ISSN 1329 - 2676

the CENTRE for LABOUR MARKET RESEARCH

What can happiness data tell us about the labour supply curve?

by

Dr Alfred Michael Dockery
Centre for Labour Market Research and School of Economics and Finance
Curtin University

Acknowledgement: This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA project was initiated and is funded by the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs (FaHCSIA) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the author and should not be attributed to either FaHCSIA or the Melbourne Institute.

November 2011

CLMR DISCUSSION PAPER SERIES 2011/04

the Centre *for* Labour Market Research, Curtin Business School, Curtin University,
GPO Box U1987 Perth WA 6845 Australia
Tel: 61 8 9266 1744 Fax: 61 8 9266 3026
Email: patricia.madden@cbs.curtin.edu.au
http://www.business.curtin.edu.au/clmr

The Centre wishes to acknowledge the support of The Western Australian Department of Training and Workforce Development

ABSTRACT

This paper proposes an alternative empirical approach to the estimation of the labour supply function, one that is based upon subjective wellbeing data and potentially addresses several limitations of the standard approach. The labour supply function in neoclassical economics is based on the idea of an 'implicit utility function' which assumes individuals are free to choose their 'utility maximising' hours of work. Clearly, this is not the case for many workers. Further, it precludes by assumption the possibility of workers choosing to work more hours than is optimal for their wellbeing, despite widespread evidence of 'overwork' in modern economies. Instead, it is proposed that data on hours worked, wages and subjective wellbeing be used to explicitly estimate the utility function from which the labour supply schedule and elasticity of supply with respect to wages can be derived. The approach is demonstrated using data for unpartnered male employees from the Household, Income and Labour Dynamics in Australia Survey. The results are consistent with people working longer hours than is compatible with the maximisation of wellbeing. A number of extensions and applications of the approach are suggested.

1. INTRODUCTION

In much of the developed world people are now more affluent than ever, yet one of the growing social concerns relates to problems arising from 'overwork'. This has variously been described in terms of the incidence of people working long- or very long-hours; the phenomenon of 'burn-out'; the challenge of balancing work and family life; and time-stress or the 'time crunch', even a problem of 'presenteeism' - people attending work even they are not well. Applying the neo-classical paradigm, it seems something of a paradox that a well-to-do, rational *homo economicus* should choose to work so many hours each week that it detracts from their quality of life, and even impacts negatively upon their physical and mental health.

The neo-classical model of the choice of working hours - of labour supply - views an individual's preferences as being described by an underlying utility curve that is a function of two arguments: hours of leisure and the consumption of goods and services. Consumption is made possible through income earned from working. The shape of that utility curve is not measured directly, but inferred from the model's assumption that the individual chooses the number of hours of work to maximise their utility. Thus observed hours of work is taken to be the optimising point at which the slope of the utility curve equates to the marginal rate of substitution between leisure and consumption: that is, the real wage. One of the key parameters of interest to economists and policymakers in this analysis is the elasticity of labour supply with respect to changes in the real wage. As is well known, the model's prediction regarding the sign of that parameter is ambiguous due to the opposing income and substitution effects associated with a change in the real wage.

There have been many extensions to the basic model, such as the incorporation of household production, intra-family decision making, habit formation and life-cycle considerations. However, essentially all empirical approaches rely on an implicit utility curve and the assumption that observed hours of work represent the utility maximising solution (see, Blundell and McCurdy, 1999).

There are at least two reasons to expect that actual hours worked do not represent the worker's utility optimising solution. Most obviously, individuals are not free to choose any number of hours of work at the given wage rate. Taking into consideration the demand side of the labour market, workers are more likely to be faced with a limited number of discrete choices: perhaps to work part-time or full-time, or to work full-time or not at all. More importantly, evidence emerging from 'happiness' research and experimental economics gives weight to the proposition that individuals may not make decisions that maximise their utility (or wellbeing). In particular, they may systematically overestimate the utility gained from consumption and status, and systematically underestimate 'intrinsic' benefits from more leisure and experiential values gained from time with friends and family (Frey, 2008: 127-137).

This paper explores some of these issues by proposing a completely different way of estimating the labour supply function. Rather than rely on an implicit utility function and assume optimality of the observed hours choice, a 'happiness' model is used to directly estimate the utility function conditional on leisure and income, and the maximum of that function with respect to hours worked is then derived.

The next section provides a brief background to the issues. The dataset used in the empirical analysis, the Household, Income and Labour Dynamics in Australia (HILDA) Survey, is then introduced and some descriptive data on work hours and preferences presented. Section 4 develops the model and section 5 reports the results of the empirical estimation for the relatively homogenous group of single male employees. Section 6 offers a number of suggestions for future improvements and applications, and provides an applied example by demonstrating the importance of job satisfaction in shaping preferred working hours. In the concluding section 7 it is argued that the approach offers some promise, and tentatively notes that the initial results are consistent with employees working more hours than are compatible with optimal wellbeing.

2. BACKGROUND

In her 1992 book, *The Overworked American*, Juliet Schor notes that productivity per worker in the US had more than doubled since 1948 so that conceivably Americans could have maintained the same standard of material wealth but be working half as much. Yet none of that 'productivity dividend' appeared to have been taken in the form of reduced workloads. "In 1990, the average American owns and consumes more than twice as he or she did in 1948, but also has less free time" (Schor, 1992: 2). Neither, do the labour

saving innovations in the home and other conveniences appear to have reduced the number of hours that full-time housewives devote to housework (Schor, 1992: 8).

In Australia, real per capita income increased by 66 per cent from 1960 to 2010, based on the Australian Bureau of Statistics' chain volume measures (ABS, 2011). That is to say, every Australian now has two-thirds more goods and services at their disposal than in 1960. Between just 1979 and 2010, real GDP per hour worked increased by 60 per cent (ABS, 2011)[1]. Based on the eight years of the HILDA data used in this paper (described below), employees reported working an average of 36.1 hours per week, a figure that remained relatively unchanged between 2001 (36.3 hours) and 2008 (36.2 hours).

Becker (1965) suggested that the full cost of a good should include the opportunity cost of time taken to consume it, just as economists treat foregone earnings as part of the cost of education. Thus some goods, such as going to the movies, are 'time intensive'. An increase in wages will make time intensive goods relatively more expensive, and lead to a shift in consumption toward income intensive, rather than time intensive consumption. An increase in non-earned income would reduce the relative prices of time intensive goods, and indirectly increase leisure through greater consumption of such goods.

Hamermesh and Lee note that it is perfectly consistent, under such a theoretical approach, for people to feel more time stress as real incomes grow over time, because the goods restraint is relaxed but the time constraint cannot be, making the time constraint more binding. They find empirical support for this in the form of wealthier people being more likely to report feelings of time stress (2007: 374). However, this 'kvetching' is not seen as reflecting a sub-optimal situation: "... we assume that they are utility maximising but are simply unhappy about the limits on their available time" (Hamermesh and Lee, 2007: 382).

Less explicable within the utility optimising paradigm are the apparent negative 'externalities' that have been associated with long work hours. In the US rising workloads have been associated with stress, in turn leading to a range of physical health problems; sleep deficits; marital breakdowns; and 'parenting deficits' negatively impacting upon the country's youth (Schor, 1992: 11-13). Concerns of the effects of long work hours have also been expressed in the Australian context, including with respect to the 'double shift' experienced by women now working longer hours in paid work without a commensuration reduction in non-market work and caring duties (see, for example, Pocock, 2003; Pocock, Skinner and Pisaniello, 2010, Strazdins *et al.* 2011). Dockery (2003) presents evidence of lower life satisfaction for persons who report working more hours than they would prefer. On the other hand, Gray, Qu, Stanton and Weston (2004) found that a significant proportion of Australian fathers working long hours were satisfied with their hours of work, and this was associated with higher levels of wellbeing. Drago, Wooden and Black (2006) likewise highlight that long hour 'volunteers' and 'conscripts' need to be treated as two distinct groups.

[1] Figures cited from the ABS National Accounts (2011) are based on the author's own calculations using chain volume series as real GDP.

These phenomena, if true, call for an explanation in which individuals work longer hours than is consistent with their optimal wellbeing. Surely, this is not something a rational and utility maximising worker would do if we accept that the hours worked by individuals in countries such as the US and Australia are primarily a matter of free choice, albeit subject to some limitation in options. Is it possible that workers make errors of judgement about the balance of work and leisure that is consistent with optimal wellbeing? Frey (2008: chapter 11) presents evidence that it may indeed be the case that individuals systematically overestimate the utility gained from income and consumption, and underestimate the utility derived from 'experiential' or intrinsic attributes of goods (and leisure). Among these:

- people may underestimate how quickly they adapt to a higher income, such that it no longer provides heightened wellbeing;
- recollections of past events upon which people must base expectations of future utility may be distorted by an emphasis on intense or 'peak' moments, to the neglect of duration effects;
- people have a need to rationalise decisions, and therefore place greater emphasis on economic payoffs or material benefits. They will, for example, choose jobs offering the higher pay over ones with other intrinsic benefits such as flexibility in work hours or shorter commutes;
- finally, as has been discussed extensively in the 'happiness' literature (see, for example, Frank, 1999) consumption may be rivalrous such that it is one's income relative to others that matters for wellbeing. One person can increase their income and status by working longer hours, but the end result of such competition may be a zero-sum game of longer hours of work for no, or little, net increase in wellbeing.

Each of these perspectives suggests individuals may choose and/or prefer to supply more hours of labour than is consistent with maximising utility. This could not be detected in the neoclassical model that is based on the assumption that hours worked is both a choice variable and a point of utility maximisation, but may well be revealed through direct estimation of the utility function.

3. HILDA AND SOME DESCRIPTIVE DATA ON AUSTRALIAN WORKING HOURS

HILDA is Australia's first nationally representative household panel survey (see, http://melbourneinstitute.com/hilda/ for details on the survey and sampling frame). The panel was established through the 'Wave 1' interviews commencing in August 2001 and this paper uses data from the first 8 waves, covering annual surveys from 2001 through 2008. HILDA is based on a representative sample of Australian private households, and all individuals aged 15 and over in selected households come into its scope as responding persons. For each household, three different survey instruments are administered: a household form completed by one adult member of the household, a person questionnaire completed through face-to-face interviews with each in-scope household member and a self-completion questionnaire which respondents complete in private. Each year the panel consists of 7,000 plus households with around 13,000 responding individuals.

In addition to a wealth of information on individuals' socio-economic characteristics, attitudes, living arrangements, family background and labour market history, among other topics, HILDA collects detailed information on current employment arrangements. This includes hours worked and earnings. Workers are further asked how many hours they would prefer to work – is it about the same as they usually work, or more or less?

For the sample of employees pooled across waves 1 to 8, including males and females and both part- and full-time workers, the average of reported usual hours worked each week was a fraction over 36 hours, with that number varying little over the eight waves. A significant proportion of employees do indeed work long hours: 11.3 per cent report usually working between 50 and 59 hours, and a further 5.7 per cent report working 60 hours or more.

As one indication of how closely actual working hours match employees' preferences, figure 1 shows the distribution of actual hours worked for all employees. Next to these, the darker shaded bars show the distribution of working hours for only those who were happy with their work hours; that is, those who indicated their preferred hours were 'about the same' as their usual hours worked. It can be seen there is a remarkably strong concordance in the distributions: it appears that the number of hours worked each week closely matches the number of hours employees would like to work each week.

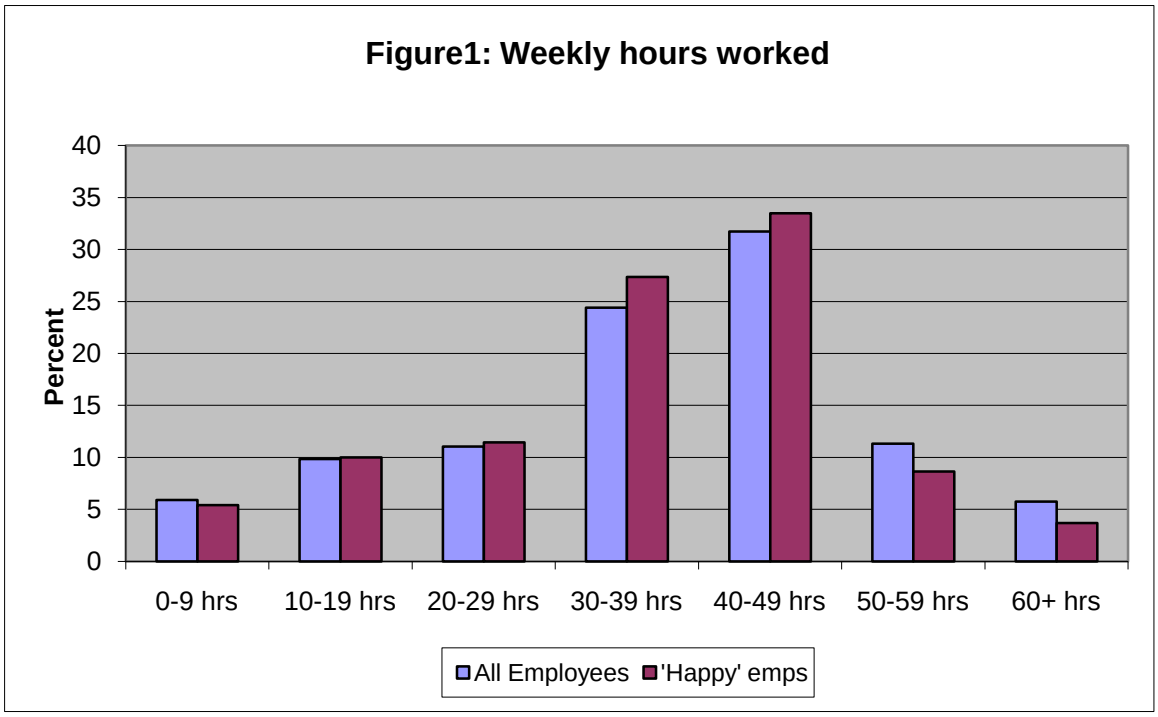

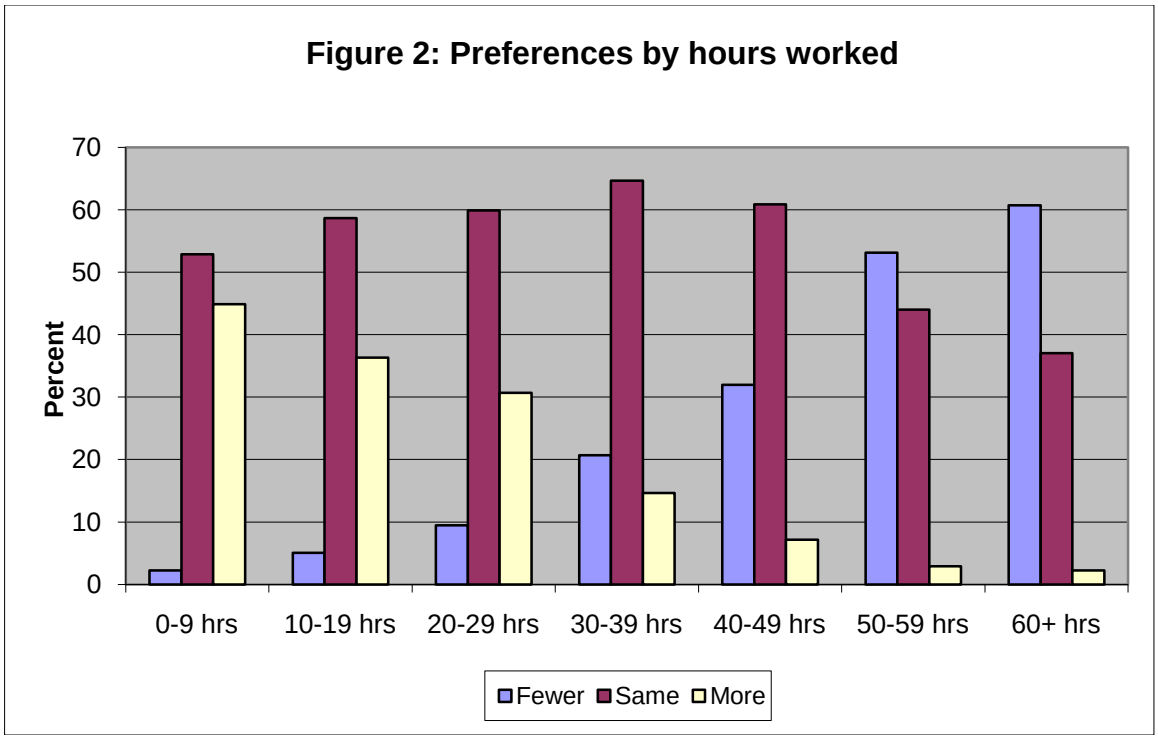

However, some evidence of mismatch emerges when we look at preferences conditional upon hours worked. In total 58 per cent reported wanting to work about the same number of hours as they do; while 26 per cent reported a preference for working fewer hours and 16 per cent more hours. Figure 2 shows that the incidence of mismatch is highest at the extremes of the hours distribution. For those usually working 0-9 hours, 45 per cent would prefer to work more hours. The majority of those working 60+ hours (61 per cent) and 50-59 hours (53 per cent) would in fact prefer to work fewer hours. This suggests some degree of inflexibility, and a divergence from 'optimal' labour supply. In particular, it suggests a substantial degree of over-work.

Finally, figure 3 draws on data from questions on how satisfied individuals were with various aspects of their jobs, in this case satisfaction with 'the hours that you work'. Responses were given on scale ranging from zero (completely dissatisfied) to 10 (completely satisfied). Satisfaction is highest for those working 30 to 39 hours. Evidence of a degree of 'over-work' can again be seen in workers' average satisfaction dropping off rapidly with hours worked beyond this category.

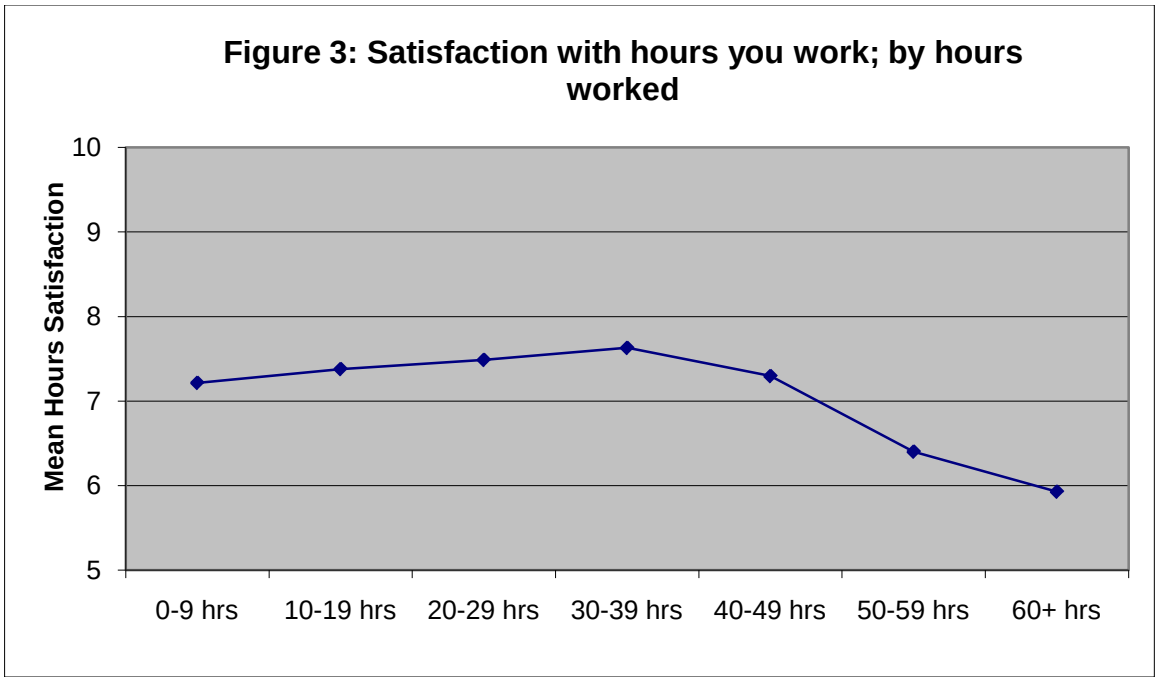

These descriptive statistics seem to suggest that it is not a lack of labour market flexibility in terms of the hours of work that employers can offer that results in Australians working non-preferred hours, since desired and actual hours are fairly well aligned in total. Instead, this disguises a problem of mismatch in which a significant proportion of people who are working long-hours are not the people who want to be doing so. This also holds for people working short working weeks, but to a lesser extent. There is evidence of both over- and under-employment. This would imply a lack of flexibility for people to adjust their hours of work to meet their preferences once they are in a job.

Note, however, that reported satisfaction with hours worked may underestimate the extent of the sub-optimality – or the loss of wellbeing - associated with excessive work hours. The arguments set out above, suggesting people may systematically overestimate the utility gained from more income and under-estimate the utility gained from leisure, would similarly imply that people would systematically over-estimate their optimal (or preferred) hours of work. Some working long hours may come to recognise this, and either adjust their hours or report a preference for fewer hours. However, others may report satisfaction with their hours or even a preference for more hours when they are already supplying more hours than is consistent with utility maximisation. The important assumption for this paper is that this is then reflected in their assessment of subjective wellbeing, even if it is not reflected in their reported preference for working hours.

4. DERIVING THE LABOUR SUPPLY CURVE BASED ON A SIMPLE MODEL OF SUBJECTIVE WELLBEING

The individual labour supply schedule maps out the number of hours of work a person is willing to supply conditional upon the wage rate. This can be derived from a simple 'happiness function' which includes income and leisure among it arguments. As in the neoclassical model, take a worker's utility (U) to be a function of a set of given individual characteristics (X), the number of hours of leisure (L) they have each week, and the quantity of market goods and services available to them, which is in turn directly dependent upon their weekly real income (Y).

$$U = X + \alpha \ln(L) + \beta \ln(Y) \tag{1}$$

Assume people need a minimum of 8 hours per day for necessities such as sleeping, eating and personal hygiene, leaving 112 potential work hours per week. Thus the weekly number of hours of leisure is (112-h), where h is hours worked per week. Income is a combination of unearned income (Y_u) and earned income (h x w, where w is the wage rate). Substitution into (1) gives the following utility function:

$$U = X + \alpha \ln(112 - h) + \beta \ln(Y_u + hw) \tag{2}$$

To identify the number of hours of work that will give the maximum level of utility, we differentiate (2) with respect to h.

$$dU/dh = \frac{-\alpha}{(112 - h)} + \frac{\beta w}{Y_u + hw}, \text{ and} \tag{3}$$

$$d^2U/dh^2 = \frac{-\alpha}{(112 - h)^2} - \frac{\beta w}{(Y_u + hw)^2} \tag{4}$$

Setting $dU/dh = 0$ in (3) and solving for h gives the utility maximizing number of hours, h*, as:

$$h^* = \frac{(112\beta w - \alpha Y_u)}{w(\alpha + \beta)}. \tag{5}$$

By assumption of utility being an increasing function of leisure and income (and confirmed empirically below) the parameters α and β are positive. Hence the second derivative given in (4) will be negative, confirming that h* is indeed a maximum.

Given data on hours worked, hourly wages, unearned income and utility for a sample of workers, the parameters α and β can be econometrically estimated. The analyst then has all the information required on the right-hand side of (5) to solve h* given the wage rate.

That is, the schedule of the utility maximising number of hours worked at each wage rate, or the labour supply curve.

This simple model treats the worker as an individual decision maker, who retains all earned and unearned income for the purposes of their own consumption. It therefore ignores possible interactions within a household and is most applicable to people who are not partnered. As noted below, it is relatively straightforward to extend the approach to couple households: the challenge lies in identifying the appropriate functional form for the utility function. Since the main purpose here is to demonstrate the basic approach, the following section generates the supply function and estimates of the elasticity of hours supplied with respect to changes in the wage using the less complex case of single male employees as an applied example.

5. EMPIRICAL ESTIMATION

A rapidly growing literature in economics studies the determinants of subjective wellbeing – happiness or life-satisfaction – and the challenges that these empirical findings pose for policy. The 'revolutionary' aspect of this development (Frey, 2008) lies in the rejection of economics' long-held belief that it is not possible to measure utility, but only to infer utility from revealed preferences. Thus the happiness literature relies strongly on the assumption that measures of subjective wellbeing can be used to make valid inferences about individuals' utility. While there are extensive and on-going arguments for and against this claim, it is not intended to revisit this debate here. The paper proceeds on the assumption that the available measure of subjective wellbeing is a valid construct which correlates with individuals' overall 'utility'.

That available measure is the rating of life satisfaction taken from the HILDA survey. The question on life satisfaction comes from the person questionnaire and is worded "All things considered, how satisfied are you with your life?" with respondents instructed to pick a number between zero (totally dissatisfied) and 10 (totally satisfied). Equation (2) above is in the form of a relatively standard 'happiness function' with the exception that the constraints on hours worked, income and hours of leisure for a given wage are explicitly imposed. Estimation of the model is possible since HILDA contains estimates of hours worked, the hourly earnings, unearned income and a wide range of other characteristics that may potentially impact upon wellbeing.

The sample is limited to unpartnered male employees (neither married nor living in a de facto relationship), who are aged 25 and over to abstract from participation in education; aged less than 65 to abstract from retirees; and without a long term health condition that limits the amount of work they can do. The vector of individual characteristics, X, includes age in 10 year categories, the presence of a disability and dummy variables for the wave of the survey the observation is drawn from. All monetary amounts are indexed by the consumer price index to be expressed in 2008 Australian dollars.

To obtain estimates of the parameters α and β in (1) and (2) above, models are estimated as random effects panel models to control for individual specific effects using two different specifications. One is the binary logit model, with the dependent variable collapsed to a dummy variable equal to one for people indicating eight to 10 on the life satisfaction scale, and zero for those indicating seven and below. Technically, this model estimates the probability of an individual indicating they are 'satisfied' as opposed to (relatively) 'dissatisfied' with their life as a whole. The other is a simple panel linear regression. Although this is clearly an inappropriate specification for a dependent variable bounded between zero and 10, results tend to be similar when such dependent variables are treated as ordinal variables as when the more technically correct logit or probit specifications are used (see, Kristoffersen, 2010). Although one specification models the dependent variable directly and the other the log of the odds of being 'satisfied', both imply a maximum level of utility when the right hand side argument of (2) is maximised with respect to hours worked. The regression results presented in table 1 show that the coefficients on the log of hours of leisure and the log of income can be estimated with some statistical precision under both specifications, and have the expected positive signs.

Table 1 - The utility function: panel regression estimates of life satisfaction for single male employees, HILDA waves 1-8

	Linear regression			*Logit model*		
	Coefficient		*P>|z|*	*Coefficient*		*P>|z|*
Intercept	5.206	***	0.00	-5.166	***	0.01
wave1	—			—		
wave2	0.068		0.30	0.047		0.79
wave3	0.146	**	0.03	0.472	***	0.01
wave4	0.094		0.16	0.302	*	0.09
wave5	0.175	***	0.01	0.296	*	0.10
wave6	0.095		0.17	0.099		0.59
wave7	0.055		0.44	-0.030		0.87
wave8	0.150	**	0.04	0.126		0.50
Aged:						
25-34 yrs	0.139	**	0.02	0.273	**	0.06
35-44 yrs	—			—		
45-54 yrs	0.123		0.07	0.133		0.43
55-64 yrs	0.571	***	0.00	1.441	***	0.00
Has a disability	-0.055		0.38	-0.498	***	0.00
Leisure (log hrs)	0.323	***	0.01	0.700	**	0.03
Income (log)	0.093	*	0.08	0.304	**	0.02
Observations	4684			4684		
Individuals	1606			1606		
Observations per individual						
Average	2.9			2.9		
Minimum	1			1		
Maximum	8			8		
Wald Chi-sq	55.3	***	0.00	62.9	***	0.00
R-squared:						
Within	0.008					
Between	0.019					
Overall	0.023					

Notes: ***, ** and * denote the estimated coefficient is significantly different to zero at the 1, 5 and 10 per cent level, respectively.

Table 2 presents the solutions for h* as given by (5) using the estimates for α and β from each of the models (that is, the estimated coefficients for the log of hours of leisure and log of income, respectively). These are calculated using the sample means for the hourly wage and unearned income. Table 2 also presents the sample mean for actual hours worked. The estimates imply a utility maximising working week of 23 hours and 32 hours, respectively, using the linear regression and logit model specifications. In each case, the second derivative evaluated at the means is negative, implying the solution at h* is a local maxima. Both of these estimates are well below the actual average of hours worked by single, male employees of 43 hours per week.

Table 2 - Implied optimal hours of work (h*) and elasticity of labour supply (e), single Australian men aged 25-64

	Estimated coeff's		*Sample means*			*Derived parameters*	
	Leisure (α)	*Income (β)*	*Real hourly wage*	*Weekly unearned income*	*Weekly hours worked*	*Optimal hours (h*)*	*Elasticity (e)*
Linear regress.	0.323***	0.093*	$25.76	$69.28	43.0	**23.0**	0.09
Logit model	0.700**	0.304**	$25.76	$69.28	43.0	**32.0**	0.06

Notes: ***, ** and * denote the estimated coefficient in the relevant model is significantly different to zero at the 1, 5 and 10 percent level, respectively.

The elasticity of hours worked with respect to the real wage (e) is typically presented in the form of the percentage change in hours worked that results from a one per cent increase in the wage. One such estimate of e can be derived directly from (5) by calculating the percentage change in h* when the wage rate is evaluated at the sample mean ($25.76) and at the sample mean plus one per cent ($26.02). These estimates are presented in table 2, and are close to zero under both models, suggesting that the substitution and income effects of an increase in the real wage largely offset one another.

The elasticities presented in table 2 are calculated at mean hourly wages. Noting that the logit models generally produce estimates of h* more in line with observed hours worked, figure 4 maps out the labour supply curves (or h*) for hourly wages at $10 intervals up to $100 and based on the coefficients estimated from the logit model. Note the supply curve indicates a very inelastic supply response as wages increase beyond a wage rate of around $40 per hour.

A further way to calculate the elasticity using data from the full sample is to use equation (5) to derive h* for each individual based on their own values of the variables in the vector X and their own wage, and then recalculate h* assuming a one per cent increase in their wage. Using the estimated α and β from the logit model, this approach suggests a wage elasticity of supply of 0.24. That is, if all men in the sample were to receive a one per cent real wage rise, then on average their utility maximising hours of work would increase by 0.24 of one per cent.

6. SOME EXTENSIONS AND AN APPLIED EXAMPLE: THE IMPORTANCE OF JOB SATISFACTION

The previous sections set out a very simple model and fitted it with data from a convenient, relatively homogenous, sample to demonstrate how the individual labour supply curve can be derived through estimation of an explicit utility function. Many potential refinements and extensions to the model are readily apparent, an obvious one being theoretical and empirical development of the most appropriate functional form to use for the utility function. Some others are:

1. Developing models for persons in couple households and incorporating the effects of dependent children. In addition to assumptions on the sharing of earned and unearned income within the household, extending the analysis to households presents added complications regarding the form of the utility function. Is it related to own income-share and leisure, or does a partner or parent derive utility from the consumption and leisure enjoyed by other household members? Are there synergies in utility from sharing leisure and goods and services?[2]
2. Further moving to a fuller model of household labour supply that takes account of hours worked by partners. Breunig, Cobb-Clark and Gong (2008) find within-household correlations in preferences for working hours and leisure to be important.
3. Incorporating non-market work, such as housework, although Becker (1965) points out that the distinction between non-working time that is devoted to leisure and that to home production is not so clear-cut.
4. Workers who are not employed, including the unemployed, can be readily included within this framework, although this would require imputing a wage rate.
5. The current analysis has not accounted for tax rates – the estimates are based on gross weekly earnings. With sufficient information on the tax and benefit parameters, it would be straightforward to model tax-transfer policies as changes in earned and unearned income (own and partner's).
6. Allowing for the effects of job quality on utility and labour supply.

I take up the last of these as one further example of an application of the approach. HILDA includes a range of variables relating to individuals' self-assessed satisfaction with different aspect of their jobs. To demonstrate the importance of workers' job satisfaction upon their willingness to supply labour, the sample of single male employees was split into those with relatively high job satisfaction and those with relatively low job satisfaction. This was on the basis of their response to the question "all things considered, how satisfied are your with your job" given on a scale from zero (totally dissatisfied) to 10 (totally satisfied). Separate models were estimated for those who reported six or lower (21.5 per cent of the sample) and those who reported seven or above (78.5 per cent).[3]

---

[2] Full consideration of these issues is beyond the scope of this paper. In preliminary work, an arbitrary assumption was applied for couples in which all income is equally shared between the two, and utility is simply a function of own leisure time and one's own share of income. This can be simply be reflected by incorporating partner's earned income as part of Y_u and rewriting the utility function given in (2) as:

$$U = X + \alpha \ln(112 - h) + \beta \ln[(Y_u + hw)/2] \qquad (6)$$

Since $\beta \ln\left[\frac{Y_u+hw}{2}\right] = \beta \ln(Y_u + hw) + \beta \ln(0.5)$ this is the same as adding a constant to the utility function, and does not affect the derivatives. The algebraic solution for h* remains as that given in (5). For calculation of the actual value, however, this makes an important difference in that Y_u for the individual is taken to include their partner's earned income, and so varies greatly for partnered men and women. Initial estimates consequently suggest much lower preferred working hours for married women, but continue to show 'excessive' work hours for all groups.

[3] A more even split between those reporting 0-7 and 8-10 returned implausible parameters in some models. Experimentation with entering job satisfaction, and job satisfaction weighted by hours as additional arguments in the utility function has proven unsuccessful to date. In doing this, job satisfaction was

The full regression results are reported in the appendix, and table 3 presents the derived utility maximising level of hours worked.

Table 3 - Implied optimal hours of work (h*) and elasticity of labour supply (e), single Australian men aged 25-64, by level of job satisfaction

	Estimated coefficients		*Derived parameters at sample means (w=$25.76, Y_u=$69.28)*		*Mean Weekly hours worked*
	Leisure (α)	*Income (β)*	*Optimal hours (h*)*	*Elasticity (e)*	
Linear regression					
Low job satisfaction	0.612**	0.038	**3.9**	0.6	43.1
High job satisfaction	0.143	0.041	**23.0**	0.1	43.0
Logit model					
Low job satisfaction.	2.192**	0.194	**6.6**	0.4	43.1
High job satisfaction	0.486	0.191	**29.6**	0.1	43.0

Notes: ***, ** and * denote the estimated coefficient in the relevant model is significantly different to zero at the 1, 5 and 10 percent level, respectively.

There is now considerable imprecision in the estimates of α and β, and consequently the derived parameters for h* and e need to be treated as being largely for illustrative purposes. However, the broad pattern of the results are in accordance with expectations in several important respects, offering general support for the approach, and suggesting a very significant role of job satisfaction in determining labour supply. Surprisingly, there is almost no difference in the average number of actual hours worked each week between male employees with relatively low job satisfaction and those with relatively high job satisfaction – both groups report working 43 hours per week. However, under both the linear and logit specifications, the estimated α coefficient implies that persons who are relatively dissatisfied with their jobs have a much greater distaste for hours worked or, put another way, a much greater preference for hours of leisure. This could arise through either individual traits (these individuals do not like work or place a high value upon leisure) or a job specific effect (these individuals are working in low quality, unpleasant jobs).

When h* is derived using these coefficients, the result is that those with higher job satisfaction are estimated to have a much higher number of working hours consistent with utility maximisation. To abstract from other effects, h* is calculated using the mean values of the wage and unearned income for the full estimation sample of single male employees. For those with low job satisfaction, optimal working hours are calculated to

measured *relative* to individual's mean satisfaction rating with six other aspects of their lives (their home, financial situation, how safe they feel, feeling part of the local community, health and the neighbourhood in which they live), to avoid endogeneity associated with more 'positive' people and/or those in positive moods reporting higher satisfaction with their jobs and with their life overall (the dependent variable). Thus, the measure was constructed as satisfaction with their job relative to average satisfaction in other life domains.

approximately five hours per week.[4] This compares to estimates of 23 hours and 30 hours, depending upon model specification, for workers reporting high job satisfaction. Further, as would be expected, the labour supply of those with low job satisfaction is estimated to be much more responsive to changes in the wage rate. The estimated labour supply curve remains quite inelastic with respect to changes in the wage (figure 4). For both groups, however, observed hours worked remain in excess of that consistent with welfare maximisation – dramatically so in the case of those with low job satisfaction.

Figure 4 - Utility maximising labour supply curves: All employees, those with low job satisfaction and those with high job satisfaction (based on logit model estimates)

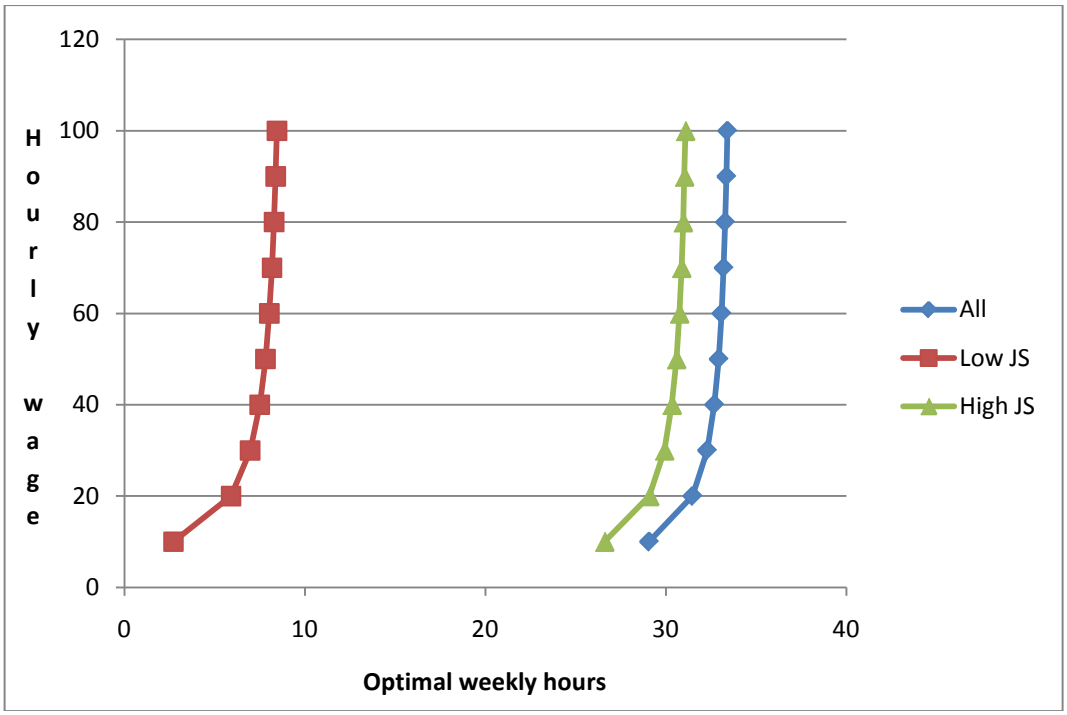

7. CONCLUSION

To the best of the author's knowledge, this paper proposes a new empirical approach for estimating the individual labour supply curve and the elasticity of labour supply with respect to changes in the wage. In contrast to the normal assumption of an implicit utility function and the individual choosing their utility maximising labour supply, the approach taken here is to explicitly estimate the utility function conditional upon income and hours of leisure. It can account for two potential shortfalls of the traditional neo-classical

[4] While this seems very low, it should be noted that due to the interaction in the tax and benefits system in Australia, many people in low paying jobs – and those which are likely to be of poor quality – can face very high replacement rates, and hence suffer relatively minor financial penalty from not working or from working few hours as opposed to working full-time (see, Dockery, Ong and Wood, 2011).

approach that might lead to observed hours deviating from utility maximising hours: 1. individuals generally do not have the flexibility to decide on the number of hours they work; and 2. Individuals may not accurately identify their utility optimising work hours. To address this second problem, it is assumed that 'sub optimal' choices will be reflected in lower reported subject wellbeing. In particular, there are reasons to believe that individuals systematically choose to work longer hours than is consistent with optimal wellbeing.

The analysis contained in this paper should be considered very much exploratory, presenting a very basic model. Still, the initial results suggest considerable promise. For single male employees the results imply an elasticity of hours worked with respect to a change in the real wage of close to 0.1, well within the realm of empirical estimates found in other studies (see, Blundell and MaCurdy, 1999, table 1). It implies a relatively inelastic labour supply curve in which the income effects of a rise in wages (greater demand for leisure) largely offset the substitution effects of a rise in wages (the higher opportunity costs of an hour of leisure). The result is also consistent with the expectation that, in the long run, rising real incomes should not lead to increased hours of work.

The empirical results for the sample of single, male Australian employees support the common perception that a problem of 'over-work' exists in today's Australian society. The estimates for optimal hours per week: 23 hours or 32 hours depending upon specification, compare to average reported actual working weeks in excess of 40 hours. The estimates for optimal hours are also below the duration of standard working weeks, which in Australia tend to be 35 hours, 37.5 hours or 40 hours per week. Preliminary estimates for other groups - single women, married men and married women - all return estimates of optimal hours below actual hours worked and below the standard working week. Hence limited flexibility in hours choice may well be one significant cause of excess work. However, it may also be the case that Australians are going the way of the 'Overworked American', pursuing consumption and status at the expense of intrinsic rewards that would contribute more to their own happiness, and the happiness of their families. Evidence is also presented to show that individuals' satisfaction with their job has a far greater effect on the hours of labour they would prefer to supply than does the wage. An implication of this finding is that to increase work effort, it may be far more profitable for many employers to invest in policies to improve the quality of jobs and employee satisfaction than to pay higher wages.

The approach could well be challenged with the charge that, for policy relevance, the estimated labour supply curve should approximate what people actually do, not what they would like to do. It is their realised choices that matter, not some theoretical optimum. In part, this can be treated as a purely empirical challenge: which approach fits data better and has better predictive power? Importantly, the approach presented here may be more consistent with long run equilibrium changes in the labour market, rather than short-term responses to changes in the wage. Further, we should not lose sight of the point that maximising people's wellbeing should be the objective of policy.

APPENDIX

Utility functions: panel regression estimates of life satisfaction for single male employees by level of self-assessed job satisfaction, HILDA waves 1-8

	Linear reg.				*Logit model*			
	Low Job Satisfn.		*High Job Satisfn.*		*Low Job Satisfn.*		*High Job Satisfn.*	
	Coef.	*P>|z|*	*Coef.*	*P>|z|*	*Coef.*	*P>|z|*	*Coef.*	*P>|z|*
Intercept	3.391	0.05	6.619	0.00	-12.328	0.02	-3.104	0.14
wave1	—		—		—		—	
wave2	0.269	0.11	-0.021	0.77	-0.031	0.95	-0.035	0.86
wave3	0.330	0.04	0.102	0.15	0.103	0.82	0.522	0.01
wave4	-0.121	0.48	0.132	0.06	-0.765	0.13	0.467	0.03
wave5	0.084	0.64	0.106	0.14	-0.559	0.28	0.332	0.11
wave6	0.259	0.14	0.026	0.72	-0.521	0.30	0.131	0.54
wave7	-0.153	0.41	0.035	0.64	-0.743	0.16	-0.042	0.85
wave8	0.318	0.09	0.066	0.38	-0.129	0.81	0.073	0.73
Aged:								
25-34 yrs	0.275	0.04	0.102	0.10	0.128	0.73	0.312	0.06
35-44 yrs	—		—		—		—	
45-54 yrs	0.276	0.08	0.016	0.82	0.339	0.44	0.009	0.96
55-64 yrs	0.627	0.03	0.469	0.00	1.273	0.09	1.362	0.00
Has disability	-0.009	0.96	-0.056	0.40	-0.490	0.32	-0.541	0.00
Leisure (log hrs)	0.612	0.03	0.143	0.28	2.192	0.02	0.487	0.17
Income (log)	0.037	0.75	0.041	0.45	0.194	0.55	0.191	0.18
Observations	998		3686		998		3686	
Individuals	588		1428		588		1428	
Obs per indiv.								
Average	1		1		1		1	
Minimum	1.7		2.6		1.7		2.6	
Maximum	8		8		8		8	
Wald Chi-sq	31.8	0.00	39.8	0.00	15.6	0.27	55.4	0.00
R-squared:								
Within	0.046		0.007					
Between	0.020		0.017					
Overall	0.031		0.017					

REFERENCES

Australian Bureau of Statistics (ABS) (2011), *Australian National Accounts: National Income, Expenditure and Product*, Catalogue No. 5206.0, on-line time series spreadsheets, accessed 4 July 2011.

Becker, G.S. (1965), 'A theory of the allocation of time', *Economic Journal*, 75, 299: 493-517.

Blundell, R. and McCurdy, T. (1999), 'Labor supply: A review of alternative approaches', in Ashenfelter, O. and Card, D. (eds.) *Handbook of Labor Economics*, 3A chap 27, Amsterdam: Elsevier Science.

Breunig, R., Cobb-Clark, D. and Gong, X. (2008), 'Improving the modelling of couples' labour supply', *Economic Record*, 84(267), 466-486.

Dockery, A.M. (2003), 'Happiness, life satisfaction and the role of work: evidence from two Australian surveys, School of Economics and Finance Working Paper No. 03.10, Curtin Business School, November.

Dockery, A.M., Ong, R. and Wood, G. (2011), 'Measuring work disincentives: Taxes, benefits and the transition into employment', *Australian Journal of Labour Economics*, 14(3), in-press.

Drago, R., Wooden, M. and Black, D. (2006), 'Long work hours: Volunteers and conscripts', Melbourne Institute Working Paper Series No. 27/06, Melbourne Institute of Applied Economic and Social Research.

Frank, R.H. (1999), *Luxury Fever*, The Free Press: New York.

Frey, B. (2008), *Happiness: A revolution in economics*, Massachusetts Institute of Technology.

Gray, M., Qu, L., Stanton, D. and Weston, R. (2004), 'Long work hours and the wellbeing of fathers and their families', *Australian Journal of Labour Economics*, 7(2), 255-273.

Hamermesh, D.S. and Lee, J. (2007), 'Stressed out on four continents: Time crunch or yuppie kvetch?', *Review of Economics and Statistics*, 89(2), 374-383.

Kristoffersen, I. (2010), 'The Metrics of Subjective Wellbeing: Cardinality, Neutrality and Additivity', *Economic Record*, 86(272), 98-123.

Pocock, B. (2003), *The work/life collision*, The Federation Press: Sydney.

Pocock, B., Skinner, N. and Pisaniello, S. (2010), *How much should we work? Working hours, holidays and working life: the participation challenge*, The Australian Work and Life Index 2010, Centre for Work + Life, University of South Australia.

Schor, J.B. (1992), *The overworked American: The unexpected decline of leisure*, Basic Books: New York.

Strazdins. L., Griffin, A., Broom, D., Banwell, C., Korda, R., Dixon, J., Paolucci, F. and Glover, J. (2011), 'Time scarcity: another health inequality?', *Environment and Planning A*, 43, 545-559.