

WORK AND WELL-BEING

ANDREW E. CLARK*

Introduction

A perhaps apocryphal story has a departmental seminar speaker starting their presentation with the line “Today I’d like to talk to you about my research on XXX...”. A hand immediately goes up for the first question from the audience: “Why?”

I suspect that all researchers have come across this kind of reaction at some point in their career. Its prevalence surely depends on the discipline in which one works. Those carrying out medical research or working on renewable energy may never have come across it; economists who suggest that asking people how happy they were is a good research strategy are likely more familiar with it.¹ In this context, I would like to consider two broad research questions in the realm of labour economics in which subjective well-being information has played a key role. The word “key” is undoubtedly over-used, and not only in economics and by me, but I think that it is warranted here, in the sense that I have difficulty in seeing how the research in question could have been carried out without appealing to this kind of subjective data.

The first of the two research questions in the world of work has to do with worker allocation on the labour market: we see some people working and others not, and amongst those who are in employment there is considerable heterogeneity in the type of jobs that they do. Why do these patterns of labour-market status pertain? One useful distinction here is between optimal choice and constraints. Is unemployment chosen by the unemployed? Do the self-employed think they have the best of all labour-market statuses, or could they not find anything better? Understanding whether choice is constrained in this way is important

for policy purposes: if individuals are making the best choice out of all labour market statuses, then all is well. If on the other hand individuals are making constrained choices then there are potential welfare gains: individuals may well be better off in other labour force statuses, and we should consider how we can loosen the constraints in question.

The second research area is quite simply that of the distribution of well-being. This distribution may be at a point in time – between individuals, occupations or even between countries – or it may be over time within a country.

Allocation on the labour market

There are a number of different statuses on the labour market: why do people end up where they do? In what follows, I am going to take it for granted that subjective well-being scores do contain useful interpersonally-comparable information.² By treating subjective well-being scores as an index of whatever it is that people are trying to maximise, we can arguably make progress in determining whether different labour market statuses result from unconstrained choices or not.

Probably the first question to be asked here concerns unemployment. This has been one of the principal areas of research in the subjective well-being literature, along with the relationship with income. We would like to know here to what extent unemployment is voluntary and to what extent involuntary, and we can use well-being scores to help answer it. If unemployment is chosen (and well-being scores are comparable between individuals) then in cross-section data, *ceteris paribus*, the employed and the unemployed should report similar well-being scores. However, if the unemployed are found to report lower well-being scores than do the employed then either unemployment is not a choice, or it is a choice and it is systematically chosen by individuals who have lower levels of well-being.

* Paris School of Economics and IZA.

¹ This is less true now than it was in the past: my most raucous seminar audiences were in the first half of the 1990s.

² There is by now a considerable amount of disparate research which suggests that this is the case: some of this is reviewed in Clark, Frijters and Shield (2008).



Figure



The results here are arguably unequivocal. The analysis of cross-section data reveals that the unemployed report significantly lower well-being scores than do the employed. The estimated regression coefficient here is both very significant and very large in absolute size. This is illustrated in the Figure, where the raw data from eight waves of the European Community Household Panel (ECHP) (1994 to 2001) have the unemployed reporting life satisfaction scores almost two points lower than those of the employed, on a scale that only goes from one to six. In this sense, unemployment is one of the most negative events to be identified in the typical panel data to which happiness economics appeals (such as the British Household Panel Survey (BHPS) and the German Socio-Economic Panel (GSOEP)).

This negative correlation between unemployment and well-being is robust to the addition of explanatory variables (including income) in multivariate regression analyses (amongst many others, see Clark and Oswald 1994 and Winkelmann and Winkelmann 1998). If we believe that well-being scores are comparable between individuals, then the conclusion is that unemployment is associated with far lower well-being, and this lower well-being is largely not explained by the associated lower income.

The issue of well-being score comparability, such that it is those with low well-being scores who end up unemployed, can be addressed by appealing to the increased development of panel data sets including well-being questions over the past 25 years. By re-interviewing individuals at regular (very often one-year) intervals, fixed psychological factors or response style systematically affecting well-being scores

can be controlled for. Even if it is true that grumpy people are more likely to end up unemployed, the economic question regarding choice or constraint can be answered by seeing whether grumpy people end up even grumpier when they enter unemployment.³

This kind of analysis is illustrated simply in the Table, using data from the first 17 waves of the BHPS. I here consider those aged 16 to 65, who are either working (employed or self-employed) or unemployed. I use two well-being measures. The first is the zero to

twelve General Health Questionnaire Caseness score (as used in Clark and Oswald 1994), reflecting overall psychological functioning, which is available in all waves. The second is a “standard” overall life satisfaction score, on a one to seven scale, which was recorded in waves 6–10 and 12–17. In both cases larger numbers reflect higher levels of well-being.

The message from the first panel of the Table is that, as in the Figure, the unemployed report lower well-being scores than do those in work. The difference in well-being by labour-force status is significant at all conventional levels for both measures. The second and third panels report “within-subject” analyses of the same relationship. These are presented as transition matrices, and measure the average change in well-being according to both last year’s and this year’s labour-force status. The diagonal terms show that there is little change in well-being from year to year for those who do not change labour-force status. However, the off-diagonal elements suggest substantial movements. Someone who was employed last year but is now unemployed has an average drop in well-being of 1.07 points; when individuals move from unemployment back to work their well-being rises by about 1.3 points. Both of these figures are similar to the simple cross-section difference in the first panel of the Table. The same conclusion, to a somewhat lesser extent, holds for life satisfaction. Were it to be the case that the unhappy are more likely to become unemployed, the within-subject work-unemployment well-being difference would be smaller than that between subjects. The Table provides only little evidence of this.

³ A useful analogous analysis of marriage using panel data concluded that most of the reason why the married were more satisfied than singles in cross-section data was that more satisfied people were more likely to get married (Stutzer and Frey 2006).

Table

Well-being, employment and unemployment: BHPS

	Employed(t)	Unemployed(t)
GHQ-12	10.36	9.19
Life satisfaction	5.24	4.61
Transition matrix: GHQ-12		
Employed(t-1)	-0.039	-1.066
Unemployed(t-1)	1.305	-0.011
Transition matrix: Life satisfaction		
Employed(t-1)	-0.020	-0.370
Unemployed(t-1)	0.358	0.025

Source: BHPS, 1991–2007. Life satisfaction measured on a one to seven scale; GHQ-12 on a zero to twelve scale.

The conclusion from the well-being analysis of unemployment is then that it is associated with significantly lower levels of well-being than working, both between individuals and within-individual. This is not consistent with unemployment being a choice. Second, most of the well-being impact of unemployment seems to be non-pecuniary. Both of these findings should influence labour-market policy.

Unemployment is not the only labour-force status of interest. There is also a small literature which has looked at well-being and self-employment.⁴ Although the self-employed do worse than the employed with respect to many aspects of the job, they systematically (at least in rich countries) report higher levels of overall job satisfaction than do the employed. If, as well-being data suggest, self-employment is better than employment, why aren't more of us self-employed? Two broad explanations have been proposed. One is in terms of matching on the labour market. Some individuals may really like autonomy but not be particularly risk-averse: with these preferences, they will prefer self-employment to employment. Others may have different preferences that lead them to prefer employment. Certain parameterisations of the utility will yield greater satisfaction for those who choose self-employment than for those who choose employment, even though both are happy with their choice.

The second explanation is in terms of barriers to entry into self-employment: some people really want to be self-employed but can't, so that their choices are constrained. These barriers are often imagined in terms of capital constraints. A well-known contribu-

tion by Blanchflower and Oswald (1998) appeals to satisfaction data to provide some evidence in favour of this capital constraints story. The authors use the National Child Development Study (NCDS) data to show that those who receive an inheritance are more likely to become self-employed (and control for endogeneity by instrumenting inheritances with the death of a parent). They also show that the self-employed in their sample report higher job satisfaction on average than do the employed. However, they then separate the sample to show that this gap is only found for the self-employed who did not receive an inheritance. The self-employed who have received inheritances actually report the same level of job satisfaction as do employees. This finding is consistent with capital constraints, in the sense that self-employment is a choice for those who have enough resources, but is simply not available for some others. Without the analysis of subjective data, not only would the "problem" of the satisfied self-employed not have been apparent, but the interpretation via capital constraints would likely have been more difficult to show.

The same comparison of labour-force statuses can also be carried out with respect to retirement. One crucial piece of information in terms of the debate over different retirement ages is the effect of retirement on both health and overall well-being. While there is an enormous literature on well-being and age – and on older individuals and retirees specifically – that on changes in well-being attendant on retirement is not particularly large. The analysis here is complicated by serious problems of endogeneity (those in worse health and with lower levels of well-being tend to choose to retire earlier). These are dealt with by looking for exogenous movements in retirement age or benefits. Some contributions along these lines are Charles (2004) who finds that, when retirement is instrumented, it reduces both depression and loneliness, Coe and Lindeboom (2008) who find no overall effect of retirement on health, and Dave et al. (2008), who suggest that retirement is associated with worse mental health. In general, the results in this small literature are ambiguous, and more work on this particular aspect of well-being and labour supply, especially in the context of the greying population, would seem of great use.

Subjective well-being and job quality

Following on from the broad idea of subjective well-being as reflecting the relative attractiveness of dif-

⁴ Two examples are Blanchflower and Oswald (1998) and Bianchi (2010).

ferent labour-market outcomes, it seems reasonable that individuals should leave lower-satisfaction for higher-satisfaction jobs. This is in fact what they are observed to do in panel data: job satisfaction at time t has been shown to be a good predictor of future quits (Clark 2001 and Lévy-Garboua et al. 2007, amongst others). This role of job satisfaction in predicting labour-market mobility continues to hold regardless of the observable job characteristics which are added to the quit equation. This underlines that both jobs consist of more than just wages and hours and other observables, and that job satisfaction does a good job of picking up these unobservable job characteristics (which are important to workers, as they are observed to seek them out).

This brings me to the second topic, that of the distribution of well-being at work. As the quits literature has intimated, not all jobs are the same. Furthermore, they are not the same in ways that are not only reflected in wages and hours of work. This allows us to make some progress in talking about job quality. This is a topic that has interested me since the mid-1990s. The central questions here are whether some jobs seem to be systematically better than others, and whether jobs have been getting worse over time.

There has in the past been a tendency to examine job quality as purely a function of income, and perhaps hours of work. There is no guarantee that these on their own will provide an adequate description of job quality, however. The job values questions in the three “Work Orientations” waves of the International Social Survey Programme (ISSP) (1989, 1997 and 2005), suggest that the job aspects that both men and women consider to be the most important are job security and the interest of the job (Clark 2010). These are followed by independence and whether the job allows the individual to help others. Only around a quarter of ISSP respondents consider having a job with a high income as very important. These most important aspects of the job would seem to be very difficult to measure objectively, and any evaluation of job quality based only on income and hours of work will miss out on many job domains that workers value.

One approach, followed in Clark (2010), is then to construct a series of indicators using the ISSP data to evaluate worker’s subjective evaluations of pay, hours of work, promotion and job security, the difficulty of the job, job content, and interpersonal relationships at work. These are complemented with a

job satisfaction score which arguably summarises all of the aspects that workers appreciate or dislike about their job, some of which are explicitly measured in the ISSP and some of which are likely not.

The analysis of movements in these subjective job quality measures over time reveals that while some of them dropped between 1989 and 1997, there has been a recovery up to 2005. This is particularly the case for promotion opportunities and job security. Concentrating on the overall measure, job satisfaction rose for both men and women. The analysis shows that this was partly, but far from entirely, due to the movements in income and hours of work over the period in question.

The regression analysis of both the six specific domains of job quality, as well as the overall measure of job satisfaction, underline that there are systematic differences between both types of individual and types of job in terms of well-being at work. In particular, men are less satisfied with their jobs than women, *ceteris paribus*, and older employees and the married are more satisfied. A number of the country dummy variables are also significant: workers in Hungary, Japan and France are relatively miserable, while the most satisfied on average are found in Denmark, Switzerland and the US. Last, it is very typically found that both industry and occupation are significantly correlated with subjective measures of job quality. As such, jobs are not homogeneous, and in ways other than just earnings and hours of work. When sufficient numbers of different waves of data are available, it can also be shown that the trends in job satisfaction are not the same across different groups of workers (Clark 2005 carries out this exercise using BHPS data). Subjective well-being data thus allows us to identify which groups of workers are doing worse than others, and those who are falling behind.

Let’s put it all together

Understanding the distribution of well-being across society, and here across workers, is surely a useful piece of information in its own right. Yet we can also consider it as a complement to the analysis of worker behaviour sketched out in the section on the labour market. Empirical analysis shows that the well-being of the unemployed is lower than that of the employed, on average. But of course no-one is average. The systematic heterogeneity in job quality discussed above will translate into heterogeneity in the employ-

ment-unemployment well-being gap. Put simply, some will lose more psychologically from unemployment than others. Sceptics may claim that changes over time in life satisfaction scores are going to be too noisy to reflect this phenomenon. Yet the drop in well-being on entering unemployment is correlated with both the search intensity of the unemployed, and the speed at which they consequently find a new job (see Clark 2003, using BHPS data, and Clark et al. 2010, in GSOEP data). This seems like an important piece of information. As is arguably the case with all individual choices, it is not the level of utility that counts, but the level of utility compared to the other options. In this context, shorter unemployment spells will not result from making unemployment a miserable experience, but from making sure that employment is attractive relative to unemployment: it is the delta that counts (and the latter also sounds rather more palatable as a policy proposition).

Taking heterogeneity seriously may also help us to understand why the results regarding retirement are so mixed: some may experience a drop in well-being on retirement, and others a rise. This is obviously linked to the type of job that the individual had; however, it is also linked to the type of retirement that the individual can expect too. As the results in Clark and Fawaz (2009) demonstrate, a simple taxonomy of “good jobs” doesn’t necessarily do a good job in predicting the change in well-being on retirement. Those who had bad jobs may equally have poor pensions: as such, they may not always experience much of a well-being benefit from stopping work. Again, behaviour is driven by a comparison of available options on the labour market, and subjective well-being information would seem like a useful way of summarising how attractive these different options are.

Despite the increasing amount of work using well-being data, in particular over the past ten years or so, there are still enormous gaps in our knowledge. Taking well-being data seriously as both an index of the distribution of well-being and a driver of individual behaviour will likely remain a very fruitful area of research for the foreseeable future, not only in economics but across all of the social sciences.

References

- Bianchi, M. (2010), “Financial Development, Entrepreneurship and Job Satisfaction”, *Review of Economics and Statistics*, in press.
- Blanchflower, D. G. and A. J. Oswald (1998), “What Makes an Entrepreneur?”, *Journal of Labor Economics* 16, 26–60.
- Charles, K. K. (2004), “Is Retirement Depressing? Labour Force Inactivity and Psychological Well-Being in Later Life”, in S.W. Patachek, ed., *Research in Labour Economics* 23, Elsevier, Amsterdam.
- Clark, A. E. (2001), “What Really Matters in a Job? Hedonic Measurement Using Quit Data”, *Labour Economics* 8, 223–42.
- Clark, A. E. (2003), “Unemployment as a Social Norm: Psychological Evidence from Panel Data”, *Journal of Labor Economics* 21, 323–51.
- Clark, A. E. (2005), “Your Money or Your Life: Changing Job Quality in OECD Countries”, *British Journal of Industrial Relations* 43, 377–400.
- Clark, A. E. (2010), “Work, Jobs and Well-Being Across the Millennium”, in E. Diener, J. Helliwell and D. Kahneman, eds., *International Differences in Well-Being*, Oxford University Press, Oxford.
- Clark, A. E. and Y. Fawaz (2009), “Valuing Jobs via Retirement: European Evidence”, *National Institute Economic Review* 209, 88–103.
- Clark, A. E., P. Frijters and M. Shields (2008), “Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles”, *Journal of Economic Literature* 4, 95–144.
- Clark, A. E., A. Knabe and S. Rätzel (2010), “Boon or Bane? Others’ Unemployment, Well-being and Job Insecurity”, *Labour Economics* 17, 52–1.
- Clark, A. E. and A. J. Oswald (1994), “Unhappiness and Unemployment”, *Economic Journal* 104, 48–59.
- Coe, N. and M. Lindeboom (2008), “Does Retirement Kill You? Evidence from Early Retirement Windows”, *IZA Discussion Paper* no. 3817.
- Dave, D., I. Rashad and J. Spasojević (2008), “The Effects of Retirement on Physical and Mental Health Outcomes”, *Southern Economic Journal* 75, 497–523.
- Lévy-Garboua, L., C. Montmarquette and V. Simonnet (2007), “Job Satisfaction and Quits: Theory and Evidence from the German Socioeconomic Panel”, *Labour Economics* 14, 251–8.
- Stutzer, A. and B. S. Frey (2006), “Does Marriage Make People Happy, Or Do Happy People Get Married?” *Journal of Socio-Economics* 35, 32–47.
- Winkelmann, L. and R. Winkelmann (1998), “Why Are the Unemployed So Unhappy? Evidence from Panel Data”, *Economica* 5, 1–15.