the independent variable(S). Thus we cannot assume that the lagged dependent variable represents any kind of fixed effect. We can only say that this variable tells us the influence of the past, and the results above suggest that there is an overhang from the past.

To check whether this is just an artefact of this particular sample a comparison was made with individuals in their 50s over the same time period, and again lagged happiness is statistically significant and positively associated with current happiness. This is the case with both minimum and default instrumentation, and the diagnostics are similar to those for the twenties age range, statistically supporting both specifications. Taken together it seems that a small amount of happiness is persistent over time, and that this persistence is not age dependent (approximately 0.09 for the 20s, 0.08 for the 50s). However, the slightly lower coefficient regarding the influence of past happiness may be consistent with Carstensen's theory (mentioned in section 2) where older people are said to feel more connection to the present. The coefficient on lagged happiness for individuals aged between 15 and 60 is around 0.095, offering yet more support for results found for the 20s and 50s individually.

A further robustness check involves using the alternative to the life satisfaction one-item measure, the General Health Questionnaire. Frequently used as a proxy for psychological well-being, it is available in every wave of the BHPS thus providing 17 waves of data, and avoiding the problem presented by the life satisfaction question not being asked in the middle of the data set. The table below reports the coefficient on the lag of both Caseness and Likert GHQ, when they have, respectively been used as a 'happiness' proxy, for the twenties.

Table 5 How GHQ well-being is influenced by its lag, twenties age range

	Minimum instruments with collapse (58)	Minimum instruments (397)	Maximum instruments (1268)	Sample size
Caseness	0.081***	0.073***	0.047***	Obs = 17537
Likert	0.081***	0.081***	0.060***	Avg obs/group = 3.3

Table 5: How GHQ well-being is influenced by its lag, fifties age range

	Minimum instruments with collapse (60)	Minimum instruments (448)	Maximum instruments (1504)	Sample size
Caseness	0.102***	0.132***	0.121***	Obs = 13126
Likert	0.060***	0.087***	0.108***	Avg obs/group = 4.04

Caseness for the whole age range using minimum instrumentation with 'collapse'[13] – the only one Stata could calculate in less than twenty four hours – gives the average impact of lagged satisfaction as 0.084. This is also highly significant, and the diagnostics are, like all of the others above, as expected offering further support for the finding of a small but persistent effect of past happiness on current happiness. Or, from another angle, that current happiness is largely determined by current events. Not displayed here due to space considerations, but similar life satisfaction estimates with the German SOEP to those undertaken here with the BHPS were made and an approximate 0.1 result is found for the coefficient on lagged happiness. A result that is similarly statistically significant at the 1% level.

One further advantage of dynamic panel analysis is the ability to separately quantify the short-run impact and the long-run effect of various independent variables. The short-run impacts are given by the estimations above, and the long-run effects are calculated by dividing these short-run coefficients by 1 minus the coefficient on lagged happiness. In the case of lagged happiness, this means, in effect dividing each coefficient by approximately 0.9. Because the coefficient on lagged happiness is positive but small, this means that the LR effects (which include the short-run impact) are only slightly higher than the short-run impact. This seems reasonable: the happiness associated with being married in one particular year has very little carryover to long-run happiness. (This, for

---

13

As the Stata help file for *xtabond2* declares, the collapse command greatly curtails the width of the instrument matrix and helps keep the matrix within Stata's size limit. It does this by specifying that *xtabond2* should create one instrument for each variable and lag distance, rather than one for each time period, variable, and lag distance.

example, is unlike increased spending on schools or hospitals and its subsequent impact on a measure of performance.) The results are not reported here because they are not qualitatively different from the short-run impacts estimated in tables 3 and 4. In other areas, where the coefficient on the lagged dependent variable is higher, the long-run impacts can be quite different from the short-run impact, both in terms of size but also significance. Happiness, with its low level of persistence, is likely to be a special case.

One more reason for happiness being a special case is with respect to the argument that dynamics should be modelled as a matter of course. This argument is made by Bond (2002) who argues that "even when coefficients on lagged dependent variables are not of direct interest, allowing for dynamics in the underlying process may be crucial for recovering consistent estimates of other parameters" (p.1; see also p.20). We suggest that for happiness estimates (because of the low influence of the past), where the variables of interest are historic rather than current, the impact will be caught in the 'black box' of the lagged dependent variable, and thus dynamic estimation may not be particularly informative.[14] Yet Bond's general advice is helpful. If serial correlation is found in the idiosyncratic error term, then estimating happiness equations with dynamics should be considered on an investigation by investigation basis.

5 Concluding remarks

14 In this case, to account for the omitted dynamics one can model them in the residual as long as the common factor restrictions hold. For a more detailed discussion, and an example using happiness and education see Piper (2012).

The analysis of this paper both supports and extends recent research. The focus on the twenties is novel, and many justifications both theoretical and methodological have been provided above for the concentration on both a specific age range and this particular range. The static estimations provide support for previous studies, with the qualitative impact on happiness of various controls being similar to that found in many of the studies discussed throughout the paper. The underlying U-shape relationship between well-being and age is supported by the analysis here: happiness declines over the twenties, a result found by comparing the coefficients of the age dummies.

The diagnostic tests on the static estimations, reported in section 3, indicate the presence of omitted dynamics, and thus the analysis here extends previous work by estimating a dynamic panel model in section 4. The central finding is robust: a statistically significant positive coefficient for lagged happiness. There is a clear, if small (approximately 0.09), impact of past happiness on this year's happiness. This finding suggests that there is an element of persistence in happiness, at least for UK data. This is the case for those in their 20s, 50s, and over the whole of the lifecycle. This finding is also robust to an alternative measure of well-being, the General Health Questionnaire. Estimates with the German SOEP also demonstrate a similar effect. Furthermore, another study (Bottan and Perez-Truglia 2011) finds the same result in Germany, Japan, and Switzerland, though as discussed above these authors do not seem to know what it means or implies for happiness. Rather than being a 'puzzle', this result should perhaps be unsurprising: most of what influences current happiness appears to be contemporaneous.

This finding of a robust positive coefficient for the impact of lagged happiness on current happiness suggests that the static models of the previous empirical literature are misspecified, potentially invalidating their results. However the small coefficient on lagged happiness is indicative of current happiness being largely due to current circumstances, which offers some support for the results found in the previous static literature. Indeed, dynamic estimations can offer a robustness check for static regressions. One further benefit of dynamic analysis is the ability to distinguish between short- and long-run effects, and for happiness regressions it confirms what seems reasonable: most of the effect comes from short-run effects (or the current circumstances of the individual). Applied researchers can also decide whether to estimate variables as potentially endogenous or exogenous. The results above suggest that this choice is important, modelling income and unemployment as endogenous results in an influence on happiness that is not significantly different from zero. Much more research is necessary regarding the potential endogeneity of independent variables with respect to happiness and GMM estimation can help with such an investigation.

Dynamic analysis is unusual in this literature, but the analysis of this paper suggests that consideration of dynamic methods should become the default position. There are clear dynamic effects in satisfaction estimations and a failure to model these can, even when the lagged dependent variable is not of interest (as Bond, 2002, highlights), lead to inconsistent estimates of the other parameters.

References

Arellano, M. and Bond. S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations, *The Review of Economic Studies* 58, pp. 277 – 297.

Arnett, J. J. (2000) Emerging Adulthood: A Theory of Development from the Late Teens through the Twenties *American Psychologist* 55, 5: pp. 469-480. (http://www.jeffreyarnett.com/articles/ARNETT_Emerging_Adulthood_theory.pdf)

Bee, H. (1997) *Lifespan Development,* 2nd Edition. Harlow: Addison Wesley Longman.

Blanchflower, D. G. (2001) Unemployment, Well-being and Wage Curves in Eastern and Central Europe, *Journal of Japanese and International Economies* 15, pp. 364-402.

Blanchflower, D. G. (2010) *The Wellbeing of the Young*, University of Dartmouth Working Paper. Available: http://www.dartmouth.edu/~blnchflr/papers/bjir%20youth%20attitudes%20draft.pdf

Blanchflower, D. G., and Oswald, A. J. (2000) The Rising Well-being of the Young in *Youth Employment and Joblessness in Advanced Countries.* Eds Blanchflower, D. G. and Freeman, R. B. National Bureau of Economic Research, pp. 289-328.

Blanchflower, D. G. and Oswald, A. J. (2008) Is Well-Being U-Shaped over the Life Cycle? *Social Science and Medicine* 66, 6: pp. 1733-1749.

Bond, S. R. (2002) Dynamic Panel Models: A Guide to Micro Data Methods and Practice. *Institute for Fiscal Studies / Department of Economics, UCL, CEMMAP (Centre for Microdata Methods and Practice) Working Paper* No. CWPO9/02.

Bottan, N .L. and Perez-Truglia, R.N. (2010) Deconstructing the Hedonic Treadmill: Is Happiness Autoregressive? *Working Paper Series Harvard University.* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1262569

Bottan, N .L. and Perez-Truglia, R.N. (2011) Deconstructing the Hedonic Treadmill: Is Happiness Autoregressive? *Journal of Socio-Economics* 40, 3: pp. 224-236.

Buss, D. M. (2000) The Evolution of Happiness, *American Psychologist* 55, pp. 15-23.

Clark, A. E. (2007) Born to be Mild? Cohort effects don’t (fully) explain why well-being is U-shaped in age, *IZA Discussion Paper* No. 3170, Institute for the Study of Labor (IZA).

Clark, A.E., Diener, E., Georgellis, Y. and Lucas, R. (2008b) Lags and Leads in Life Satisfaction: A Test of the Baseline Hypothesis, *Economic Journal* 118, pp. F222–F243.

Clark, A. E., Frijters, P. and Shields, M. A. (2008a) Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles, *Journal of Economic Literature* 46, 1: pp. 95-144.

Clark, A. E. and Oswald, A. J. (1994) Unhappiness and unemployment, *Economic Journal*, 104, pp. 648-659.

De Neve, J. and Oswald, A. J. (2012) Happiness Pays: Using Sibling Fixed Effects to Estimate the Influence of Psychological Well-Being on Later Income, *Royal Economic Society 2012 Annual Conference paper*, University of Cambridge, March 26-28, 2012.

Di Tella, R., MacCulloch, R. J. and Oswald, A. J. (2003) The Macroeconomics of Happiness, *Review of Economics and Statistics* 85, pp. 809-827.

Drukker, D. M. (2003) Testing for serial correlation in linear panel-data models. *Stata Journal* 3, pp. 168-177.

Gilbert, D. (2006) *Stumbling on Happiness.* London: Harper Press.

Greene, W. (2002) *Econometric Analysis,* Fifth Edition. Upper Saddle River, New Jersey: Prentice Hall.

Hansen, L. P. (1982) Large Sample Properties of Generalized Method of Moments Estimators, *Econometrica* 50, 4: pp. 1029-54.

Kamvar, S., Mogilner, C. and Aaker, J. L. (2009) The Meaning of Happiness *Stanford Graduate School of Business Research Paper* No. 2026.

King, G. and Roberts, R. (2012) How Robust Standard Errors Expose Methodological Problems They Do Not Fix, *Presented at the 2012 Annual Meeting of the Society for Political Methodology*, Duke University.

Layard, R. (2005) *Happiness: Lessons from a New Science.* London: Penguin.

Lelkes, O. (2008) Happiness Across the Life Cycle: Exploring Age Specific Preferences *MPRA Paper* No. 7302 http://mpra.ub.uni-muenchen.de/7302/

Loevinger, B. L. (1976) *Ego Development.* San Francisco: Jossey-Bass.

Loevinger, B. L. (1984) On the Self and Predicting Behaviour, in *Personality and the prediction of behaviour* Eds Zucker, R. A., Aronoff, J. and A. I. Rabin, A. I. New York: Academic Press.

MacKerron, G. J. (2012) Happiness economics from 35,000 feet *Journal of Economic Surveys*, 26, 4: pp. 705-735.

Mogilner C., Kamvar, S. and Aaker, J. (2011) The Shifting Meaning of Happiness, *Social Psychological and Personality Science* 2, 4: pp. 395-402

Myers, D. G. and Diener, E. (1995) Who is happy? *Psychological Science* 6, pp. 10-19.
Neugarten, B. L., Havighurst, R. J. and Tobin, S. S. (1961) The Measurement of Life Satisfaction, *Journal of Gerontology* 16, pp.134-143.

Nickel, S. J. (1981) Biases in Dynamic Models with Fixed Effects, *Econometrica*, 49, pp. 1417-1426.

Powdthavee, N. (2009b) Estimating the Causal Effects of Income on Happiness, *Department of Economics, University of York Discussion Papers* 09/02.

Rindfuss, R. R. (1991) The Young Adult Years: Diversity, Structural Change and Fertility, *Demography* 28, pp. 493-512.

Roodman, D. (2006) How to do xtabond2: An Introduction to “Difference” and “System” GMM in STATA, *Center for Global Development, Working Paper* No. 103.

Roodman, D. (2007) A Note on the Theme of Too Many Instruments, *Center for Global Development, Working Paper* No.125.

Roodman, D. (2009) How to do xtabond2: An Introduction to Difference and System GMM in Stata, *The Stata Journal*, 9(1) pp.86-136.

Verbeek, M. (2000) *A Guide to Modern Econometrics.* Chichester, UK: Wiley.

Wooldridge, J. M. (2002) *Econometric Analysis of Cross Section and Panel Data.* Cambridge, MA: MIT Press.