

Money and Happiness: Evidence from the Industry Wage Structure

Jörn-Steffen Pischke*
LSE

August 2010

Abstract

There is a well-established positive correlation between life-satisfaction measures and income. This paper attempts to provide some evidence on whether this correlation reflects causality running from money to happiness. I use industry wage differentials as instruments for income. This is based on the idea that at least part of these differentials are due to rents, and part of the pattern of industry affiliations of individuals is random. To probe the validity of these assumptions, I compare estimates for life satisfaction with those for job satisfaction, present fixed effects estimates, and present estimates for married women using their husbands' industry as the instrument. All these specifications paint a fairly uniform picture across three different data sets. IV estimates are similar to the OLS estimates suggesting that most of the association of income and well-being seems to be causal.

JEL Classification D1, J31

Keywords Life satisfaction, well-being

*I thank Georg Graetz for excellent research assistance, and Andrew Clark, Larry Katz, Alan Manning, Guy Mayraz, Guy Michaels, Andrew Oswald, Justin Wolfers, and seminar participants for helpful comments.

Money cannot buy you happiness, and happiness cannot buy you money. That might be a wise crack, but I doubt it. **Groucho Marx**

1 Introduction

Economists have taken an increasing interest in the study of life satisfaction data during the past two decades (witness for example the surveys by Frey and Stutzer, 2002 and Di Tella and MacCulloch, 2006). Much of this literature is concerned with the impact of income on happiness (e.g. Easterlin, 1974; Deaton, 2008; Stevenson and Wolfers, 2008a; Layard, Mayraz, and Nickell, 2008). Disentangling causality between these two variables is difficult, something that has not escaped a keen observer of human nature like Groucho Marx. But research has not made too much progress on the causality question to help Groucho out.

This paper tries to provide some evidence on the question of whether “money can buy happiness” using industry wage differentials as instruments for family income. The motivation for these instruments stems from the literature on the reasons for industry wage differentials, which concluded that these wage differentials most likely reflect rents rather than differences in unobserved skills and worker sorting. Workers in a high wage and a low wage industry may therefore not be very different in terms of other characteristics but workers in high wage industries will generally have higher wages, earnings, and family incomes. This approach follows Shea (2000), who has used industry affiliation as an instrument for income in regressions of intergenerational persistence. Using surveys from various countries I find that workers in high wage industries tend to be happier than those in low wage industries. IV estimates of the effect of income are of similar

magnitudes as OLS estimates.

There are various complications with the interpretation of these results. First, even if industry wage differentials mostly reflect rents rather than unobserved skills, there could still be sorting of workers into industries. This could still lead to the same type of reverse causality problem as when simply regressing happiness on income. Secondly, there may be third factors which affect both industry affiliation and hence income, and happiness. A good example are mission oriented industries. These industries offer low pay but might attract relatively happy individuals. More generally, industries may differ in job attributes in addition to the different wages they offer. The theory of compensating differentials suggests that job attributes, which are likely to affect happiness directly, will be systematically related to wages. I will discuss these problems in the next section.

In order to overcome these complications I offer various alternative approaches. I start by controlling for potentially omitted factors. But this approach is also problematic as these controls all tend to be choice variables of the individuals. One particular conditioning approach is individual fixed effects, and studying industry switchers. I also contrast results using life satisfaction and job satisfaction, since I would expect both to be subject to similar biases. An alternative approach is to look at married couples. Here, I instrument the happiness of married women with the industry affiliation of the husbands. With random mating, the wives should not be affected directly by the potential biases identified above. Even with assortative mating, these biases should be reduced in the wives sample. Neither of these approaches suggests that the effect of income on happiness is overestimated in the simpler specifications.

While the happiness literature has investigated many aspects of the re-

relationship between income and life satisfaction, there are only a handful of papers which have tried to address the causality question. One approach is to control for individual effects as in Frijters, Haisken-DeNew, and Shields (2004), who find positive income effects for east Germans after unification. But individual fixed effects are unlikely to be a complete solution as omitted factors may also be time varying. More promising are the attempts by Gardner and Oswald (2007) and Apouey and Clark (2009) to use lottery winners. Both papers find positive effect. In order to produce reasonable sized samples of winners, these studies have to rely on relatively small shocks to income as most wins are small, so that results tend to be imprecise. In a similar vein, Engelhardt and Gruber (2005) look at the effect of the Social Security notch on various variables, including income and happiness. Unfortunately, their estimates are too noisy to be informative.

The closest study to mine is Luttmer (2005), who instruments income with industry x occupation interactions of the respondent and spouse, and finds IV estimates three times as high as corresponding OLS estimates. Luttmer's interest is primarily in comparison income of a reference group, and the instrument for individual income in his paper is not particularly strongly motivated. I will revisit his results below, showing that the occupation instruments are responsible for the high IV estimates. I argue that industry is a more plausible instrument than occupation.

Clark (2003) investigates a somewhat different question. He studies how industry and occupation effects in job satisfaction are related to industry and occupation effects in wages. Clark thinks about these relationships as telling us something about rents. His regressions essentially correspond to the reduced forms I analyze. I find similar results to Clark (2003) when I analyze job satisfaction but different results when analyzing life satisfaction.

Chevalier and Lydon (2002) directly regress job satisfaction on wages. They instrument the wage with the wage of the spouse, finding IV estimates about twice the size of the OLS effects.

Finally, Oswald, Proto, and SgROI (2009) directly investigate the reverse causal channel running from happiness to productivity by studying individual output in a laboratory setting. They use variation in happiness coming from two sources. The first is a manipulation of mood by showing some of the subjects a short comedy clip. These individuals report higher satisfaction before starting a task which takes effort. The second approach uses variation in life satisfaction which is due to adverse circumstances like bereavement or illness of family members. In both cases they find that individuals who report higher satisfaction are more productive on the lab task. While these results are intriguing, it is an open question to what degree they carry over to settings outside the lab.

2 Motivating the instrument

I am interested in estimating regressions of the following form:

$$LS_i = \alpha + \beta y_i + \gamma x_i + e_i$$

where LS_i is a measure of life-satisfaction, y_i is the logarithm of household income, and x_i is a set of possible covariates. This is a happiness equation typically estimated in the literature on cross-sectional data. The worry is that there may either be reverse causality or other factors which affect both life satisfaction and productivity and hence income. Extreme examples of this would be health problems or disability which limit the ability to work. But it is easy to think of less extreme examples, like not having your life together, mid-life crisis, etc. I expect these omitted variables to affect

happiness and income in the same direction, so the coefficient on income in the happiness equation should be biased up.

It is easy to imagine that these factors could be time varying, so individual fixed effects wouldn't solve the problem. Frijters, Haisken-DeNew, and Shields (2004) look at income changes of east Germans after reunification. They argue that post-unification income growth in east Germany was more exogenous than other income changes. But they control for year effects, hence taking out any aggregate income growth. It is unclear why individual level income changes for east Germans should be any more exogenous than those for other individuals.

Finding an instrument for income is an obvious alternative. There are two big challenges in this undertaking. First, data sets with LS variables tend to be relatively small by today's standards in fields like labor economics: a few 10,000 observations. This implies that any candidate instrument needs to move income a lot in order to give any reasonable level of precision. The second challenge lies in the fact that most instruments exploit some more aggregate information about individuals or sub-groups of them. But lots of things including macroeconomic variables could affect life satisfaction directly.

Industry wage differentials are a candidate instrument which has the potential to overcome both problems. These differentials are large: in the US there is a 50 log point spread in household incomes between those working in the top and bottom 2-digit industries. Despite this level of variation, my IV estimates tend to be imprecise, and my instruments are in the weak instrument discomfort zone. This suggests that much more refined instruments, relying on small amounts of variation in the data, are unlikely of any use.

My argument that income variation due to industry wage differentials is informative in this context rests on a series of assumptions. The first is that at least part of the variation in industry wage differentials is due to rents rather than employee characteristics. Many researchers in the 1980s concluded that a large part of the industry wage structure is indeed due to rents (e.g. Dickens and Katz, 1987; Katz and Summers, 1989; and Holzer, Katz and Krueger, 1991; but see Murphy and Topel, 1987 and 1989, for a dissenting view). Even if industry wage differentials only reflect rents, there could be much sorting into industries. The best evidence on this issue is probably the study by Gibbons and Katz (1992) who compare industry switchers who lost their jobs in plant closings (and who are therefore likely to choose their new industry more randomly than those who quit) to other industry switchers. Their results are largely inconclusive.

As a result it is important to consider the possibility that the industry wage structure is correlated with unobserved skills. If happier workers are also more productive this will imply that I should find happier workers in higher paying industries and my IV strategy does not solve the reverse causality problem. To address this issue I compare the estimates for life satisfaction to similar estimates for job satisfaction. If the causality runs from happiness to productivity I would expect the cheerful workers also to report higher job satisfaction. I would therefore expect to see similar effects of income, instrumented by industry affiliation, in equations for life satisfaction and job satisfaction. If the causality runs from income to satisfaction, I would expect an effect on life satisfaction but not job satisfaction if the answers to the job satisfaction question only refer to non-wage aspects of the job.¹ I will show that there is no strong relationship between industry

¹The wording of the questions is not particularly informative on whether individuals will think of the answers as including only non-monetary job attributes or wages as well.

income differentials and job satisfaction. If anything, workers in high wage industries tend to have lower job satisfaction. This seems to be more consistent with a story where the IV estimates truly work through income, and income only affects life satisfaction but not job satisfaction.

Even in the extreme case where industry wage differentials are entirely due to rents, there could still be sorting of individuals into these industries. Holzer, Katz, and Krueger (1991) demonstrate that jobs in high wage industries tend to attract longer job queues. As a result, the employers in these sectors have more choice. Even if applicants look similar in terms of their productive characteristics, employers may prefer to hire more cheerful or happier workers. This would also invalidate the exclusion restriction. In order to address this problem, I consider married men and women. I find similar effects in the subsample of married men. Repeating the same exercise for married women, using their husband's industry affiliation as instruments for their family income, again yields fairly similar results. Even if there is assortative mating on the basis of happiness or if there are spillovers from the happiness of the men on the happiness of the women, I would generally expect these effects to be weaker for the women than the direct effects for the men. The reverse causality stories would therefore suggest more attenuated IV estimates, which is not what I find. This also seems to point towards a story where income is causal.

A particular sorting story has to do with compensating differentials. The jobs in certain industries may be more pleasant than in others. Industry wage differentials may, at least partly, reflect these compensating differentials. The theory of compensating differentials in competitive labor

The GSS question is: "On the whole, how satisfied are you with the work you do?" while the GSOEP question is: "How satisfied are you today with the following areas of your life? [...] How satisfied are you with your job?"

markets suggests that wages and amenities are negatively related in equilibrium. Since amenities should enter life satisfaction this would lead to a downward bias of any income effects on happiness in my IV estimates. Of course, one premise of my investigation is that there are rents in the labor market. Labor markets where firms have wage setting power will generally lead to a tendency of wages and amenities to be more positively correlated. Firms will use both higher wages and more amenities to attract workers (see e.g. Manning, 2003). This leads to a countervailing force to the standard association between wages and amenities in compensating differentials models. As a result, the equilibrium correlation between wages and amenities may be positive or negative. If this correlation is positive, my IV estimates would be biased upward.

Comparing the results for life and job satisfaction should again help in this context as job satisfaction should reflect workplace amenities. The negative association between job satisfaction and income differentials across industries points towards the traditional compensating differentials story dominating (and this is more pronounced in the US data than in the German data). As a result, it seems that the estimates of income on happiness are more likely biased downwards than upwards.

An example of compensating differentials is mission oriented sectors. Some individuals may both be relatively content and not care as much about income but put a lot of emphasis on doing a meaningful job. They will tend to work in mission oriented sectors, which tend to be low paying on average. One interpretation is that workers in these sectors pay a compensating differential for a particular amenity they value.

Figure 1 illustrates that this is likely an important issue in my data. IV estimation with an exhaustive set of dummy variables, like industry dum-

mies, is equivalent to grouping the data by industry (after partialling out covariates) and regressing industry mean happiness on industry mean income. Figure 1 shows this “visual IV” graph (after removing effects due to education, age, and race). It shows clearly that higher paying industries are associated with happier individuals on average (the slope of regression line in this graph is 0.19 with a standard error of 0.10 and corresponds to the estimate in column (1) of Table 1 below). One notable outlier on the top left is “welfare and religious services”. Priests and social workers are among the most content individuals despite the fact that they work in the lowest income industry. Most other industries line up reasonably well around the regression line. Mining and legal services are the highest income industries, and they tend to have very happy workers. However, these two industries are located much above the regression line as well.

Figure 2 repeats the same exercise for job satisfaction instead of happiness. Higher incomes are not related to higher job satisfaction. In fact, the relationship is weakly negative. “Welfare and religious services” are again an outlier. Workers in this sector are highly satisfied with their low paying jobs. This is in line with the competitive version of the compensating differentials model. Ignoring “welfare and religious services,” there is little relationship between income and job satisfaction. This suggests that amenities and compensating differentials probably do not play a large role in biasing the estimates for income. To the extent there is a bias it will tend to reduce my IV estimates.

3 Data

The data come from three sources: The US General Social Survey (GSS), the European Social Survey (ESS) and the German Socio Economic Panel

(GSOEP). The GSS is a relatively consistent repeated cross-sectional survey. I use the 26 waves from 1972 to 2006. The basic life satisfaction measure in the GSS asks “Taken all together, how would you say things are these days—would you say that you are very happy, pretty happy, or not too happy?”, i.e. it allows answers in three categories. In order to make results from this and other measures, which come from different scales, comparable I follow an approach of van Praag and Ferrer-i-Carbonell (2008), and transform the categories to the means implied by an ordered probit fitted to the raw sample fractions. All regressions are then run on the transformed values (van Praag and Ferrer-i-Carbonell refer to this as probit-adapted OLS). Other scalings give qualitatively very similar results. The question on job satisfaction allows four answers, which I transform in the same way.

My second main regressor, family income, is also a bracketed variable. I assign midpoints to the brackets as in Layard, Mayraz, and Nickell (2008). I use Hout’s (2004) conventions for the assigning values to the top bracket, and use a value of \$187,500 for 2006.

The main sample consists of employed men aged 20 to 64. I use a consistent sample without missing values for marital status, education, occupation, industry, happiness, or job satisfaction. Like Stevenson and Wolfers (2008b), I delete the black oversamples and the Spanish language sample, and I weight all regressions by the weight WTSSALL. The male sample has 12,121 observations.

I use the industry and occupation affiliation of the respondent. These are coded in 3-digit codes following the 1970 and 1980 Census classifications. In order to reduce the number of instruments and controls, I aggregated the industries into 33 and the occupations into 22 consistent categories. Details are given in the appendix. For the wives sample, I selected married females

aged 20 to 64 with employed husbands. For the wives sample, the industry and occupation variables I use refer to the husband, while other controls (age, education, race) refer to the respondent.

The ESS is also a repeated cross-section survey which has been carried out in 2002, 2004, 2006, and 2008 in 24 European countries. The set of countries differs somewhat between waves of the survey so that there are between 15 and 20 countries in each particular wave. One advantage of the ESS is that it has two well-being questions, one on happiness and one on life-satisfaction. The former is similar to the question in the GSS and the latter to the one in the GSOEP. The happiness question is “Taking all things together, how happy would you say you are?” while the life satisfaction question asks: “All things considered, how satisfied are you with your life as a whole nowadays?” The answers to both questions are on an 11 point scale.

Household income is also bracketed in the ESS. I follow Layard, Mayraz, and Nickell (2008) again in assigning bracket midpoints. For the lowest bracket, I assign $2/3$ of the bracket boundary, and for the highest $3/2$ of the bracket boundary. Industries in the ESS are classified according to NACE codes (revisions 1 and 1.1). I grouped the NACE categories into 30 sectors, which are relatively commensurate with the GSS classification, although a precise correspondence is not possible. Similarly, I grouped occupations, which are classified according to ISCO-88 into 22 groups commensurate with the GSS. Details are in the appendix.

The GSOEP is a longitudinal survey of households in Germany that started in 1984. The well-being question is similar to the life-satisfaction question in the ESS: “How satisfied are you with your life, all things considered?” Answers are on an 11 point scale, also as in the ESS. Unlike in the

GSS and ESS, household income is a continuous variable. I use the same sample selection criteria as in the other data sets. I also restrict the sample to west Germans (eliminating foreigners), drop the high income oversample, and the East German respondents who were added after German unification in 1991. Satisfaction for east Germans might have been influenced by many factors special to the transition period during the initial phase after unification.

Because of attrition, item non-response, and refreshments, the panels I obtain are unbalanced. They consist of individuals aged 20 to 64 who were interviewed between 1984 and 2007. The pooled male sample consists of 58,120 observation on 9,310 individuals. For the wives sample, I matched employed husbands with their wives to obtain a sample of 6,059 married couples with 37,815 observations. Regressions are weighted by the individual cross-sectional sampling weights PHRF. I use the two-digit CNEF equivalent industry variable provided in the GSOEP which has 28 categories. Occupation is also coded according to ISCO-88, and I recode it similar to the ESS into 22 categories.

4 Results

Table 1 displays the results for the sample of men from the GSS. Each column shows a different specification, and each specification is estimated by OLS, 2SLS, and using Akerberg and Devereux's (2009) Improved Jackknifed IV estimator (IJIVE). All estimates control for a basic set of covariates: age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Controlling for these variables is relatively innocuous as they are not choice variables except for education. Controlling for education is important as more educated workers are typically more satisfied with both

their jobs and their life, and they tend to work in particular industries. Oreopoulos (2007) uses compulsory schooling laws to investigate whether the relationship between life satisfaction and education is causal. He finds little evidence for a bias. Hence, education appears to be a valid control in this regression.

The dependent variable in columns (1) to (4) is the ordered probit scaled trichotomous happiness variable. The OLS estimate of $\ln(\text{family income})$ is 0.16, i.e. a 10% increase in income is associated with an increase of 1.6 percent of a standard deviation of happiness.²

The first stage of the IV regression involves running family income on 31 industry dummies as well as the covariates. This is a fairly substantial number of instruments. Although the industry wage structure predicts wages and family incomes well, there is a concern about weak instruments in these regressions. I present the F -statistic for the excluded instruments at the bottom of each column. In column (1) the F is 13.3, which is somewhat borderline (Stock, Yogo and Wright, 2002). I therefore repeat all the 2SLS estimates using IJIVE. This estimator tends to have much better bias properties and coverage in small samples with weak instruments as illustrated by Akerberg and Devereux (2009). They find that the estimator performs similar to LIML or better. I am using IJIVE instead of LIML here because the estimator is linear, and specifically accommodates large covariate sets and fixed effects, which are difficult to use in LIML. The IJIVE estimates are generally very similar to the 2SLS estimates although standard errors are slightly higher.

The IV estimates in column (1) are very similar to the OLS estimates, slightly above 0.16. However, the IV estimates are not particularly precise

²Running an ordered probit model on the original happiness variable directly controlling for all the covariates yields a slightly higher estimate of 0.218.

given the sample size and the strength of the instrument. While the 2SLS and IJIVE estimates are significantly different from zero, the confidence intervals include both values substantially below and above the OLS estimate.

Column (2) controls for four marital status dummies, a covariate used by many researchers in the happiness literature. But marital status is certainly a more dubious covariate. There is a reverse causality problem as happier men may be more attractive to women and hence more likely to be married (Stutzer and Frey, 2006). Moreover, higher earnings due to a more favorable industry affiliation may affect marriage prospects or marital stability, so that marital status would be a “bad control” in the language of Angrist and Pischke (2009, section 3.2.3). I include this specification for comparability with the previous literature. I will also rely on subsamples of married individuals below, which are, of course, subject to similar problems. While both the OLS and IV estimates in column (2) are lower the main conclusion is that OLS and IV results are still very similar.

The next set of controls entered in column (3) are 21 occupation dummies. The motivation for this is that industry may proxy for many personal attributes beyond the basic regressors like age and education. Occupation should pick up a lot of this variation. On the other hand, occupation is a choice variable of the individual and similar concerns as for marital status apply. The OLS estimates are very similar to column (2) and the IV estimates are somewhat larger again. It is difficult to know why the IV estimates go up in this case. A large part is probably that occupation controls for some of the effects like happy social workers and priests working in low paying industries. Since the precision of the estimates is not that high this change may well also be due to sampling variability. Most importantly, there is no evidence that the previous IV estimates were too high because

of omitted attributes of the individuals.

Industry affiliation is an outcome of the matching process between workers and employer. As a result, the IV estimates may still be subject to reverse causality or third factors if industry affiliation is correlated with other relevant attributes. For example, happier workers may be more productive, and employers in high wage industries may be more likely to hire these more productive workers (or, because they have a greater pick of the crop, they may simply want to enjoy the more cheerful colleagues). If the IV results are due to happier workers being more productive I would expect this effect also, or particularly, to reveal itself when I look at job satisfaction, rather than life satisfaction.

Results for job satisfaction are in columns (5) to (7). These are very different from those in columns (1) to (3). In particular, the OLS results are similarly positive but the IV estimates are now negative, mirroring the findings by Clark (2003). Although the first stage is the same, 2SLS and IJIVE estimates are more dissimilar, often an indication for problems with weak instruments. But IJIVE estimates should still be approximately unbiased. However, the IJIVE estimate in column (5) is strongly negative and actually significant (although even IJIVE standard errors may be too small with weak instruments). Including occupation dummies in column (7), the estimate remains negative but becomes insignificant. The picture painted by these results is certainly very different from that emerging from the results for the happiness variable. Higher paying industries do not seem to be associated with higher job satisfaction.

Under the assumption that job satisfaction summarizes job amenities (but not wages), a simple way to deal with compensating differentials is to include job satisfaction as a control in the happiness regression. This is

done in column (4) and is subject to the same caveats as with the earlier covariates. Given that the industry level correlation between incomes and job satisfaction is negative it is not surprising that using job satisfaction as a control raises the income coefficients in the happiness regressions, albeit only slightly from the specification in column (3) which already controls for occupation. Basically, occupation and job satisfaction seem to play a very similar role here.

How are the results with industry dummies as instruments related to the findings by Luttmer (2005)? Luttmer used industry and occupation interactions as instruments, and found large effects of income. Table 2 replicates these types of results with the GSS data. Column (1) repeats the baseline results from Table 1. Column (2) replaces industry with occupation as instruments. The occupation instruments generate IV estimates about twice the size of OLS. This is may be not surprising in light of the findings by Krueger and Schkade (2008) that higher earnings occupations tend to be more pleasant. But it seems unlikely that this reflects the causal effect of income rather than simply the sorting of workers into occupations. Column (3) uses the actual Luttmer instrument, the interaction of industry and occupation. One problem here is that there are a lot of industry-occupation cells, and hence a lot of instruments. As a result of the large degree of overidentification, the instruments tend to be very weak. While the 2SLS estimate is lower than in column (2), this is less true for the IJIVE estimate, indicating that 2SLS may be biased downward. The conclusion from this exercise is that it is occupation which causes the higher estimates. I consider occupation a much more dubious instrument than industry.

The results so far mostly show a fairly consistent pattern and point towards a story where income is indeed causal. But it is important to probe

further whether the association between higher income due to industry wage differentials and happiness is not driven by some other factor which affects both happiness and industry affiliation. I therefore turn to a sample of married women with employed husbands in Table 3. Here, I instrument the family income for these women by the industry affiliation of their husband. Even if there is some relationship between the industry of the male and that man's happiness, this should not fully spill over to his wife. Even with some spillovers, and with assortative mating I would expect any bias for the women to be less than it is for the men. On the other hand, family income is strongly influenced by the husband's earnings, and industry affiliation continues to work as instrument in the female sample.

Table 3 starts by repeating the regressions for happiness from Table 1 for the subsample of married men. This sample is smaller and the instruments are slightly weaker now, with first stage F -statistics of around 6 to 9. Consequently, LJIVE standard errors are quite large but the estimated coefficients on income are not terribly different from Table 1. Columns (3) and (4) display the results for wives. Notice that the GSS samples individuals, so these are not the wives of the men in columns (1) and (2) but a random sample of married women with working husbands. The instruments are actually slightly stronger in the wives sample, with F s of 9 to 12. The OLS coefficient on income for women is only about 60% the size of the coefficient for men. The IV coefficients tend to be a bit lower yet.

The remainder of Table 3 splits the wives sample into those working, and those not employed. I would like to rule out that any of the results for women are driven by the employment pattern of the wives. For example, one worry might be that couples are more likely to work in the same industry as two random individuals, and even the results for wives still reflect some

other relationship between industry affiliation and happiness. In columns (5) and (6), instead of controlling for covariates for the wives, I look at the subsample of non-working wives. The instruments are yet weaker and standard errors higher than before in this smaller sample. The IV results for this particular subsample look close to zero, and are definitely much lower than the OLS results (which are not changed much from columns (3) and (4)). This is one of the few specifications which suggests little effect of income on happiness.

Columns (7) to (10) in Table 3 look at the sample of employed wives. Columns (7) and (8) replicate the specifications from columns (3) and (4) on this subsample. The results do not look fundamentally different from those for the husbands in columns (1) and (2). Together with the results for non-working wives this might actually suggest that the results could be driven by marital sorting, the fact that couples work in similar industries, and non-income industry attributes could be driving the results. However, in the working wives' sample we can control for the wives' industry affiliation directly. This is done in columns (9) and (10). F -statistics are remarkably unaffected by this (although they are very low); in fact only about 15 percent of women work in the same two digit industry as their husband. The IV estimates are higher than in columns (7) and (8), not lower. This suggests a different story than the non-working women. In fact, the result seems more consistent with a causal effect of income and the result for non-working women being an outlier.

Table 4 presents similar results with data from the European Social Survey (ESS). The ESS has two satisfaction variables, happiness and life satisfaction. Both are recorded on an 11 point scale, much finer than in the GSS. But the ESS has also many drawbacks. Income is recorded on the

same 12 bracket scale for all countries. Because the sample countries differ in their average income levels, and the brackets span the income scale across countries, the within country variation in income is small. Spouses' industry is not recorded, so the analysis here is limited to the male sample.

Columns (1) to (3) in table 4 present the regressions for happiness. The OLS coefficient is 0.22, a bit higher than in the US GSS. The IV coefficients in column (1) are slightly below the OLS result, the opposite as with the GSS. The instruments here are yet weaker, with an F of only 9. While the regressions control for country x wave effects, I pool the estimates for all countries and waves and only use one set of industry dummies to gain more power. The rationale for this is the fact that the industry wage structure is very similar across countries (see Katz and Summers, 1989).³ Despite the low F , IJIVE is very close to 2SLS.

Controlling for occupation in column (3) further weakens the power of the instruments and the IV estimates are lower than OLS. In addition to the non-working wives in the GSS, this is yet another specification where the IV estimates are low. I would also discount this result somewhat because it is not replicated with the life satisfaction variable instead of happiness in columns (4) to (6). For life satisfaction, the IV estimates change relatively little depending on the specification. One likely explanation of the column (3) results may simply be sampling variation.

Table 5 presents results for west German men from the GSOEP. The German question refers to life satisfaction, and is most similar to the life satisfaction question in the ESS. The OLS estimate for Germany of 0.33

³The individual country samples are rather small, and using country specific industry effects would result in too large a number of instruments. There is does not seem to be enough of a simple pattern to industry wage effects across countries which could be exploited for some more parsimonious specification allowing for some heterogeneity across countries.

is about a third higher than the ESS estimate.⁴ Yet more curious, the IV estimates are about twice the size of the OLS estimates in column (1) with the baseline set of covariates. This pattern remains when other covariates are added in columns (2) to (4). The F -statistic for the German data is very low, in the range of 3 to 4, which sheds some doubt on the results. One reason for this might be the fact that industry wage differentials tend to be narrower in more regulated labor markets like Germany than in the US.

The pattern for job satisfaction is again different, just like in the GSS. However, the IV estimates on job satisfaction are positive now rather than negative, although they are much closer to zero than the life satisfaction estimates. One explanation for this pattern of results is that frictions are more important in the German labor market than in the US. Workplace amenities and wages are therefore positively correlated, and the IV estimates for life-satisfaction in Germany are biased up. For the US, the standard compensating differentials story holds, and the income effects for the US are biased down. A puzzling aspect of this explanation is that Germany seems to be different from the other countries in the ESS although other European countries seem to have labor market institutions closer to those in Germany's than in the US.⁵

One of the main advantages of the GSOEP is the longitudinal nature of the data. This allows the introduction of individual fixed effects. While the fixed effects absorb a lot of variation in the life satisfaction variable, it is an open question whether the remaining variation in industry affiliation due to

⁴The OLS estimate on the ESS subsample for Germany is 0.352 (0.035), which is slightly above the GSOEP estimate. The ESS estimate excluding Germany is 0.211 (0.017).

⁵It is worth noting that the ESS includes a fair number of eastern European countries, which tend to have relatively flexible labor markets, as well as the UK and Ireland. Looking at subgroups of countries doesn't suggest any particular pattern according to labor market institutions.

industry switchers is more random than the cross-sectional variation. In any case, the fixed effects estimates, shown in Table 6, are a useful comparison. F -statistics are between 1 and 2, and standard errors are high. Both OLS and IV results for life satisfaction tend to be slightly lower than in the pooled estimates in Table 5 but the basic pattern is still the same. IV results are still well above the OLS estimates. The results for job satisfaction change more: the IV coefficient is now negative. Curiously, the fixed effects results for job satisfaction mirror the cross-sectional GSS much more than the cross-sectional GSOEP results.

Table 7 turns to the results for husbands and wives. The specifications are comparable to those for the GSS in Table 3. The results on husbands look very much like those for all men, and all the results for wives are very similar to those for husbands. In particular, there is no evidence in the German data that the IV coefficients are low for non-working wives, as in the GSS. This suggests that the GSS results may be an outlier. I don't show any fixed effects estimates for these smaller samples as these are simply too imprecise to be useful.

Despite some lack of precision a somewhat consistent story does emerge from the different samples and specifications. IV results are typically similar to OLS or larger. The few cases which are different seem to be outliers rather than systematic deviations, something that is fully expected given the sizeable confidence intervals.

5 Discussion and other instruments

The results presented in the previous section suggest that the OLS estimate might indeed be close to the causal effect of income on happiness. In this section I will explore whether this interpretation is consistent with other

findings in the literature, consider other potential instruments, and alternative explanations for the IV results.

I consider four main issues. First, I investigate some alternative instruments for income. Other employer wage differentials are prime candidates, particularly those due to union status and firm size (union status is also part of Shea's, 2000, instrument set). I will discuss results using these instruments in the GSS, which are noisy but in line with the earlier industry results. Another possibility is to exploit the changes in wage inequality which have taken place in the US over the past 40 years. These changes have also been exploited in the literature on intergenerational transmission, see Acemoglu and Pischke (2001), while Stevenson and Wolfers (2008b) document inequality trends in the GSS happiness data. I link these trends to income, and try to argue that it is more difficult to learn anything on the question at hand from the GSS data.

The second issue is the relationship between the results presented here, and those found in the literature using cross-country regressions, or the within country time series evidence. I will argue that aggregating to the country level is IV, and hence these results should form a comparison for the IV estimates presented here.

The third issue is that permanent income may matter more for life satisfaction than transitory income. The IV estimates isolate relatively permanent differences in income, and this may lead to larger estimates. I address this issue directly by using income over multiple years in the GSOEP, which is a panel.

Finally, I will take up the issue of how to interpret the IV results presented earlier if relative, rather than absolute income matters, as has been argued in much of the literature on happiness. While it is beyond the scope

of this paper to disentangle the effects own income and that of comparison groups, I will show that the presence of relative income in a happiness equation does not lead to an upward bias in the estimates I present.

Start with income differentials due to union status and firm size. I have argued that individuals end up in different industries at least to some degree by chance. The same argument can be applied to firms of different size and union status. Hence, these variables have as much claim as valid instruments as industry. The GSS asks about union status and firm size in some but not all years. As a result an analysis with these variables has to rest on much smaller samples. Union status and firm size, although they predict income well, are not quite as powerful as industry. Consequently, standard errors tend to be about twice the size of those in Table 1. Results are displayed in Table 8. Although the estimates are more variable across specifications, the estimates are also clustered around the OLS value, corroborating the previous findings using industry affiliation.⁶

The changes in wage inequality in the US have led to pronounced changes in the family income distribution. There are different ways of exploiting these changes in the present context. One way would be to use the changing returns to education as an instrument for family income, controlling for education and year main effects. But much of the change in wage inequality happened within narrow groups. An alternative is therefore to exploit the changes in overall wage inequality: Acemoglu and Pischke (2001) show that this is most easily done by controlling for families' rank in the income distribution.

The paper by Stevenson and Wolfers (2008b) on happiness inequality

⁶IV estimates in the sub-sample using both union and firm size instruments tend to significantly larger than the OLS estimates, and larger than in the union or firm size only samples. However, IV estimates are also higher in this sub-sample when just one of the instruments is used in isolation.

essentially provides the reduced forms for both these exercises, and the literature on wage inequality the first stages. As a result, it is simple to put together what to expect from these exercises. Stevenson and Wolfers find that overall happiness inequality has decreased since the 1970s, while education differentials have widened. This suggests that using changing returns to education as an instrument would lead to positive estimates of the income effect, while using overall inequality would rather lead to zero or negative estimates (the latter has also been noted by Layard, 2003).

For completeness, I report results of these two exercises in Table 9. The first line simply reports the standard OLS estimate. The second line reports OLS estimates controlling for rank in the income distribution. This is effectively the same as using the interaction of income rank and year as an instrument for income. The estimates are positive but much smaller than OLS. Of course, the happiness literature has long stressed the potential importance of relative income comparisons (e.g. Luttmer, 2005). No student of this literature will therefore be surprised to find rank to do most of the heavy lifting in these regressions, with less of an effect left for the level of income. Nevertheless, the point estimates are still small and positive, maybe surprising in light of the Stevenson and Wolfers (2008b) findings. They do a more careful job in making sure the GSS satisfaction measure is consistent over time, something that might matter for this particular exercise exploiting time series variation.

The next two rows show the estimates using the interaction of survey year and years of schooling as an instrument for income. As expected, this yields estimates larger than the OLS ones. Nevertheless, it works poorly in the GSS data because there is basically no existent first stage here. The GSS annual samples are small, returns to schooling bounce around from

year to year in these samples but don't show a pronounced upward trend. One could presumably make some headway on this by using better data on income, for example from the CPS in conjunction with the GSS reduced form. I don't pursue this particular avenue here since the main point of this exercise is to illustrate that wage inequality trends are difficult to exploit in this context, and one can obtain different answers depending on the component of inequality used. Clearly, in order to pursue either of these strategies it is necessary to think more carefully about the potential role of relative income.

A large literature discusses country level correlations between income (or GDP) and happiness, and similar estimates for single countries over time, most notably a series of papers by Richard Easterlin (1974, 1995). Aggregating income to the country level is the same as instrumental variables estimation with a set of country dummies as instruments. A similar argument applies to the within country time series regressions. Arguably, country dummies (or year dummies) are useful instruments in this context. I would expect most of the reverse causality or omitted variables bias issues to operate at the level of the individual. Aggregating will filter out that variation. Differences in GDP per capita seem unlikely to be due to the fact that some countries have happier and hence more productive citizens, although this possibility cannot be ruled out completely. Hence, country level studies form another natural comparison for my results here.

The most comprehensive and thorough analysis of the currently available cross-country data is Stevenson and Wolfers (2008a). They find that the income coefficients using individual level micro data are remarkably similar in size to the estimates from the cross-country regressions. In a large sample of countries, including poorer ones, they also find that the within

country time series relationship mirrors the cross-country and micro data estimates. They also analyze a number of single country experiences, particularly Japan and the US. Taking account of breaks in the Japanese data, they conclude, unlike previous authors, that the period of high growth in Japan was mirrored by increasing life-satisfaction, followed by a leveling off in both variables after 1980, and a decline in life-satisfaction recently. Like others, they conclude that there was no increase in happiness in the US over the past 35 years. Stevenson and Wolfers (2008a) point out that standard errors for single country time series estimates are large. As a result, they do not want to draw strong conclusions from the US time series evidence alone.

Other authors, e.g. Layard, Mayraz, and Nickell (2009) argue strongly for the value of the US experience, and debate some of the conclusions to be drawn from the evidence for European countries. This argument focuses on a few small subsamples because these are particularly pertinent for the debate about the relevance of relative income comparisons. Easterlin and Angelescu (2009) find little evidence of a time series relationship of income and happiness in a sample of 37 countries with relatively long time series. They stress the difference between the impact of long run income growth and fluctuations due to the business cycle and argue that the long run growth seems to bear little relation with happiness while the business cycle does, a point also raised by Layard, Mayraz, and Nickell (2009). Nevertheless, most of the macro evidence also seems fairly consistent with my results. The main deviation, just as in the inequality case, may be due to some of the time series variation. Whether this implies that happiness is related to relative income is an issue which remains open to debate.

Instrumental variables estimates invariably isolate more permanent com-

ponents of income. Consumption smoothing by forward looking individuals implies that the distinction between current, permanent, and transitory income may be quite pertinent for the evaluation of income differences. It is also unclear exactly what time horizon survey respondents use when answering well-being questions. Most happiness questions actually mention that they refer to “these days” or “nowadays.” Some of the psychology literature also raises issues as to whether the answers to these questions are actually unduly influenced by current circumstances and mood (Schwarz and Clore, 1983, Schwarz, 1987). Whether transitory or permanent income matters more for life-satisfaction, at the end of the day, is an open empirical question.

Stevenson and Wolfers (2008a) also worry about the permanent income issue. They look at results using education as instrument for income in large cross-country samples, arguing that education will proxy for permanent income. They find IV results which are about 50-100% higher than OLS in the Gallup World Poll and Pew Global Attitudes Survey. In the World Value Survey, on the other hand, IV results are about the same size as OLS estimates. Of course, it is far from clear whether education is a useful instrument in this context as education may possibly affect life-satisfaction directly (I use it as a control variable instead!). It is difficult to know whether the education effect works through income, and permanent income is more important than transitory, or whether there is a direct role of education in affecting happiness.

Oswald and Powdthavee (2008), using panel data from the British Household Panel Survey, instead instrument income with lagged income, which seems preferable to education. Similar to Stevenson and Wolfers (2008a), they find a 50-70% increase in the income coefficient. I follow a similar ap-

proach here by using averages of income over three years as the instrument. Table 10 presents the results. This table has three main rows. The first row shows OLS results. The second row shows 2SLS estimates, using three year average income as instrument for current income. The third row shows 2SLS estimates using industry dummies as instruments for income as before. I repeat these estimates because the sample here is smaller than before (only men with three consecutive observations contribute). The upshot from the table is simple: the 2SLS results using industry dummy instruments are well above the OLS results, while the 2SLS results using average income as instruments are only 5-10% higher than the OLS estimates. To the degree that the IV estimates are above the OLS estimates there is little evidence in the GSOEP data that this is because permanent income matters more for life satisfaction than current income.

Much of the debate in the happiness literature centers around the importance of relative income, as this is one of the possible explanations for the absence of a secular increase in happiness in the most advanced economies. The empirical evidence I have presented here is not particularly well suited to speak to this issue. Nevertheless, it is important to consider how the presence of relative income concerns might affect the interpretation of the IV estimates.

Suppose we are interested in the following regression

$$LS_{ig} = \alpha + \beta y_{ig} + \gamma \bar{y}_g + \delta w_{ig} + \varepsilon_{ig} \quad (1)$$

where $g \in G$ denotes some group, \bar{y}_g is average group income, and w_{ig} is a confounder. Our conjecture is $\beta > 0$, i.e. income has a causal effect on life-satisfaction, and $\gamma < 0$, relative income comparisons matter. The regression I run is instead

$$LS_{ig} = a + by_{ig} + e_{ig}. \quad (2)$$

So the question arises, if we are interested in regression (1) and we run (2), what does this imply for the comparison of OLS and IV applied to (2)? Or to put it differently, could my IV results be biased up because of relative income concerns?

There are two sources of the OLS - IV difference. The first is the presence of $\gamma\bar{y}_g$ in regression (1). That is the confounder we are interested in here. The second is δw_{ig} , the standard omitted variables bias problem. While this is the focus of the current paper, I will shut down this channel by setting $\delta = 0$ for the following discussion in order to highlight the effect of relative income.

Start with the OLS regression. The population regression coefficient from running (2) is

$$b_{OLS} = \beta + \gamma \frac{\text{cov}(\bar{y}_g, y_{ig})}{\text{var}(y_{ig})} < \beta.$$

Also note that

$$\frac{\text{cov}(\bar{y}_g, y_{ig})}{\text{var}(y_{ig})} < 1.$$

This comes from the fact that individual i is member of group g , so that $\text{cov}(\bar{y}_g, y_{ig}) = \text{var}(\bar{y}_g) < \text{var}(y_{ig})$.

Compare this to the IV estimator using an instrument z_s which varies at an aggregate level $s \in S$ possibly different from g . The population IV coefficient is

$$b_{IV} = \beta + \gamma \frac{\text{cov}(\bar{y}_g, z_s)}{\text{cov}(y_{ig}, z_s)}.$$

In order to analyze this, consider two cases. In case 1 the partition S is the same or a coarser set than G . In this case $\text{cov}(y_{ig}, z_s) = \text{cov}(\bar{y}_g, z_s)$ and hence

$$b_{IV} = \beta + \gamma < b_{OLS}.$$

Now consider case 2 that S is finer than G . For simplicity, I will also assume that S partitions G . Then we have $cov(\bar{y}_g, z_s) \propto var(\bar{y}_g)$ and $cov(y_{ig}, z_s) \propto var(\bar{y}_s)$ so that

$$b_{IV} = \beta + \gamma \frac{var(\bar{y}_g)}{var(\bar{y}_s)}.$$

With $\gamma < 0$, and since

$$\frac{var(\bar{y}_g)}{var(y_{ig})} < \frac{var(\bar{y}_g)}{var(\bar{y}_s)} < 1,$$

it follows that $b_{IV} < b_{OLS}$.

As a result $b_{IV} < b_{OLS}$ independent of whether S or G is coarser. In either case, IV affects the covariance involving y_{ig} in the denominator more than the already smaller covariance involving \bar{y}_g . Hence, if relative income matters, my IV estimate should be smaller than the OLS estimate. Moreover, both will be underestimates of β . Finally, the difference between OLS and IV should be large when the comparison groups are larger groups and vice versa.

Mayraz, Wagner, and Schupp (2009) analyze data from the German SOEP where individuals were asked directly whom they compare themselves with. They find that average income of the same sex, individuals in the same profession, and co-workers are the main comparison groups. Clark and Senik (2010) similarly finds that co-workers are the main comparison group using data from wave 3 of the ESS. Since my regressions are for either men or women, and include (country-) year effects, any comparison group at the national level would be absorbed. The findings by Mayraz et al. (2009) imply that comparison groups may be relatively broad. As a result, I should find IV results which are substantially smaller than OLS. This is not systematically the case. I conclude that relative income comparisons alone

do not explain the pattern of the IV results particularly well. Moreover, if we are interested in the presence of an omitted variable biasing the OLS income coefficient up, then the IV results provide even stronger support for a causal effect of income in the presence of relative income concerns.

6 Conclusion

This paper attempts to provide some evidence on whether the cross-sectional association between well-being measures in survey data (happiness or life satisfaction) and family income is causal or not. Compelling and viable instrumental variables for income are hard to come by in this setting because the data sets with life satisfaction questions are small, and the idiosyncratic variation in happiness is large. As a consequence, the results presented here are suggestive at best.

I have used industry wage differentials as predictors for family income. While industry differentials are large, particularly in the US, they are also far from ideal as an instrument in this context. This is highlighted, for example, by the outlier industry “welfare and religious services” which is poorly paid but has very happy workers. This combination is likely due to the sorting of a very particular group of individuals into this sector. The approach in this paper rests on the assumption that sorting of this type is the exception, and industry affiliation is to a large part actually more or less random. In order to probe this assumption, I push the data in various directions: with the comparison of life satisfaction versus job satisfaction results, using individual fixed effects, and looking at the life satisfaction of the wives, using husbands industry as the instrument. While IV standard errors are large, and the results bounce around to some degree, to me at least, they seem to be pointing in a remarkable consistent direction. The

IV results tend to be very similar to the OLS results, and for the most part not smaller.

Of course, the similarity of 2SLS and OLS estimates could simply stem from the fact that overidentified 2SLS estimators with weak instruments are biased towards OLS. However, IJIVE estimates, which tend to be much better behaved point in the same direction. The results are also fairly consistent with the existing literature on this topic, including the cross-country studies, which I argue can be interpreted as useful IV studies in this context. The results are also not simply explained with the IV estimates proxying for relatively permanent income, or by the presence of relative income comparisons.

None of the individual IV estimates presented in this paper are very precise. But it is important to keep in mind that the estimates stem from four or five independent data sets (men and women in the GSS, men in the ESS, and men and women in the GSOEP, although the latter are actual couples and hence may not be independent). Treating husbands and wives as in the GSOEP as independent, the meta estimate across the IJIVE estimates for the baseline cross-sectional specifications from the five samples is 0.165 with a standard error of 0.047. This compares to an OLS meta estimate of 0.186 with a standard error of 0.008. Moreover, the 95% confidence interval for the IJIVE meta estimate is 0.07 to 0.26, ruling out both very low and very high estimates. Overall, the provisional evidence presented here points in the direction that the income-life satisfaction relationship is mostly causal rather than driven by reverse causality or omitted factors.

References

- [1] Acemoglu, Daron and Jörn-Steffen Pischke (2001), "Changes in the Wage Structure, Family Income, and Children's Education," *European Economic Review*, 45, (May), 890-904.
- [2] Akerberg, Daniel A., and Paul J. Devereux (2009), "Improved JIVE Estimators for Overidentified Linear Models with and without Heteroskedasticity," *Review of Economics and Statistics*, 91, No. 2 (May), 351-362.
- [3] Angrist, Joshua D. and Jörn-Steffen Pischke (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton, NJ: Princeton University Press.
- [4] Apouey, Bénédicte and Andrew Clark (2009), "Winning Big but Feeling No Better? The Effect of Lottery Prizes on Physical and Mental Health," mimeographed, PSE (December).
- [5] Chevalier, Arnaud and Raymonn Lydon (2002), "Estimates of the Effect of Wages on Job Satisfaction," CEP Discussion Paper 531 (May).
- [6] Clark, Andrew (2003), "Looking for Labour Market Rents Using Subjective Data," mimeographed, PSE (September).
- [7] Clark, Andrew and Claudia Senik (2010), "Who compares to whom? The anatomy of income comparisons in Europe" *Economic Journal*, 120, N0. 544 (May), 573-594.
- [8] Deaton, Angus (2008), "Income, Health and Wellbeing Around the World: Evidence from the Gallup World Poll," *Journal of Economic Perspectives*, 22, No. 2 (Spring), 53-72.

- [9] Di Tella, Rafael and Robert MacCulloch (2006), “Some Uses of Happiness Data in Economics,” *Journal of Economic Perspectives*, 20, No. 1 (Winter), 25-46.
- [10] Dickens, William T. and Lawrence F. Katz (1987), “Inter-Industry Wage Differences and Industry Characteristics,” in: Kevin Lang and Jonathan Leonard, editors, *Unemployment and the Structure of Labor Markets*, New York: Basil Blackwell, 48-89.
- [11] Easterlin, Richard A. (1974), “Does Economic Growth Improve the Human Lot? Some Empirical Evidence,” in: David, Paul A. and Reder, Melvin W., Editors, *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz*, Academic Press, Inc., New York.
- [12] Easterlin, Richard A. (1995), “Will raising the incomes of all increase the happiness of all?” *Journal of Economic Behavior and Organization*, 27, No. 1, (June), 35-47.
- [13] Easterlin, Richard A., and Laura Angelescu (2009), “Happiness and Growth the World Over: Time Series Evidence on the Happiness-Income Paradox,” IZA Discussion Paper No. 4060.
- [14] Engelhardt, Gary V. and Jonathan Gruber (2005), “Social Security and Elderly Well-Being in the United States,” mimeographed, MIT (May).
- [15] Frey, Bruno S. and Alois Stutzer (2002), “What Can Economists Learn from Happiness Research?” *Journal of Economic Literature*, 40, No. 2 (June), 402-435.
- [16] Frijters, Paul, John P. Haisken-DeNew, and Michael A. Shields (2004), “Money Does Matter! Evidence from Increasing Real Income and Life

- Satisfaction in East Germany Following Reunification,” *American Economic Review*, 94, No. 3 (June), 730–740.
- [17] Gardner, Jonathan and Andrew J. Oswald (2007), “Money and mental wellbeing: A longitudinal study of medium-sized lottery wins,” *Journal of Health Economics*, 26, No. 1 (January), 49-60.
- [18] Gibbons, Robert and Lawrence F. Katz (1992), “Does Unmeasured Ability Explain Inter-Industry Wage Differentials?” *Review of Economic Studies*, 59, No. 3 (July), 515-535.
- [19] Holzer, Harry J., Lawrence F. Katz, and Alan B. Krueger (1991), “Job Queues and Wages,” *Quarterly Journal of Economics*, 106, No. 3 (August), 739-768.
- [20] Hout, Michael (2004), “Getting the Most Out of the GSS Income Measures,” GSS Methodological Report 101 (July).
- [21] Katz, Lawrence F., and Lawrence H. Summers (1989), “Industry Rents: Evidence and Implications,” *Brookings Papers on Economic Activity. Microeconomics*, 209-290.
- [22] Krueger, Alan B. and David Schkade (2008), “Sorting in the Labor Market: Do Gregarious Workers Flock to Interactive Jobs?” *Journal of Human Resources*, 43, No. 4, 859-883.
- [23] Layard, Richard (2003), “Happiness: Has Social Science A Clue?” Lionel Robbins Memorial Lectures delivered at LSE on March 3-5, 2003.
- [24] Layard, Richard, Guy Mayraz, and Stephen Nickell (2008), “The marginal utility of income,” *Journal of Public Economics*, 92, No. 8-9, (August), 1846-1857.

- [25] Layard, Richard, Guy Mayraz, and Stephen Nickell (2009), “Does Relative Income Matter? Are the Critics Right?” CEP Discussion Paper 938, LSE (July).
- [26] Luttmer, Erzo F. P. (2005), “Neighbors as Negatives: Relative Earnings and Well-Being,” *Quarterly Journal of Economics*, 120, No. 3 (August), 963-1002.
- [27] Manning, Alan (2003), *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton, NJ: Princeton University Press.
- [28] Mayraz, Guy, Gert G. Wagner, and Jurgen Schupp (2009), “Life Satisfaction and Relative Income: Perceptions and Evidence,” CEP Discussion Paper 918, LSE (March).
- [29] Murphy, Kevin M. and Robert H. Topel (1987), “Unemployment, Risk and Earnings: Theory and Evidence from a Model of Equalizing Wage Differentials,” in: Kevin Lang and Jonathan Leonard, editors, *Unemployment and the Structure of Labor Markets*, New York: Basil Blackwell, 103-140.
- [30] Murphy, Kevin M. and Robert H. Topel (1989), “Efficiency Wages Reconsidered: Theory and Evidence,” in: Yoram Weiss and Gideon Fishelson, editors, *Advances in the Theory and Measurement of Unemployment*, London: The Macmillan Press, 204-240.
- [31] Oreopoulos, Philip (2007), “Do dropouts drop out too soon? Wealth, health and happiness from compulsory schooling,” *Journal of Public Economics*, 91, No. 11-12 (December), 2213-2229.

- [32] Oswald, Andrew J., and Nattavudh Powdthavee (2008), “Death, Happiness, and the Calculation of Compensatory Damages,” *Journal of Legal Studies*, 37 (June), S217-S251.
- [33] Oswald, Andrew J., Eugenio Proto, and Daniel Sgroi (2009), “Happiness and Productivity,” IZA Discussion Paper No. 4645 (December).
- [34] Schwarz, Norbert (1987) *Stimmung als Information: Untersuchung zum Einfluss von Stimmungen auf die Bewertung des eigenen Lebens*, Lehr- und Forschungstexte Psychologie, Vol. 24, Berlin: Springer Verlag.
- [35] Schwarz, Norbert and Gerald L. Clore (1983), “Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states,” *Journal of Personality and Social Psychology*, 45, No. 3, (Sept), 513-523.
- [36] Shea, John (2000) “Does parents’ money matter?” *Journal of Public Economics*, 77, No. 2, (August), 155-184.
- [37] Stevenson, Betsey and Justin Wolfers (2008a), “Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox,” *Brookings Papers on Economic Activity*, No. 1, 1-87.
- [38] Stevenson, Betsey and Justin Wolfers (2008b), “Happiness Inequality in the United States,” *Journal of Legal Studies*, 37 (June), S33-S79.
- [39] Stock, James H., Jonathan H Wright, and Motohiro Yogo (2002), “A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments,” *Journal of Business and Economic Statistics*, 20, No. 4 (October), 518-529.

- [40] Stutzer, Alois, and Bruno S. Frey (2006), “Does Marriage Make People Happy, or Do Happy People Get Married?” *Journal of Socio-Economics*, 35, No. 2 (April), 326-247.
- [41] van Praag, Bernard M.S. and Ada Ferrer-i-Carbonell (2008) *Happiness Quantified: A Satisfaction Calculus Approach*, revised edition, Oxford: Oxford University Press

Appendix

Coding of industries

Sector	GSS		ESS, GSOEP
	1972 – 1988 1970 Census codes	1989 – 2006 1980 Census codes	NACE codes
1 Agriculture, forestry, fishery	17-28	10-31	1-5
2 Mining	47-57	40-50	10-14
3 Construction	67-77	60	45
4 Lumber, wood, furniture	107-118	230-242	20, 36
5 Stone, clay, glass	119-138	250-262	26
6 Metal	139-169	270-301	27-28
7 Machinery, exc. electrical	177-198, 258	310-332	29-30
8 Electrical machinery	199-209	340-350	31-32
9 Transportation equipment	219-238	351-370	34-35
10 Professional equipment	239-259	371-382	33
11 Food and tobacco	268-299	100-130	15-16
12 Textile, apparel, leather	307-327, 388-397	132-152, 220-222	17-19
13 Paper	328-337	160-162	21
14 Printing	338-339	171-172	22
15 Chemicals	347-369	180-192	24
16 Petroleum and rubber	377-387	200-212	23, 25
17 Other manufacturing	259, 398	390-392	37
18 Transportation	407-429	400-432	60-63
19 Communication	447-449	440-442	64
20 Utilities	467-479	460-472	40-41, 90
21 Wholesale trade	507-588	500-571	51
22 Retail trade	607-698	580-691	50, 52
23 Finance, insurance, real estate	707-718	700-712	65-71
24 Business services	727-748	721-742	72-74
25 Repair services	749-759	750-760	
26 Personal services	769-798	761-791	95-97
27 Recreation services	807-809	800-802	92
28 Health	828-848	812-840	85
29 Legal services	849	841	
30 Education	857-868	842-860	80
31 Religious and welfare services	877-879	861-871, 880	
32 Other services	869, 887-897	872, 881-892	91, 93
33 Public administration	907-937	907-937	75, 99

Coding of occupations

	Occupation	GSS		ESS, GSOEP
		1972 – 1988 1970 Census codes	1989 – 2006 1980 Census codes	ISCO-88 codes
1	Administrative and managerial	1, 56, 201-246	3-37	1000-1319, 2400, 2410-2419, 2470, 3440-3449
2	Engineers	2, 6-26	43-63	2000-2100, 2140-2149
3	Math and computer scientists	2-5, 34-36, 55	64-68	2120-2139
4	Natural scientists	42-54	69-83	2110-2114, 2210-2213
5	Health professionals	61-73	84-89	2200, 2220-2222
6	Health treatment occupations	74-76	95-106	2223-2230, 3220-3223, 3230-3232
7	Post-secondary teachers	102-140	113-154	2310
8	Teachers, exc. post-secondary	141-145	155-159	2300, 2320-2359
9	Counsellors, librarians, archivists	32-33, 174	163-165	2430-2432
10	Social scientists, urban planners	91-96	166-173	2440-2445
11	Social and religious workers	86-90, 100, 101	174-177	2446, 2460, 3460, 3480
12	Lawyers and judges	30-31	178-179	2420-2429
13	Writers, artists, athletes	175-194	183-199	2450-2455, 3470-3478
14	Technicians and support occupations	80-85, 150-173	203-235	3000-3213, 3224-3229
15	Sales occupations	260-296	243-285	3400-3429, 5000, 5200-5220
16	Clerical and admin. support occupations	301-396	303-389	3300-3340, 3430-3434, 4000-4223
17	Private household workers	980-986	403-407	5121, 5131, 5133
18	Protective services workers	960-976	413-427	3450, 5160-5169
19	Service workers, exc. 17 and 18	901-954	433-469	5100-5120, 5122-5130, 5132, 5139-5149
20	Farming occupations	801-846	473-499	6000-6154
21	Crafts and repair workers	401-586	503-699	7000-7442
22	Operators and laborers	601-796	703-889	8000-9330

Figure 1
Visual IV
Happiness Against ln of Family Income for Men, GSS, 1972 – 2006




Figure 2
Visual IV
Job Satisfaction Against ln of Family Income for Men, GSS, 1972 – 2006

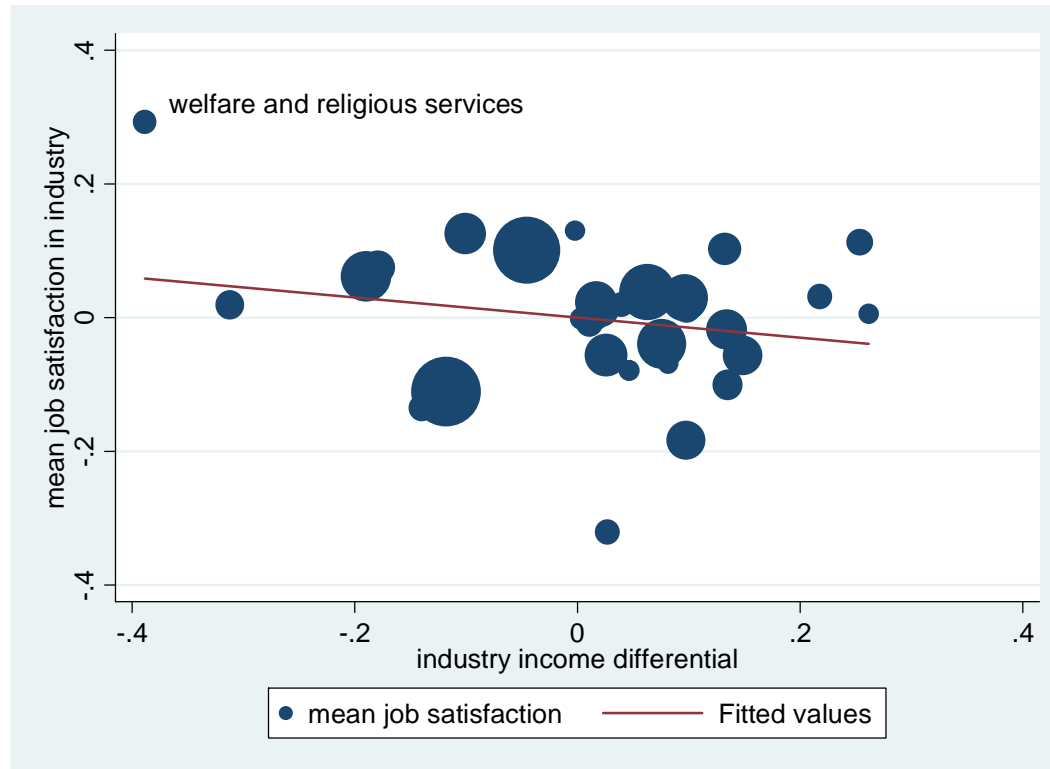


Table 1
Regressions of Happiness on ln of Family Income for Men
General Social Survey, 1972 – 2006
(Standard errors in parentheses)

Estimation method	Dependent Variable						
	Happiness				Job Satisfaction		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OLS	0.163 (0.014)	0.119 (0.014)	0.114 (0.014)	0.089 (0.014)	0.121 (0.014)	0.115 (0.014)	0.099 (0.014)
2SLS	0.192 (0.069)	0.112 (0.072)	0.208 (0.088)	0.210 (0.085)	-0.150 (0.069)	-0.189 (0.073)	-0.046 (0.089)
IJIVE	0.196 (0.075)	0.111 (0.079)	0.222 (0.100)	0.239 (0.096)	-0.174 (0.075)	-0.219 (0.081)	-0.066 (0.101)
First stage <i>F</i> -statistic	13.29	12.20	8.02	8.05	13.29	12.20	8.02
Baseline controls	✓	✓	✓	✓	✓	✓	✓
4 marital status dummies		✓	✓	✓		✓	✓
21 occupation dummies			✓	✓			✓
4 job satisfaction dummies				✓			

Weighted regressions using GSS sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Instruments are 32 industry dummies. Number of observations is 12,121. Heteroskedasticity robust standard errors in parentheses.

Table 2
Comparison with Luttmer (2005)
Men, General Social Survey, 1972 – 2006
(Standard errors in parentheses)

Estimation method	(1)	(2)	(3)
OLS	0.163 (0.014)	0.163 (0.014)	0.163 (0.014)
2SLS	0.192 (0.069)	0.332 (0.065)	0.261 (0.044)
IJIVE	0.196 (0.075)	0.347 (0.069)	0.306 (0.063)
First stage <i>F</i> -statistic	13.29	24.06	3.51
Instruments	Industry	Occupation	Ind*Occ

Weighted regressions using GSS sampling weight. The coefficient on ln(family income) is displayed. All regressions include controls for age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Instruments are 32 industry, 21 occupation dummies or their interactions. Number of observations is 12,121. Heteroskedasticity robust standard errors in parentheses.

Table 3
Regressions of Happiness on ln of Family Income for Married Men and Women
General Social Survey, 1972 – 2006
(Standard errors in parentheses)

Estimation method	Sample									
	Husbands		Wives		Wives, not working		Wives, working			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	0.159 (0.021)	0.153 (0.022)	0.095 (0.017)	0.089 (0.018)	0.105 (0.025)	0.095 (0.026)	0.090 (0.025)	0.086 (0.026)	0.092 (0.026)	0.086 (0.026)
2SLS	0.121 (0.101)	0.256 (0.122)	0.056 (0.077)	0.097 (0.094)	0.035 (0.097)	0.028 (0.123)	0.098 (0.107)	0.168 (0.119)	0.147 (0.122)	0.194 (0.132)
IJIVE	0.118 (0.114)	0.278 (0.143)	0.052 (0.083)	0.098 (0.104)	0.019 (0.114)	0.008 (0.153)	0.100 (0.124)	0.186 (0.141)	0.160 (0.147)	0.222 (0.164)
First stage <i>F</i> -statistic	9.46	6.49	12.18	9.13	6.41	4.80	7.61	6.30	5.98	5.16
Baseline controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
21 occupation dummies		✓		✓		✓		✓		✓
32 industry dummies (wives)									✓	✓
Number of observations	7,737	7,737	8,450	8,450	3,505	3,505	4,945	4,945	4,945	4,945

Weighted regressions using GSS sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Instruments are 32 industry dummies for husband's industry affiliation. Heteroskedasticity robust standard errors in parentheses.

Table 4
Regressions of Satisfaction on ln of Family Income for Men
European Social Survey, 2002 – 2008
(Standard errors in parentheses)

Estimation method	Dependent Variable					
	Happiness			Life Satisfaction		
	(1)	(2)	(3)	(4)	(5)	(6)
OLS	0.215 (0.016)	0.180 (0.015)	0.169 (0.016)	0.243 (0.015)	0.219 (0.015)	0.204 (0.016)
2SLS	0.135 (0.091)	0.098 (0.091)	0.081 (0.115)	0.155 (0.090)	0.131 (0.090)	0.179 (0.113)
IJIVE	0.124 (0.102)	0.085 (0.102)	0.060 (0.136)	0.145 (0.101)	0.118 (0.102)	0.172 (0.133)
First stage <i>F</i> -statistic	9.40	9.25	6.80	9.40	9.25	6.80
Baseline controls	✓	✓	✓	✓	✓	✓
4 marital status dummies		✓	✓		✓	✓
21 occupation dummies			✓			✓

Weighted regressions using the product of the population and design weights as sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, six education dummies, and a full set of interactions of wave and country dummies. Instruments are 29 industry dummies. Number of observations is 27,740. Heteroskedasticity robust standard errors in parentheses.

Table 5
Regressions of Happiness on ln of Family Income for West-German Men
German Socio-economic Panel (GSOEP), 1984 – 2007
(Standard errors in parentheses)

Estimation method	Dependent Variable						
	Happiness				Job Satisfaction		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OLS	0.333 (0.026)	0.292 (0.025)	0.292 (0.024)	0.176 (0.019)	0.264 (0.024)	0.254 (0.024)	0.244 (0.024)
2SLS	0.644 (0.207)	0.623 (0.218)	0.696 (0.203)	0.596 (0.161)	0.157 (0.186)	0.166 (0.196)	0.199 (0.199)
IJIVE	0.658 (0.218)	0.640 (0.232)	0.721 (0.218)	0.626 (0.173)	0.144 (0.197)	0.153 (0.208)	0.189 (0.213)
First stage <i>F</i> -statistic	3.65	3.83	3.87	3.99	3.65	3.83	3.87
Baseline controls	✓	✓	✓	✓	✓	✓	✓
4 marital status dummies		✓	✓	✓		✓	✓
22 occupation dummies			✓	✓			✓
10 job satisfaction dummies				✓			

Weighted regressions using GSOEP sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for nationality, nine education dummies, nine state dummies, and 23 year dummies. Instruments are 28 industry dummies. Number of observations is 57,271. Standard errors, clustered by individual, in parentheses.

Table 6
Fixed Effects Regressions of Happiness on ln of Family Income for West-German Men
German Socio-economic Panel (GSOEP), 1984 – 2007
(Standard errors in parentheses)

Estimation method	Dependent Variable						
	Happiness				Job Satisfaction		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OLS	0.183 (0.028)	0.166 (0.026)	0.163 (0.026)	0.125 (0.023)	0.123 (0.025)	0.128 (0.025)	0.126 (0.025)
2SLS	0.584 (0.340)	0.583 (0.353)	0.431 (0.379)	0.506 (0.349)	-0.074 (0.413)	-0.103 (0.429)	-0.279 (0.452)
IJIVE	0.670 (0.452)	0.690 (0.480)	0.507 (0.556)	0.650 (0.515)	-0.181 (0.564)	-0.226 (0.594)	-0.505 (0.680)
First stage <i>F</i> -statistic	1.69	1.66	1.41	1.44	1.69	1.66	1.41
Baseline controls	✓	✓	✓	✓	✓	✓	✓
4 marital status dummies		✓	✓	✓		✓	✓
22 occupation dummies			✓	✓			✓
10 job satisfaction dummies				✓			

Weighted fixed effects regressions using GSOEP sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age squared and 23 year dummies. Instruments are 28 industry dummies. Number of observations is 57,271, number of individuals is 9,212. Standard errors, clustered by individual, in parentheses.

Table 7
Regressions of Happiness on ln of Family Income for Married Men and Women (West-Germany)
German Socio-economic Panel (GSOEP), 1984 – 2007
(Standard errors in parentheses)

Estimation method	Sample									
	Husbands		Wives		Wives, not working		Wives, working			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	0.309 (0.038)	0.298 (0.037)	0.294 (0.033)	0.261 (0.034)	0.298 (0.056)	0.228 (0.057)	0.305 (0.037)	0.299 (0.038)	0.292 (0.037)	0.287 (0.038)
2SLS	0.623 (0.296)	0.683 (0.268)	0.738 (0.189)	0.744 (0.237)	0.706 (0.263)	0.450 (0.317)	0.608 (0.202)	0.697 (0.243)	0.579 (0.208)	0.673 (0.257)
IJIVE	0.644 (0.320)	0.713 (0.292)	0.758 (0.196)	0.781 (0.253)	0.751 (0.289)	0.500 (0.368)	0.630 (0.218)	0.729 (0.261)	0.601 (0.224)	0.708 (0.279)
First stage <i>F</i> -statistic	3.44	3.32	6.84	4.44	4.97	2.87	5.55	4.92	5.06	4.08
Baseline controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
22 occupation dummies		✓		✓		✓		✓		✓
28 industry dummies (wives)									✓	✓
Number of observations	37,376	37,376	37,376	37,376	14,044	14,044	22,198	22,198	22,198	22,198

Weighted regressions using GSOEP sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for nationality, nine education dummies, nine state dummies, and 23 year dummies. Instruments are 28 industry dummies. Standard errors, clustered by individual, in parentheses.

Table 8
Alternative Instruments Using Employer Differentials
Men, General Social Survey, 1972 – 2006
(Standard errors in parentheses)

Estimation method	Instruments							
	Union status		Firm size		Union and firm size		Union, firm size and industry	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OLS	0.169 (0.017)	0.122 (0.017)	0.169 (0.021)	0.108 (0.022)	0.163 (0.026)	0.101 (0.026)	0.163 (0.026)	0.101 (0.026)
2SLS	0.083 (0.158)	0.130 (0.121)	0.144 (0.142)	0.131 (0.151)	0.262 (0.153)	0.251 (0.137)	0.333 (0.097)	0.297 (0.106)
IJIVE	0.081 (0.160)	0.129 (0.122)	0.146 (0.153)	0.136 (0.164)	0.274 (0.166)	0.259 (0.145)	0.384 (0.121)	0.354 (0.137)
First stage <i>F</i> -statistic	88.06	148.55	14.28	8.34	11.31	14.51	6.37	4.76
Baseline controls	✓	✓	✓	✓	✓	✓	✓	✓
4 marital status and 21 occupation dummies		✓		✓		✓		✓
Number of observations	8,418	8,418	4,987	4,987	3,467	3,467	3,467	3,467

Weighted regressions using GSS sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Instruments are a dummy for union status, six dummies for firm size categories, and 32 industry dummies. Heteroskedasticity robust standard errors in parentheses.

Table 9
Alternative Instruments Using Changes in Wage Inequality
Men, General Social Survey, 1972 – 2006
(Standard errors in parentheses)

Estimation method	(1)	(2)
OLS	0.163 (0.014)	0.119 (0.014)
OLS controlling for rank	0.071 (0.035)	0.034 (0.034)
2SLS	0.258 (0.219)	0.167 (0.233)
IJIVE	0.329 (0.771)	0.128 (1.081)
First stage <i>F</i> -statistic	1.10	0.93
Baseline controls	✓	✓
21 occupation dummies		✓

Weighted regressions using GSS sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for black and other race, eight education dummies, and 25 year dummies. Instruments is a linear indicator for years of schooling interacted with 25 year dummies. Number of observations is 12,121. Heteroskedasticity robust standard errors in parentheses.

Table 10
Role of Permanent Income: Regressions of Happiness on ln of Family Income for West-German Men
German Socio-economic Panel (GSOEP), 1984 – 2007
(Standard errors in parentheses)

Estimation method	Pooled		Fixed Effects	
	(1)	(2)	(5)	(6)
OLS	0.343 (0.034)	0.355 (0.031)	0.204 (0.032)	0.202 (0.032)
2SLS, IV: 3 year avg. inc.	0.374 (0.039)	0.390 (0.036)	0.213 (0.041)	0.211 (0.041)
2SLS, IV: ind. dummies	0.423 (0.243)	0.614 (0.225)	0.228 (0.380)	0.237 (0.363)
First stage <i>F</i> -statistic (ind.)	2.89	3.10	1.67	1.15
Baseline controls	✓	✓	✓	✓
22 occupation dummies		✓		✓

Weighted regressions using GSOEP sampling weight. The coefficient on ln(family income) is displayed. Baseline controls are age, age squared, dummies for nationality, nine education dummies, nine state dummies, and 23 year dummies. Number of observations is 32,732. Standard errors, clustered by individual, in parentheses.