Deriving the labor supply curve from happiness data[1]

Alfred M. Dockery
Centre for Labour Market Research, Curtin University

Abstract: An alternative empirical method to estimating the labor supply function is proposed, based upon subjective wellbeing data. It potentially addresses limitations of the standard neo-classical approach by allowing workers' observed hours worked to deviate from their utility maximizing point.

JEL Codes: J22 (Time allocation and labor supply)

Key words: Subjective wellbeing, labour supply, working hours

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.

[1] Corresponding author: Alfred M. Dockery, Centre for Labour Market Research, Curtin University, GPO Box U1987, Perth WA 6845, Australia. Phone: +61 8 9266 3468; Fax: +61 8 9266 3026; Email: m.dockery@curtin.edu.au;

Deriving the labor supply curve from happiness data

1. Introduction

The neo-classical model of labor supply views an individual's preferences as being described by an underlying utility curve that is a function of hours of leisure and the consumption of goods and services. Consumption is made possible through income earned from working. The utility curve is not estimated directly, but inferred from the assumption that individuals choose the number of hours of work to maximize their utility. While there have been many extensions to the basic model, such as the incorporation of household production, and life-cycle considerations, all empirical approaches essentially rely on an implicit utility curve and the assumption that observed hours of work represent the utility maximizing solution (see Blundell & McCurdy 1999).

There are at least two reasons to expect that actual hours worked may not represent the worker's utility optimising solution. Most obviously, taking into consideration the demand side of the labor market, individuals are not free to choose any number of hours of work. More importantly, evidence emerging from 'happiness' research (Frey and Stutzer 2002) gives weight to the proposition that individuals make decisions that are inconsistent with utility maximization. In particular, they may systematically overestimate the utility gained from consumption and status, and systematically underestimate 'intrinsic' benefits from leisure and time with friends and family (Frank 1999; Frey 2008: 127-137).

This paper proposes an alternative approach to estimating the labor supply curve that accounts for both these limitations. Rather than rely on an implicit utility function, the general proposition is to directly estimate a 'happiness model' (utility function) conditional upon leisure and income. For each wage level, the parameters from that model can be used to calculate the utility maximizing number of working hours.

2. A simple model

Following the neoclassical model, take a worker's utility (U) to be a function of a set of individual characteristics (X), weekly hours of leisure (L), and the consumption of goods and services, which is in turn determined by weekly real income (Y).

$$(1) \quad U = X + \alpha \ln(L) + \beta \ln(Y)$$

Assume people need a minimum of 8 hours per day for necessities such as sleeping, eating and personal hygiene, leaving 112 hours per week to be divided between work (h) and leisure (L=112-h). Income comprises of unearned income (Y_u) and earned income (h x w, where w is the real wage rate). Substituting into (1) gives:

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

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

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

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

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

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

By assumption of utility being an increasing function of leisure and income (confirmed empirically below) the parameters α and β are positive. Hence the second derivative given in (4) is negative, confirming h^* is 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 obtained through econometric estimation of the utility function (2). This provides all the information required on the right-hand side of (5) to solve h^* for a given wage rate. This gives the schedule of the utility maximizing number of hours of work at each wage rate - the labor supply curve.

3. Empirical results: a worked example

The model set out above is estimated using data from the first 10 waves (2001-2010) of the Household, Income and Labour Dynamics in Australia survey (HILDA), Australia's first nationally representative household panel survey (see http://melbourneinstitute.com/hilda/).

Equation (2) 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, earnings, income from other sources (unearned income) and subjective wellbeing. The key measure of subjective wellbeing is the individual's response to the question: "All things considered, how satisfied are you with your life?" on a 0 (totally dissatisfied) to 10 (totally satisfied) scale. Underpinning the happiness literature is the belief that such measures of subjective wellbeing – happiness or life-satisfaction – can be used to make valid inferences about individuals' utility. This paper proceeds on that assumption, though it is acknowledged that there are arguments for and against this claim (see Layard 2003).

To obtain estimates of the parameters α and β, the sample is limited to unpartnered male employees (neither married nor living in a de facto relationship); 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, age-squared and the presence of a disability. All monetary amounts are indexed by the consumer price index to be expressed in 2010 Australian dollars.

Equation (2) is estimated as a fixed effects panel model to control for unobserved heterogeneity between individuals. For ease of exposition, the specification is simple linear regression. Although this is technically inappropriate for an ordinal dependent variable bounded between 0 and 10, results tend to be very similar whether such dependent variables are treated as cardinal or the more technically correct ordered logit or probit specifications are used (see Kristoffersen 2010, Ferrer-i-Carbonell and Frijters 2004). The regression results presented in Table 1 show that the coefficients on (log) hours of leisure and income can be estimated with some statistical precision and have the expected signs.

For the estimation sample, the mean number of hours worked each week is 43 hours. As shown in Table 2, the solution for the utility maximizing hours of work is around two-and-a-half hours lower, at 40.7 hours. This solution for h* is calculated from (5) using the estimated coefficients α and β and evaluated at the sample means for the real wage and unearned income. The elasticity of hours worked with respect to the real wage can be derived directly from (5) by calculating the percentage change in h* when the wage rate is evaluated at the sample mean ($27.75) and at the sample mean plus 1 per cent ($28.02). The very small estimate of 0.04 suggests that the substitution and income effects of an increase in the real wage largely offset one another.

Figure 1 maps out the labor supply curve (or h*) for hourly wages at $5 intervals. It indicates a very inelastic supply response as wages increase beyond around $30 per hour and, no matter how high the wage rate, the estimated utility maximising hours does not reach the mean hours actually observed. The results are of course sensitive to the estimated coefficients α and β. Estimation of a random effects model, which exploits variation in hours of work across individuals, but may suffer from bias due to unobserved heterogeneity, returns a markedly lower estimate for h* of around 31 hours.

Table 1: The utility function: fixed-effects panel regression estimates of life satisfaction, single male employees

	Coefficient	P>|z|	Variable mean
Intercept	5.320	0.00	—
Age	-0.041	0.21	39.41
Age-squared/100	0.070	0.07	16.60
Has a disability	-0.047	0.47	0.09
Leisure (log hrs)	0.310	0.07	4.22
Income (log)	0.188	0.01	7.01
Observations	5904		
Individuals	1850		
Observations per individual			
Average	3.2		
Minimum	1		
Maximum	10		
F(5,1849)	3.86	0.00	
R-squared:			
Within	0.007		
Between	0.007		
Overall	0.014		

Table 2: Implied optimal hours of work and elasticity of labor supply, based on sample means.

Estimated coefficients		Sample means			Derived parameters	
Leisure (α)	Income (β)	Real hourly wage	Weekly unearned income	Weekly hours worked	Optimal hours (h*)	Elasticity
0.310	0.188	$27.75	$72.10	43.03	40.69	0.04

Figure 1: Utility maximizing labour supply curve: single men

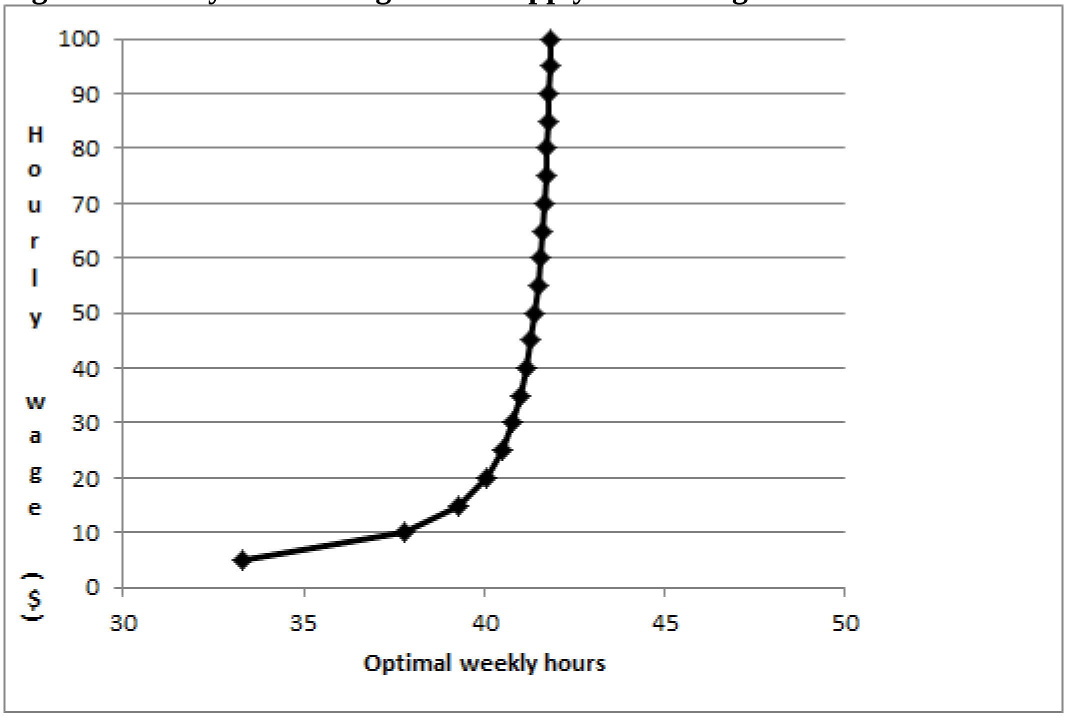

3.1 Potential extensions

For illustrative purposes, a simple model has been set out and estimated using a conveniently homogenous sample. However, the logic of the approach is general: estimate an explicit utility function that embodies the trade-off between leisure and income, and solve for the utility-maximizing hours of work. Many potential extensions are readily apparent, including theoretical and empirical development of the most appropriate functional form for the utility function. For complex functional forms in which the first and second derivatives are difficult to derive algebraically, maxima can be recovered numerically from data. Other potential extensions include:

- Developing models for persons in couple households and with dependent children.
- Incorporating non-market work, such as housework.
- The unemployed and persons outside of the labour force can readily be included within this framework, although this would require imputing a wage rate.
- Accounting for the effects of job quality on labor supply.

4. Conclusion and discussion

This paper proposes a new empirical approach for estimating the individual labor supply curve. It can account for two potential shortfalls of the standard approach that might lead to observed hours deviating from utility maximizing hours: inflexibility in the choice of hours worked and workers' inability to identify their utility optimising work hours.

It is in addressing this second problem that the approach offers considerable potential. There are reasons to believe that individuals systematically choose to work longer hours than is consistent with optimal wellbeing. In addition to the insights from happiness research, there is considerable evidence of ‘overwork’ leading to poor health outcomes and other negative externalities (see for example Schor’s 1992 *The Overworked American*). That such ‘excessive’ working hours are observed in advanced economies in which real incomes have doubled and even trebled over recent decades is a paradox difficult to reconcile within the neo-classical model. As with single males above, preliminary estimates for other groups imply ‘excessive’ working hours, and yield intuitively appealing results: workers reporting high job satisfaction have a much higher h*, and married women a low h*, due largely to their higher unearned income in the form of partners’ wages.

The approach could well be challenged with the charge that, for policy relevance, the estimated labor 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 labor market, rather than short-term responses to changes in the wage. Further, we should not lose sight of the point that maximizing people’s wellbeing should be the objective of policy.

References

Blundell, R., McCurdy, T., 1999. Labor supply: A review of alternative approaches, in: Ashenfelter, O.C., Card, D. (Eds), Handbook of Labor Economics. Elsevier Science, Amsterdam, pp. 1559-1695.

Ferrer-i-Carbonell, A., Frijters, P., 2004. How important is methodology for the estimates of the determinants of happiness? The Economic Journal 114(497), 641-659.

Frank, R. H., 1999. Luxury Fever, The Free Press, New York.

Frey, B.S., 2008. Happiness: A revolution in economics. The MIT Press, Cambridge, Massachusetts.

Frey, B.S., Stutzer, A., 2002. What can economists learn from happiness research? Journal of Economic Literature 40(2), 402-435.

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

Layard, R., 2003. Happiness: has social science a clue? Lionel Robbins Memorial Lecture 2002/3 – Lecture 1. London School of Economics, March.

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