INFOMETRICS

Department of Labour

Labour Market Policy Group

LABOUR SUPPLY IN NEW ZEALAND

– A Preliminary Analysis Using Unit Record Data from the Household Labour Force Survey

Adolf Stroombergen and Tom O'Brien

Infometrics

September 2003

ACKNOWLEDGEMENT: Access to the data used in this study was provided by Statistics New Zealand under conditions designed to give effect to the confidentiality provisions of the Statistics Act 1975.

DISCLAIMER: The results presented in this paper are the work of Infometrics not Statistics NZ.

Executive Summary

This paper presents some preliminary findings of a short study of factors that affect labour force participation in New Zealand. It is intended to be a stimulus for further research as it leaves many questions that are beyond the scope of this project unanswered. The study uses data from the Household Labour Force Survey (Income Supplement), and has the associated aim of ascertaining the usefulness of this database for future research into the New Zealand labour market.

From a brief overview of the relevant literature, it is clear that the decision to participate in the labour market and the decision on how much to participate (hours of work) are affected by different factors, and/or affected in different ways by similar factors. The literature also suggests that the former decision is likely to show the greater sensitivity to the usual factors that have been shown to influence participation, particularly amongst women. For this reason, and taking into account the constrained scope of this project, we have chosen to focus on the former decision.

Confirming previous studies, our research demonstrates the importance of factors such as age, ethnicity, dependent children and so on. These all have effects in accordance with theory. What is somewhat surprising is the high degree of explanatory power that these non-monetary variables contain in determining whether an individual is in or out of the labour force. Certainly variables such as wages, non-labour income and benefits are statistically significant, but their contribution to explaining labour market status is relatively small.

While further testing is required, it seems safe to infer that the use of unit record data allows the non-monetary factors in particular to deliver an improved signal-noise ratio compared to the use of averaged data, with which aggregation bias is a well-known problem. The second advantage of the unit record dataset is that is has enabled the identification of people in the same household and thus to allow (albeit in a relatively crude way) for non-independent observations – because of joint decision making. Fortunately (for research where this is not possible) and in accordance with theory, it seems that the presence of correlated errors does not seriously bias the coefficients, although it does lead to inefficiency. Allowing for non-independence does, however, improve the overall explanatory power of the model, in this case quite markedly.

There is considerable scope for further research, including testing the models over other time periods, refining the specification of joint household decision making with respect to labour force participation, and investigating the effects of factors such as industry and occupation. The issue of how many hours to work, given a decision to work, is another desirable extension. In fact there are so many avenues for the research which are potentially both interesting and fruitful, that we would welcome comment on what would be of most value for a future research programme.

The paper has three sections. Section 1 provides an overview of the literature, section 2 proposes a reduced-form system of equations, and section 3 presents a preliminary attempt at estimating the equations

Preface

The focus of this research is on identifying the factors that affect the decision on whether or not to supply labour, and how they differ across different individuals. Unit record data in cross-sectional form is the basis of the investigation. The data is derived from the Household Labour Force Survey (Income Supplement), undertaken by Statistics New Zealand.

Due to budget constraints and unexpected data issues this study is by no means intended to constitute a seminal step forward in the analysis of labour supply in New Zealand. It is a merely a preliminary step in understanding the usefulness of the HLFS data source for understanding New Zealand labour market issues. It will hopefully act a stimulus for further research.

1. Literature Overview

Studies using unit record data are less common than aggregate-based time series studies. Many that do exist are rather dated. While this does not matter much from the perspective of placing labour supply into some theoretical framework, the earlier datasets tended not to be as rich as the LMPG dataset,[1] and some of the econometric estimation was of a lower quality than might now be expected. Hence we can probably expect only limited guidance on how to specify and estimate a labour supply function from other studies.

We have perused a number of articles and publications, but we limit the review below to a selection of about a dozen studies, split evenly across New Zealand research and international research. Some studies have been selected for their empirical content and others for theoretical value.

1.1 New Zealand Research

A useful introduction to the whole field from a New Zealand perspective is provided by Prebble and Rebstock (1992) which contains five articles from a number of authors. The first article introduces the reader to labour supply theory – budget constraints, replacement rates etc, and briefly discusses some empirical research findings. We return to these below. The second article is primarily a description of the Treasury's Taxmod microsimulation model. While interesting it is not particularly relevant to our study.

Effective marginal tax rates are the focus of the third article. It confirms that some EMTRs in New Zealand are around 90-100% and discusses how such rates may be a serious disincentive to move off a benefit. While this is undoubtedly true, the authors seem to place too much emphasis on the EMTR relative to the replacement rate which is arguably a more important concept, as it reflects the whole budget constraint, not just a number of selected points. This issue is taken up in the fourth article which calculates replacement rates for a number of different combinations of employment status, benefit income and demographic characteristics. The key findings are that replacement rates differ quite markedly with different circumstances, that they are generally higher than in Australia, and that the replacement rate

[1] Having said that, our ability to study how people behave over time in response to changing labour market conditions cannot be examined with the LMPG dataset in its current form. In theory it is possible obtain some longitudinal information from the dataset. However, the possibility of doing this has been removed from the data that was made available to us.

for people leaving a benefit tends to be higher than for people leaving employment (and entering a benefit). This is because productivity falls with time out of employment. However, it is also possible that people who spend proportionately long amounts of time out of work are inherently less productive; that is the causation runs the other way.

The final paper presents some econometric estimation. A discrete choice model is estimated using a multinomial logit approach (discussed below), incorporating a correction for selection bias. Unfortunately the different labour market choices are defined in terms of annual income, not hours of work. Also, wage rates are constructed as annual earnings divided by an assumed 40 hours per week. This can introduce large measurement errors. Furthermore, wages for people not in employment are estimated on the basis of personal characteristics, given an estimated relationship between wages and these characteristics for people in employment. This is where the selection bias correction factor is relevant, but it does nothing to overcome the problem of using 'actual' wages and estimated wages in the same equation, which can lead to inefficient and biased estimates. Overall then, the estimates are unlikely to be robust, but the paper does clearly exemplify the difficulties in estimation labour supply functions.

Dixon (1996) examines quarterly data from the Household Labour Force Survey (HLFS) for the period 1986-1996. The methodology of the paper consists of graphing data and discussing results. Employment and participation are graphed by sex, ethnicity and age. Also examined is participation by marital status, time since last worked, and family circumstance.

The article highlights the fact that age, sex and ethnicity play a large role in determining whether an individual seeks work. Participation of those aged 15-19 and 55+ was lower than in the prime working years of 20-54. Education and retirement explain the lower participation at these ages.

Male participation rates were high and fairly constant between the ages of 20-54. Female participation rates were on the whole lower than for males, except for the 15-19 and 60-64 cohorts. In the child-bearing ages of 25-34, the female participation rate dips, before rising to a peak in the 45-49 range and then quickly tapering off.

Those with post-school qualifications were more likely to be participating in the workforce than those with no qualifications.

Men who are married or living with de facto partners are more likely to participate in the labour force than men who are divorced, separated or single. Among females, those who were never married had the highest labour force participation rate. This suggests that marital status and the work status of the partner (if there is one) are important factors in determining whether to look for work.[2] We look at this in section 3.

Maori participation rates were lower than European rates across all cohorts, for both sexes. Of those who were of prime age (25-54) and not in the work force, most had been out of work for over two years. Lifestyle choices and attitudes to work, rather than simply not being able to get a job, could be a factor here, but it should be noted that education and other characteristics (such as region), were not controlled for.

Brooks (1990) estimates an error correction model of aggregate labour force participation using data from the Quarterly Employment Survey and data on registered unemployment.

[2] The 1991 Census showed that 55% of females who were a parent in a two parent family did not work (and were aged 25-54).

Employment is defined as full-time equivalent employment (FTE) and the labour force is defined as full time equivalent employment plus registered unemployed. The data covered the period 1965-1988, but a structural break meant that the period tested was 1978-1988.

Brooks estimates separate "labour force participation rate" (LPFR) functions for males and females. However his data series means that he cannot distinguish changes in participation due to existing workers changing their hours from changes in participation due to workers entering or leaving the labour force. The key results are summarised below.

For male labour supply, in the long run:

- A 1% reduction in government income leads to a 0.37% increase in LFPR;
- A 1% reduction in non-wage or non-govt income leads to a 0.11% increase in LFPR;
- A 1% reduction in the real/after tax total wage leads to a 0.25% decrease in LFPR;
- A 1% reduction in the unemployment rate leads to a 0.05% increase in LFPR.

In the short run:

- The ECM was significant, but was found to adjust slowly to the long term trend (0.3). Non-salary and government income was not significant in the ECM.
- The share of the population under five was most significant in the short term (but was not significant over the long term) and was negatively correlated with LFPR.
- Next most significant in the ECM was the unemployment rate, followed by total wage income.
- Seasonal dummies were found to be significant.

For female labour supply, in the long run:

- Most significant was the proportion of females at a tertiary educational institution. However, this variable acted like a time trend which matches the rise in female participation. It may be significant only over the period tested.
- Next most significant was share of the population aged under 5, followed by the unemployment rate.

The short run ECM was significant, but the only significant variable was the unemployment rate.

Comparing the two sets of results, female participation adjusted more quickly to the long run equilibrium, reflecting a greater proportion of part-time workers and social acceptance of moving in and out of the workforce. In fact Brooks' insight here was more significant than he might have realised as later work (see the discussion on Heckman below) reveals that most of the elasticity around participation is driven by the decision to participate at all, rather than by how many hours to work. The effect of unemployment on participation was larger for females, while the effect of the wage rate was larger for males.

One of the themes in the Prebble and Rebstock (1992) compendium was the effect of (then) recent changes to the welfare system, especially those of the December 1990 policy package. This theme was revisited five years later in Maloney (1997). Maloney uses both the HLFS and the Household Economic Survey (HES) to estimate equations for the aggregate labour force participation rate, employment rate and unemployment rate. The HES data is used to estimate wages functions, the results of which are then used to construct replacement rates for individuals in the HLFS sample. Interestingly, Maloney finds that non-workers do not face lower wages than workers if other factors are held constant. As in Prebble and Rebstock, the equation is corrected for self selection bias.

The main results from the estimated labour market equations are that participation and employment are positively correlated with marriage and children, and negatively correlated with the replacement rate, being simultaneously female and married, Maori or Pacific Island. Age has a roughly quadratic effect. In a subsequent regression Maloney separates the numerator and denominator in the replacement rate and finds that the effect of the wage rate is nearly three times bigger (absolutely) than the benefit rate in the employment equation and twice as big in the participation equation. He also finds that broadening the definition of participation to include studying, raises the coefficient on benefits by 40-60%

Maloney also explores two versions of the replacement rate; one where other family income is endogenous if a given member changes labour market status, and one where it is exogenous. He finds that the estimated coefficient on the log of the replacement rate is almost identical under these two specifications, implying that the assumption of exogeneity is reasonable for cross-sectional analysis.

In a well presented paper by Winkelmann and Winkelmann (1997) a multinomial logit approach is used to ascertain whether differences in the labour force status of Maori and non-Maori can be explained by characteristics such as age, education and marital status; and whether the strength of the effect of these variables differs between Maori and non-Maori.

The authors' main finding is that education is a major factor in determining labour market status, and that its effect is stronger for Maori than and non-Maori. They caution, however, against inferring that more education for Maori will necessarily raise participation. In particular, Maori who seek higher education may be more strongly motivated in this regard (and motivation is an unobserved characteristic in the models). In addition, differences in labour market dynamics (industry, location etc) may also be biasing the measured effect of education.

The multinomial approach effectively assumes that individuals are faced with a limited and discrete choice of options. In this case there are four states; full time employment, part time employment, unemployment and not in the labour force. The approach is a useful way of managing the estimation problems that arise when the budget constraint is non-convex, as arises under participation in welfare programmes for example. However, one potential issue with multinomial models of this type is that the relative pair-wise probabilities are independent of the existence of the other possibilities. Clearly this may not be true.

1.2 Overseas Research

Da Vanzo et al (1976) use cross sectional data to estimate a variety of labour supply functions for males, although their prime interest is in analysing how the results vary with different assumptions.

One of their first points is that variables which appear to be exogenous may in fact be the result of past endogenous decisions within a life cycle theory of labour market participation. Current wage rates, non-wage income and the number of children are obvious examples. Clearly with cross sectional data there is only limited potential to study the life cycle hypothesis.

The research starts with a narrow sample of workers for whom all of the required information is available, and uses either weeks per year, hours per week or annual hours as the dependent variable. The results are mixed with many incorrect signs and insignificant coefficients. Broadening the sample to include cases where missing data can be imputed yields coefficients on wages that become less negative and even positive, with particular sensitivity to the inclusion of non-workers. The authors postulate that this might be due to such persons usually having both low hours and low wages. While this may be true, the real problem is that the equation has been mis-specified by failing to differentiate between the decision at the extensive margin on whether to work at all, and the decision at the intensive margin on how many hours (weeks) to work.

The paper does clearly illustrate the importance of including education as an explanatory variable; otherwise biased wage coefficients result. Also demonstrated is that differences in results between studies that use imputed data and those that use reported data, may in fact be caused more by who is included in the sample. The authors sum up by noting that of all the various models, in only two cases are the income and substitution elasticities.

Overall the database is rather poor with little variation in hours worked, leading to small coefficients on wages. The low variation in working hours is probably a reflection of a lack of ability to change one's hours rather than an inherent insensitivity to other factors.

Jumping forward by seven years, the book by Killingworth (1983) effectively presents a summary of labour supply research to that date. Key findings are that:

- Pecuniary variables – wage rate and property income – are important;
- Female participation is more wage/property income elastic than male participation.
- Female participation is more sensitive to the fixed costs of working because females tend to have lower wages and work fewer hours than males.
- Unobservable factors heavily influence labour supply.

While rising wage rates have attracted women into the work force over the last 100 years, they have not had the same effect on men. Of course men's participation was already higher, but it has in fact declined over the last 100 years, albeit mostly amongst the 0-24 and 65+ cohorts. The 0-24 change reflects longer participation in education. The decline amongst the 65+ group may reflect males opting to consume leisure later in life and/or male reservation wages having risen relative to market wages because of increased transfer payments.

Male labour supply curves tend to be slowly backward bending – males appear to select more leisure at high wage rates. Females select less leisure at higher wage rates, but this may be because their wages do not tend to reach the levels obtained by males (on average).

Labour saving household devices may also be enabling women to undertake more market-based work. Women are doing less 'non-market' work and men are doing more 'non-market' work than they used to, although the total amount of work done by both males and females has not changed radically. Also they are consuming the same amount of leisure. This observation underlines the interdependent nature of family labour supply decisions, as was also evident in the work by Dixon. To what extent do such decisions depend on marital status, numbers of dependents or differences in tastes and the propensity to work? The propensity to work of people who are not in the labour force may not be the same as those who are in the labour force. But is the former group 'truly' less sensitive to wage rates, or are the observed differences in sensitivity driven more by family circumstances? These are some of the questions that Killingworth proposed for future research.

Killingworth also contributes an extensive overview article (100 pages) with James Heckman (1986) writing in Volume 1 of the *Handbook of Labour Economics* on female labour supply. The article reviews various theoretical frameworks that have been applied to labour market participation, interpreting it for example as an issue of maximising individual utility in the context of family labour supply, or as a time allocation problem. Also discussed are heterogeneity in work where wage rates and hours of work are jointly determined (meaning that just observing various combinations of wage rates and hours of work does not trace out a labour supply curve), dynamic models with exogenous wages and dynamic models with endogenous wages (particularly where periods of non-participation in order to rear children, that are presumably by choice, involve a lower wage rate on return to the work force).

The latter models in particular are extremely complex. They do yield some interesting insights such as being able to explain why women in particular demographic groups who have relatively low wages are *more* likely to be in the workforce than those with higher wages. Nevertheless the models are limited by such problems as a failure to allow for periods of non-work, imperfect credit markets, the behaviour of partners and partner's earnings (in other than a rudimentary way), and the assumption of life cycle optimisation and all that entails with regard to informational requirements. Indeed some of the theoretical frameworks fit the cliché of not being able to see the wood for the trees.

The empirical studies summarised by Killingworth and Heckman produce an unsatisfactorily wide range of results for the elasticity of labour supply with respect to wage rates. Probably the main reason for this (as noted by the authors) is the failure to allow for the bias that is caused by confining the analysis only to those in employment. We will return to this most fundamental point below.

In Volume 3 of the *Handbook of Labour Economics*, Blundell and MaCurdy (1999) present another 100+ page article that thoroughly reviews alternative approaches to modelling labour supply. The emphasis is on how different theoretical constructs contain (frequently implicit) assumptions that have critical consequences for the interpretation of estimated elasticities – usually more critical than the choice of estimation technique. Issues covered include:

- static models versus multi-period models,
- whether tax and welfare systems should be modelled as piecewise-linear constraints or via approximations using differentiable functions,
- certainty versus uncertainty,
- measurement error,
- the endogeneity of non-labour income,
- learning by doing and education/human capital models,
- home production,
- collective (family) decisions on labour supply,
- discrete choices (such as in the Maloney, and Winkelmann and Winkelmann papers),
- choices at the extensive margin.

Although much of the work deals with the inadequacy of models used up to 1999, it still provides a very useful reference for academic researchers seeking to extend the frontier of research on labour supply. Surprisingly, however, the issue of choices at the extensive margin (periods of non-participation) is not formally addressed until the penultimate section. This is surprising, both in view of the problems this causes in the specification of multi-period models which are the focus of Blundell and MaCurdy and given an excellent article six years earlier by Heckman. Heckman (1993) presents another useful guide on how labour supply should be modelled. He defines four labour supply functions:

$$E(H \mid W, Y, Z) \quad (1)$$
$$E(H \mid W, Y, H>0) \quad (2)$$
$$E(H \mid W, Y) = E(H \mid W, Y, H>0) * Pr(H>0 \mid W, Y) \quad (3)$$
$$Pr(H>0 \mid W, Y) \quad (4)$$

H is the measure of labour supply, W is earnings, Y is non-labour income and Z covers all other explanatory variables. E denotes the expected value.

Equation (1) is a standard labour supply equation, which expresses hours (or days etc) worked for given values of W, Y, and Z. Equation (2) covers only those currently employed, with no allowance for other drivers. Equation (3) augments equation (2) by specifying an additional component for the probability of being employed, which is separate function in itself, as in equation (4).

The great benefit from separating the decision to work from the decision on how much to work, is that it explicitly recognises the problem of self selection bias which is not adequately controlled for in equation (1). Heckman points out that elasticities of work effort (hours, days etc) in relation to wages are driven much more by equation (4) than by equation (2). That is, given that an individual is in work, only small changes in hours are observed in response to changes in wages. As noted above, this may just reflect the inability of many workers to change their hours.

Recognition of this explains why many studies find that female labour supply decisions are more elastic than male labour supply decisions. In fact the elasticities for females tend to reflect equation (4) decisions, whereas for males the elasticities tend to reflect equation (2) decisions.

In the empirical work described below, we look primarily at equation (4), that is at what drives the probability of being employed.

Gerfin's (1996) discussion is mainly about how different functional forms can affect estimated parameters, although he also provides examples using cross-section data for Germany and Switzerland. He looks at the following functional forms. Briefly:

1. Probit, based on the Normal distribution.

2. Semi-nonparametric, with thinner tails than a t-distribution, and which contains the Probit model as a special case.

3. A single-index model, using kernel estimation for a functional form that does not invoke normality and allows local maxima.

4. A smoothed maximum score, similar to the Probit form, but again allowing local maxima and not dependent on normality.

Overall the estimated coefficients do not differ substantially between the models, with most of the differences appearing at the ends of the distribution rather than in the middle. Thus the impact in terms of the simulated participation decisions is minimal. One interesting difference was that for the German data, model (3) predicted decreasing participation of women as the number of children increases, but then increasing participation when more than two children are present. Model (1) predicted a continuous decline while model (2) produced almost no sensitivity to the number of children. However, a modified probit model, also allowing for heteroskedasticity produced the model (2) result as well.

The other results show that female participation is positively related to education and foreign permanent residency status (the latter just for Switzerland), negatively to non-labour income, negatively (or quadratically) to the number of children, and quadratically to age. All coefficients tended to be highly statistically significant, except for education in the German model.

Gerfin does not state why the Logit model is not tested, which is somewhat surprising given the general preference for this model when simulating binary choice situations; the logistic function being less restrictive than the normal density function.

In summary, distilling the above studies, plus others that we have scanned suggests that the following factors are important in determining labour force participation.

Sex	Marital status
Age	Dependent children/others
Ethnicity	Wage rate
Health	Access to government support
Unemployment rate	Level of government support
Education / qualifications	Non-labour income (excluding benefits)
Participation in education	Tastes/desire for work

As noted at the start of this section, most of studies of labour market participation are based on aggregate time series data. The time series component is useful as it enables identification of the effects of macroeconomic variables, which cannot be accomplished with cross-section data. However, because the data tends to be aggregated rather than unit record, there is limited prospect for allowing for factors such as marital status, age, dependants and so on. Even when unit records do exist, they rarely capture detailed household information, notably the labour force status of partners. In this respect the HLFS dataset has some great advantages. Indeed, the most valuable insight from previous studies is not the list of factors that affect labour supply, but how these factors interact (for instance joint family decisions on participation), how selection bias can affect estimated elasticities, how unobservable factors (for example differences in tastes and the fixed costs of work) affect the estimates, and how models should be specified to capture these more subtle forces.

2. Labour Supply Function

Following recent work, and in particular the recommendation of Heckman (1993), the decision on whether to work at all should be modelled separately from the decision on how many hours to work. The reason for this is demonstrated below.

A simple equation of the form of Heckman's equation (1) immediately runs into two significant problems.

$$H = f(W, Y, Z) \qquad (1)$$

Firstly, if $H=0$ there are no derivatives with respect to the independent variables; that is no marginal effects can be measured. Secondly, no data exists on the wages of people who are not employed. Thus the equation suffers from both specification error and measurement error.

Early researchers thought that they could circumvent these problems by confining the analysis solely to those in employment; $H>0$ and W is observable. Unfortunately this 'solution' raises the potentially more serious problem of selectivity bias. Algebraically, the error term e in equation (5) is probably not independent of the explanatory variables.

$$H = \alpha W + \beta Y + \gamma Z + e \qquad (5)$$

The reservation wage W^* is likely to be related to the same or similar variables that explain hours worked. That is:

$$W^* = \delta Y + \varepsilon X + u \qquad (6)$$

A high propensity to work implies, *ceteris paribus,* a low u and high e. Thus u and e are likely to be negatively correlated and so Y and e are positively correlated amongst workers (even though no such correlation may exist for the whole population). Hence application of 'ordinary least squares' to equation (5) will generate a biased β coefficient. A similar argument can be applied to γ if Z shares any variables with X. Solving the estimation problem requires a two-step process, the first step estimating the probability of participation (the decision at the extensive margin), as per Heckman's equation (4); and the second step using the results from the first step in the estimation of Heckman's equation (2) for simulating the decision at the intensive margin.

The literature suggests that the former decision – to participate at all – is likely to show the greater sensitivity to the usual factors that have been shown to influence participation. For this reason, and taking into account the limited resources available to the project, we have chosen to address only this aspect of the labour supply function.

A possible reduced-form model is proposed below. The equations are reduced-form in the sense that explicit derivation from a utility function is not attempted, nor necessary. Operational issues – whether the model can be estimated with the available data – are discussed subsequently. Note that labour supply is defined as being in paid employment; or being unemployed, but available and actively looking for work.

Level 1

The top level of the model defines labour supply (the decision to participate in the labour force) for an individual as a function of seven variables:

$$LS = f(DC, HK, PE, FC, WN, YO, NV, \varepsilon) \quad (7)$$

where:

LS	labour supply decision (binary)
DC	demand conditions (eg unemployment rate) to capture the prospect of finding work soon or currently being in work
HK	human capital (actual or perceived)
PE	participation in education or training
FC	family commitments
WN	actual or expected net margin of wages and salaries over the unemployment benefit and costs of being in employment
YO	income from other sources (excluding unemployment benefit) such as NZ Superannuation or dividends.
NV	non-income value of work (intellectual stimulation, social interaction etc) or not being in work

Some of these variables are unobservable and in a life-cycle context some maybe simultaneously determined with work status; participation in education for example.

Level 2

A second tier of the model expresses some of these variables as functions of other variables.

$$HK = f(ED, OC, AG, SX, ET, HL, YZ, \varepsilon) \quad (8)$$

ED	education and skills
OC	occupation
AG	age
SX	sex
ET	ethnicity
HL	health
YZ	years in New Zealand

Human capital is not an observable variable. Also while some of the arguments in equation (8) might be considered out of place, the purpose of the equation is to capture information about an individual which a potential employer might perceive (rightly or wrongly) as being relevant to a person's suitability for the work on offer.

$$FC = f(SX, MS, CH, OD) \quad (9)$$

MS	marital status
CH	dependent children
OD	other dependants

Again family commitments do not translate into an obvious observable variable.

$$WN = WS - UB - CW \quad (10)$$

WS	wages and salary that could be earned if employed (There is the usual problem here of what value to use for people not currently employed – self selection bias.)
UB	unemployment benefit and other assistance linked to labour force status
CW	costs of work.

Equation (10) defines the net financial benefit from employment as being earnings of wages and salaries (or expected earnings on the part of those unemployed), less whatever welfare assistance would be lost or abated if employment was secured, less any costs involved with employment such as travel. Ideally this calculation should be net of tax. One might conjecture that participation would be activated if WN≥WR, a reservation wage.

Operational issues

With regard to equation (7):

- The dependent variable is binary. Hence a Logit or Probit model is an appropriate specification. That is:

$$P_i = \frac{1}{1 + \exp(-X_i \beta)} \qquad (11)$$

where X denotes the explanatory variable given in equation (7) and exp is the exponential function.

- The variable for demand conditions (DC) is the same for everyone as far as time-denominated variables are concerned, and so is of no use in cross-sectional analysis. However, one could argue that component variables such as labour demand should be, say, region specific. On the other hand demand-side variables may just be proxies for individual expectations about future wages from employment, in which case they could be correlated with current wages and so lead to potentially biased estimates.

- The non-income value of work (NV) is an unobserved variable. It may be possible to find a proxy for it, or we may have to accept it as an omitted variable and be cognisant of possible bias in the results. It is probably also linked to unemployment, as discussed below.

With regard to Equations (8) and (9):

- Both human capital and family commitments are unmeasured variables. Furthermore their component series are likely to be highly correlated. One option therefore is to use Principal Components analysis to produce a single series derived from the variables shown, for each of human capital and family commitments.

With regard to Equation (10):

No data is available on the costs of work (CW) in the HLFS Income Supplement and there is no ready proxy, so again there is the potential for omitted variable bias.

As discussed above though, the key issue with this equation is the issue of defining WS for a person who is unemployed. Various researches have tried to define what an unemployed individual would earn if they were employed, using observed characteristics such as occupation or the most recent wage earned by the individual. The latter is quite impossible with the HLFS dataset. Some studies attempt to allow for factors such as age and the length of time since the last occurrence of paid employment, but the fundamental problem of selection bias is not resolved just by imputing counterfactual data.

The definition of unemployment used here is that the individual is actively searching for work and is available for work. This rules out factors such as family circumstances which might be keeping people out of the labour force. Thus there would seem to be three reasons why someone is unemployed:

1. The individual does not possess the required skills.
2. The work on offer does not meet the individual's reservation wage.
3. The work on offer is not desirable.

In the first case the individual will either have to accept lower paying work, implying that the individual is over-priced (the second case) or withdraw from the labour market.

The second and third reasons are not independent, with higher wages being demanded for less desirable work, as depicted by the curve in the diagram below. Over time we might expect the reservation wage to contract towards the origin and the definition of desirable work to be less discriminatory. In other words, as the period of unemployment lengthens, the reservation wage drops and the willingness to consider less appealing jobs increases, implying that more job offers are considered.

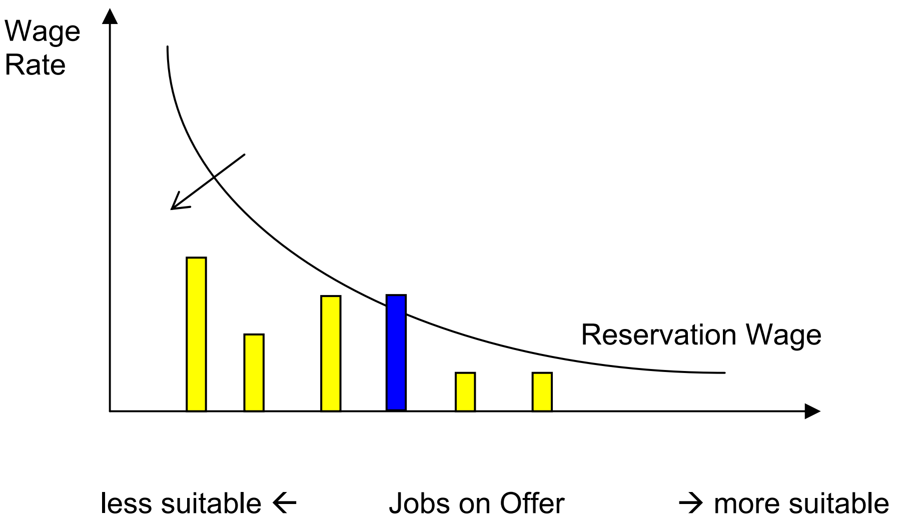

At any point in time the unemployed individual will be observing a collection of jobs on offer, as represented by the shaded bars in the above diagram. Only when the wage on offer exceeds the reservation wage curve will employment be taken up – the darker bar. It is clear then that imputing a single reservation wage, even if feasible, is conceptually unsatisfactory. For this research, in the meantime, we set aside the concept of the reservation wage.

More generally, and as discussed in the literature review, selection bias means that we cannot be confident that the various factors which influence the decision to participate operate in the same way for those in work and for those without work. Hence in comparing those in the labour force to those not in the labour force, we have to differentiate between the employed and the unemployed.

One option is to include a collection of dummy variables for both intercept and slope effects. A superior option, and one which the large size of the database allows us to utilise, is to run separate equations. That is:

1. LS = 1 for if an individual who is in the labour force and is in employment
   LS = 0 for an individual not in the labour force

2. LS = 1 for an individual who is in the labour force and is not in employment
   LS = 0 for an individual not in the labour force

Splitting the analysis in this way allows us to ascertain the extent to which the factors that drive participation differ between those employed and those unemployed.

Clearly employment and unemployment are mutually exclusive, but for multi-member households is this still the case? One should certainly admit that labour market participation decisions may be made jointly by household members. However, this is likely to relate to who will work (or attempt to work) and who will remain out of the labour force, rather than to who will work and who will be unemployed. In the next section we will endeavour to allow for joint decision making in terms of the decision to participate or not.

3. Econometric Analysis

As noted above, we had made available to us (on LMPG premises) the unit record data from the Household Labour Force Survey (Income Supplement), undertaken annually by Statistics New Zealand. Although the data spans the years 1997 to 2002, the structure of the data files is not identical over the six years. In addition the unique identifiers which enable individuals to be tracked over time have been removed. Consequently the longitudinal content of the dataset has been severely compromised. Thus each year needs to be treated as a distinct set of observations. For these reasons we have selected just one year, 1997 (to start at the beginning), as the basis for this preliminary econometric investigation.

The 1997 dataset has over 28,000 observations, about 84% of which relate to people who share a household with someone else in the dataset.

Because it is possible that people in employment have different attributes and factors that influence their labour market behaviour than people not in employment (but still in the labour force) we have separated the analysis into two groups:

1. Employment versus not in the labour force
2. Unemployment versus not in the labour force

This separation also allows us to test whether the presence of the wage rate as an explanatory variable affects the estimated coefficients on the other variables. In future research it may be worth investigating a multinomial specification for the three states (similar to that applied by Winkelmann and Winkelmann for example[3]) although this may require a full quota of multiplicative interaction effects if the explanatory variables have different effects on those employed versus those unemployed.

3.1 Employment v Not in Labour Force

Table 1 shows the results of three Logit equations as described in the section 2. The dependent variable takes a value of one if the person is employed and zero if the person is not in the labour force. Unemployed people were deleted from the sample.

While the series of equations may look like a dangerous ‘specific to general’ methodology this is not the case. Equation A is the result of a careful ‘general to specific’ approach after numerous insignificant variables have been deleted. The last two equations are presented in order to test two specific variables.

Looking firstly at equation A in Table 1, PC1 and PC2 are the first and second principal components for human capital, derived from the following five variables:

1. Age greater than or equal to 65 (yes=1, no=0)
2. Education (linearly ranked 1-4 for none, School Certificate, Sixth form or Bursary, and university degree. (This is a crude classification which can easily be improved.)
3. Vocational qualification (yes=1, no=0)
4. European (yes=1, no=0)
5. Married (yes=1, no=0)

[3] op cit

Principal components is a useful way of obtaining maximum information from a collection of variables that exhibit high correlation. The first principal component explains 27% of the variance, while the second adds another 23%. This is reasonable, but not great. We return to this point later. Sign-wise, however, PC1 seems plausible with a negative sign on the age variable and positive signs on the other variables.

The other two variables in equation A have the expected negative sign. The mean value of non-wage (non welfare benefit) income is $45/week, but $224 for those who have any income from that source. Thus while the variable is highly significant, it would take in increase of about 80% for it to have the same effect as studying.[4]

Table 1
Logit Employment Equations

	Equation A		Equation B		Equation C	
	coefficient	marginal effect*	coefficient	marginal effect*	coefficient	marginal effect*
PC1	110.1	0.27	162.3	0.36	132.4	0.29
PC2	4.59	0.012	10.7	0.024	23.5	0.051
Non-wage income	-0.015	-0.0036	-0.017	-0.0039	-0.017	-0.0037
Studying	-8.84	-0.63	-9.83	-0.77	-10.3	-0.87
HH member in employ.			-11.1	-0.77	-16.3	-0.87
Weekly wage					0.0061	0.0013
% successful	74.0		85.1		87.4	
McFadden R^2	0.380		0.612		0.686	
Notes	all sig, but PC2 t-value is 1.5		all sig.		all sig.	

*Because of the functional form of the Logit equation the coefficients do not directly show the effect of a change in an independent variable on the probability of being employed. The marginal effect denotes the effect on the dependent variable of a given change in an independent variable for an 'average' individual. For the PC variables which are all standardised, the marginal effect refers to the effect of a one percentage point change. For other continuous variables such as non-wage income in Equation A for example, the marginal effect of non-wage income at 0.0036 means that an increase in non-wage income of say $10/week reduces the probability of being employed by an average 3.6 percentage points. For the (0,1) dummy variables the marginal effect denotes the effect of changing from 0 to 1. For example in Equation A, changing between studying and not studying changes the probability by 63 percentage points.

Other variables that one expects to see in equation A (or in the principal components) such as sex, number of dependants, and various interactions with age, were not significant.

Equation A has a success rate of 74%. This measures the proportion of occasions where the model correctly predicts the individual's participation decision. If the estimated value of the dependent variable for an individual is 0.5 or more, that individual is classed as employed. Conversely for a value less than 0.5.

[4] Note that non-wage income could be the result of past labour force participation decisions, so in a dynamic model this might be problematic. However, Blundell and MaCurdy include such income in a participation equation (as here), but not in an hours worked equation. Also Maloney (op cit) found that whether other income is treated as endogenous or exogenous did not affect the estimated coefficient on the replacement rate.

In equation B a binary variable is added for observations where an individual has one or more other members of the same household in employment, and where the subject individual is not in employment. Hence the variable captures in a crude way the effect of joint decision making in a household as regards members' labour force participation status. A more refined approach would differentiate between spouses and other household members, and perhaps also between primary and secondary income earners.

Nevertheless even with this simple variable (which assumes a similar effect across all individuals in such a situation) the explanatory power of the model rises from 74% to 85% in terms of successful predictions, and the McFadden R^2 also rises substantially. The variable has a strong negative effect (t-value of 12.8) and is as powerful as the studying variable. The coefficients on non-wage income and studying are reasonably stable between the equations, but the human capital ones less so. Both PC coefficients rise, suggesting that there is some negative correlation between human capital and household participation decisions. That is, while higher human capital increases the probability of participation, the measured effect is muted if no allowance is made for situations where only one partner (household member) works, out of a partnership (household) where perhaps both partners have high human capital. This is standard omitted variable bias.

An important feature about equation B is that labour income does not appear. Of course for those not in the labour force it cannot sensibly appear, but the point is that even without wage income, the model provides a high degree of explanation. And there are still many other joint interaction variables which could be included, even apart from refining the one used here. For example while the equation allows for studying, this variable has not been filtered for partnership status, nor have human capital effects been considered.

Wages rates are partly endogenous, so there are unobservable factors that affect both wages and hours worked (including the decision to work), requiring the use of instrumental variables. This is something to be pursued when estimating equations for hours of work. In the meantime it is interesting to observe the effect of adding the wage rate to equation B. Will its presence bias the results?

Equation C shows the results. The coefficient on PC1 declines which is to be expected as higher human capital is generally associated with higher earnings.[5] These two variables should not appear in the same equation in this manner.

The coefficient on wage income is statistically significant, but small. The mean value of wage income across those who receive it is $663, so a doubling of its value produces an effect on the dependent variable of about the same amount as the studying variable and the household employment variable. The absolute rise in the coefficient on the latter is plausible, implying that the probability of one member of a household (partnership) not being in the labour force when the another member is employed, rises with the earnings of the working partner. This should be further investigated in future research. Overall though, the effect of including the wage variable suggests that the endogeneity of wages is not too serious an issue.

Allowing for a wage effect increases the success proportion by 2.3 percentage points to 87.4%. Thus wages are not a major determinant of the decision to participate in the labour force. Given previous studies and the structure of the equation this outcome is not surprising, although a net wage value which takes account of taxation, changes in benefits and the fixed costs of employment would probably be a more powerful explanatory variable.

[5] See Stroombergen et al (2002)

Table 2 repeats the results for Equation C and also shows the results of adding a set of regional dummy variables. Only six of the potential twelve variables are significant and only Northland has a significant negative sign, implying relatively high rates of non-participation in the labour force. The other coefficients are virtually unchanged from Equation C, suggesting no correlation between the regional variables and the other variables. Clearly, regional labour market differences do exist, but they are very small and do not undermine the effects of demographic, household and human capital variables on labour force participation.

Table 2
Logit Employment Equations

	Equation C		Equation D	
	coefficient	marginal effect	coefficient	marginal effect
PC1	132.4	0.29	131.7	0.29
PC2	23.5	0.051	23.2	0.051
Non-wage income	-0.017	-0.0037	-0.017	-0.0037
Studying	-10.3	-0.87	10.3	-0.85
HH member in employment	-16.3	-0.87	-16.3	-0.85
Weekly wage	0.0061	0.0013	0.0061	0.0013
Northland			-0.306	-0.043
Auckland			0.236	0.027
Wellington			0.155	0.018
Nelson, Tasman, Marl.			0.336	0.037
Canterbury			0.217	0.025
Otago			0.292	0.033
% successful	87.4		87.5	
McFadden R^2	0.686		0.687	
Notes	all sig.		all sig.	

The nesting approach outlined in the previous section is a way of reducing potential multicollinearity problems and minimising the chances of including variables (in a single equation model) which might have simultaneous positive and negative effects. For example, given that the marriage market is a selective one, being married might suggest relatively high human capital, thereby raising the probability of employment, but it might also suggest an increased likelihood of time off work to look after children, thereby reducing the probability of employment.

This approach worked reasonably well for the derivation of a human capital measure, but led to unsatisfactory results for a family commitments variable. A simpler option for the latter was to define a binary variable which takes the value of one for people under 65 (and over 15) who have dependants, and zero in all other cases. However, this variable is not significant in the above models, possibly because of correlation with the human capital principal component variables.

A different Logit specification was therefore tested, one which eschews the PC human capital measures and directly incorporates human capital variables, the family commitment variable and other individual characteristics. Both age and age squared are included. Following Blundell and MaCurdy, (who recommend the logarithms of hours worked and wages) this specification is a means by which cross-sectional data can yield useful information about wage elasticities within a lifecycle framework. That is, given the existence of life cycle effects, it is important to distinguish movements of the wage curve with movements along it. Age is a reasonable proxy for the latter. With the application of instrumental variables the individual characteristics would instrument for the wage rate.

The results are shown in Table 3. Overall they are very consistent with equations A-D. Education, European ethnicity and being married all raise the probability of employment, just as they do via the PC approach. The Education variable takes values of 1-4 corresponding to no qualifications, School Certificate, a higher secondary school qualification, or university degree respectively.

Weekly wages have same small effect as before, but non-wage income has a lower effect. The quadratic age specification implies a participation peak at age 22 which seems low, but may be plausible and is consistent with the negative sign on age in the first principle component in equations A-D. Non-New Zealanders have a lower probability of being in work, over and above the non-European effect. This variable was not significant in the PC human capital measure.

Two other new variables are female and 'age <65 with dependants,' both of which have negative effects. It would be worth testing an interaction term for these two variables, as well as with marriage. And of course the decision on whether to have children is probably not independent of the decision to work. This nexus illustrates the case for nested equations, or for perhaps splitting the sample by sex and age.

Table 3
Non-Nested Specification

	Equation E		Equation F	
	coefficient	marginal effect	coefficient	marginal effect
Female	-0.471	-0.043	-0.473	-0.044
Non-wage income	-0.012	-0.0017	-0.012	-0.0017
Age	0.033	0.0047	0.033	0.0047
Age^2	-0.0008	-0.00011	-0.0008	-0.00011
Age<65 with dependants	-0.217	-0.026	-0.218	-0.026
European	1.016	0.15	1.036	0.17
Education	0.328	0.047	0.328	0.047
Married	0.792	0.11	0.793	0.12
Nationality (NZ=0)	-0.188	-0.022	-0.215	-0.026
Weekly wage	-0.0083	0.0012	-0.0083	0.0012
Auckland			0.167	0.018
Nelson, Tasman, Marlborough			0.237	0.024
% successful	81.0		81.1	
McFadden R^2	0.548		0.549	
Notes	all sig,		all sig,	

The 'HH member in employment' variable was either not statistically significant or led to non-convergence of the Likelihood function algorithm. This suggests some instability or multicollinearity somewhere. It requires further investigation.

The list of statistically significant regional variables is much shorter than before, with only Auckland and Nelson/Tasman/Marlborough being significant. It may be that some of the other variables are correlated with regional effects – for example wage rates, as discussed in section 2.

While the nested PC-based approach needs some refinement, its results are tidier than those of the non-nested equations. The latter is prone to the effects of multicollinearity and ambiguous signs, leading to potentially misleading measures of statistical significance.

3.2 Unemployment v Not in Labour Force

We now look at unemployment versus not in the labour force, the aim being to ascertain whether the explanatory variables have different effects from those in the employment equations, or indeed if the explanatory variables are altogether different.

Analogous to Table1, Table 4 shows the results of three Logit equations, where the dependent variable takes a value of one if the person is unemployed and zero if the person is not in the labour force. Employed people were deleted from the sample.

The first four variables of the principal components analysis described above were used here except that marital status was deleted and nationality (foreign=1, NZ=0) was included. Again age>65 has a negative coefficient, but the coefficient on European switches from positive to negative. This suggests that combining those employed with those unemployed and analysing 'in the labour force' versus 'not in the labour force' would probably result in an insignificant coefficient on European, and thus rejection of a potentially useful variable. The other three coefficients, on education, vocational qualification and nationality are positive. The first principal component accounts for 29% of the variance and the second for an additional 24%.

Table 4
Logit Unemployment Equations

	Equation G		Equation H		Equation I	
	coefficient	marginal effect	coefficient	marginal effect	coefficient	marginal effect
PC1	14.1	-0.0025	14.3	-0.0025	12.8	-0.0022
Non-wage income	-0.025	-0.00044	-0.025	-0.00043	-0.025	-0.00043
Married	-0.494	-0.0088	-0.432	-0.0077	-0.434	-0.0080
Welfare benefit			0.00093	0.00002	0.00088	0.00002
Otago region					-0.391	-0.0048
Southland region					-0.348	-0.0044
% successful	0.825		0.826		0.826	
McFadden R^2	0.191		0.193		0.195	
Notes	all sig		all sig		all sig	

Equation G has a success rate of 82.5%, but a McFadden R^2 of only 19.1%, implying that there is still much to be explained regarding the decision to be in the labour force, but not actually in work. Via PC1 human capital still has a positive effect, but much smaller than before. Higher human capital is associated with labour force participation, even though that participation may not actually be in employment (yet).

Marital status is negative, implying that married people are more likely to be out of the labour force than in it and unemployed.

As with the wage rate variable before and consistent with expectations, addition of the welfare benefit to the explanatory variables has almost no effect on the model – equation H. The success rate barely moves and the other coefficients are very stable. The mean weekly benefit rate is $63 or $216 over those who receive a benefit. Thus the effect on the dependent variable of doubling its value is about half of the effect of marital status.

Equation I tests for regional effects. Two of the regional variables are significant; Otago and Southland both appearing with negative coefficients, indicating that people in these regions are more likely to be out of the labour force than in the labour force and unemployed, given that are not employed. In other words, residents of Otago and Southland are less likely to be unemployed than not in the labour force, compared to residents in other parts of New Zealand, but from these results alone we cannot say whether this is due to a high rate of employment or to a low rate of participation. Note also, that as in Equation D, the coefficients are very similar, suggesting that regional effects in these two regions are not significantly different from each other.

3.3 Further Work

As stated at the beginning of the paper, this investigation of the decision to participate in the labour market is just a preliminary analysis of how data from the HLFS Income Supplement may be used to explore New Zealand labour market issues. We have only been able to scratch the surface of this rich dataset. Amongst the more obvious analyses to undertake in future are the following:

- Test the equations over the other years.
- Test for heteroskedasticity around issues such as household level clustering, and undertake further testing of model stability.
- Revisit the equations for unemployment versus ‘not in the labour force’. The equation fit here was much worse than for employment versus ‘not in the labour force’. While the separation approach used in this paper may lead to some bias, the results obtained imply that even more bias is likely if a simple combined approach were to be adopted.
- Examine the effects of occupation and industry.
- Refine the specification of joint decisions about labour force participation within households and couples, including the joint effects of marriage, gender, dependants and education.
- Refine the specification of educational qualifications and training, and consider these more formally within a framework of investment in human capital.
- Look more closely at the age 65 – NZ Superannuation effect.
- Consider how the time dimension that exists in the successive snapshot datasets (albeit without longitudinal tracking at the unit record level) can be utilised in the models to allow for macroeconomic effects.

Eventually the analysis should be extended to analyse hours of work, using Heckman’s two step approach which the utilises the results of a Logit or similar model of the participation decision – from equations such as those explored above – to correct for self-selection bias.[6]

[6] We made a start on such an approach for this paper, but lack of resources prevented pursuing it to a stage where it could usefully be presented here.

3.4 Miscellaneous Data

Table 5 presents some summary statistics for the main variables of interest that have been used in the equations above.

Table 5
Summary Statistics of Financial Variables

	Over stated sample	For those with variable >0
Wage income over those employed and not in LF	$409	$663
Non-wage income over those employed and not in LF	$45	$224
Non-wage income over those unemployed and not in LF	$96	
DSW income over those unemployed and not in LF	$63	$216
% employed in employment equations	60.9%	
% unemployed in unemployment equations	10.9%	
% of entire sample with household overlap	84.1%	

References

Blundell, Richard & Thomas MaCurdy (1999), Labour Supply: A review of alternative approaches in O. Ashenfelter & D. Card (eds) *Handbook of Labor Economics*, Volume 3, Elsevier Science B.V., 1559-1695.

Brooks, Ray (1980), Male and Female Labour Force Participation in New Zealand 1978-1988, *Reserve Bank of New Zealand Discussion Paper* G90/3, April .

Da Vanzo, Julie, Dennis N. De Tray and David H. Greenberg (1976), The Sensitivity of Male Labour Supply Estimates to the Choice of Assumptions, *Review of Economics and Statistics*, August, 58(3), 313-325.

Dixon, Sylvia (19962), Labour force participation over the last ten years, *Labour Market Bulletin* 1996:2, 71-88.

Gerfin, Michael, (1996), Parametric and Semi-parametric Estimation of Binary Response Model of Labour Market Participation, *Journal of Applied Econometrics*, Vol 11, 321-339.

Heckman, James, (1993), What Has Been Learned About Labour Supply in the Past Twenty Years, *AEA Papers and Proceedings*, May, 116-121.

Killingworth, M (1983), *Labour Supply,* Cambridge University Press, New York, USA.

Killingworth, Mark & James Heckman (1986), Female Labour Supply: A Survey, in Orley Ashenfelter & Richard Layard (eds), *Handbook of Labor Economics*, Volume 1, North-Holland and Elsevier, 103-204.

Maloney, Tim (1997), *Benefit Reform and Labour Market Behaviour in New Zealand*, Institute of Policy Studies, Victoria University, Wellington.

Prebble, Mark & Paula Rebstock (1992), *Incentives and Labour Supply: Modelling Taxes and Benefits*, Institute of Policy Studies, Victoria University, Wellington.

Stroombergen, Adolf, Dennis Rose and Ganesh Nana (2002), *Review of the Statistical Measurement of Human Capital*, Report to Statistics New Zealand, November.

Winkelmann, Rainer & Liliana Winkelmann (1997), Determining the relative labour market status of Maori and non-Maori using a multinomial logit model, *Labour Market Bulletin* 1997:1, Department of Labour.