1 Introduction and overview

1.1 Introduction

The standard 'workhorse model' used in economics to study consumption, saving and labour supply decisions, and which also provides the basis for virtually all of public economics, is that of an individual decision-maker, who divides his time between market labour supply and leisure, and allocates the resulting income to consumption goods. There is a vast literature that uses this model to analyse these decisions, both in a static, timeless setting, and within a framework in which consumption and time allocations are chosen over an entire life cycle, with or without uncertainty.

Although this class of models has over the years yielded many valuable insights, household survey data, econometric investigation and theoretical analysis all suggest that it provides an inadequate basis for obtaining a satisfactory understanding of household decisions, and for estimating the behavioural parameters of households formed by two adults, especially if they have children. This therefore limits its usefulness in addressing many of the problems of public economic policy, for which we need both an adequate conceptual framework and robust and reliable estimates of behavioural parameters. In chapters 2, 3 and 4 of this book we expand upon this assertion, which of course may not be readily accepted by at least some of our economist colleagues. In these chapters we first summarise briefly the main results of the model, review the empirical evidence, which generally rejects its implied restrictions on household consumption demand and labour supply functions, and then undertake a comprehensive survey of the alternative models that have been developed over the last three to four decades.

It is fair to say that the theoretical development of these models is well in advance of their empirical application, the main reason for this being data limitations. In some cases, however, there is also a failure to use some very good data that are available, for example, time-use data. Although econometricians are past masters at making bricks without straw, our main contention is that a necessary condition for serious progress in the applications of these models, which are central to the formulation of public economic policy, is greatly improved gathering of data on what actually goes on inside households with respect to consumption, production and time-use decisions.

A first approach to the static analysis of the labour supply decisions of couples, the household utility function model, defines household utility as a function of total consumption and two types of leisure or labour supply. So far this has been the main basis for specifications of male and female labour supply functions that have been estimated for the purpose of tax and welfare reform analysis. We show in chapters 6, 7 and 8, where we apply the model to tax analysis, that it can be useful for theoretical purposes, given that the assumptions allowing its use are clearly understood. However, this model also has severe limitations when it is used for estimating the behavioural responses of two-parent households with dependent children. A key criticism is that, although it can be interpreted as a reduced form of a two-person household with household production, it yields restrictions on consumption demand and labour supply functions which are identical to those of the individual model, and which therefore are similarly rejected by the data. In chapter 3 we develop this point at some length, and in chapter 4 discuss some of the empirical work with this model.

The model has a further major limitation. The data show that in many households, following the arrival of children (though not before), there is a marked division of labour, with the female tending to specialise in home production of goods and services, especially child care. Thus, households are characterised by specialisation and exchange, suggesting that the two-parent household needs to be modelled as a small economy, using the concepts of general equilibrium theory and welfare economics.

The goods and services produced and consumed in the household have close (though in general not perfect) market substitutes, and what the data also show is that there is a marked heterogeneity across households in the extent to which female labour time is divided between producing these goods in the household and working in the market and buying them in. The model we develop should also allow us to explain this heterogeneity, since, as we show in chapters 5 to 9, it plays a crucial role in the analysis of taxation and income redistribution policies. Although it is not hard to develop hypotheses to explain the heterogeneity, going beyond the usual suspects of wage rates and demographic variables (which incidentally explain only a small part of it), much empirical work needs to be done before we really understand what causes it. Again, the problem here seems essentially to be data limitations.

In this introductory chapter we motivate the approach we have taken to modelling the household by presenting data on four countries – Australia, Germany, the UK and the US – that support the three main elements of our modelling approach. These are:

- The redefinition of the categories of time use from the standard two – market work and leisure, defined as time used directly for one's own consumption – to three: market work; leisure; and time spent in household production, i.e. in producing goods and services within the household for consumption by the members of that household. At many points we further refine this to distinguish between time spent on child care and on general household activities that are carried out whether or not children are present in the household.

- The recognition of the role played by wage rates and employment possibilities outside the household in determining the within-household division of labour, rather than restricting attention to individual preferences and the comparative advantage in household production of one partner relative to the other. A wage gap between partners may not only make it rational for the lower-wage partner, typically the female, to specialise in household production, but may also influence the distribution of real income or utility within the household. At the same time, across-household variation of productivities in household production, together with the price of market substitutes, is also given an important role in explaining variation in the utility possibilities of households.
- The redefinition of the life cycle in terms of the phases that the typical household goes through – what we could call the 'family life cycle' – rather than in terms of the age of the 'head of the household', as in the standard life cycle literature. The importance of this extension is suggested by time-use data that show the dramatic changes that take place, with strongly persistent effects, after the arrival of the first child.

The enormous increase in female labour force participation that took place in the developed economies between the early 1950s and the late 1980s makes the importance of these extensions to the standard models self-evident. Economic and social historians may still be debating the causes of this transformation, but its consequences are clear. The 'traditional model' of the household, in which the male head specialises in market work and the female in work within the household, now represents only around a third of families with dependent children in most OECD countries, and fewer in some cases. Moreover, as noted above, households have become highly heterogeneous in respect of the labour supply decision of the female partner. In the UK and the US, for example, roughly 30 per cent of households with dependent children continue to conform to the 'traditional model', and in roughly a further 25 per cent in the UK, and 45 per cent in the US, both partners work full-time in the market. The majority of the remaining households have one partner, not always the male, as the primary earner in full-time work, while the second works part-time. As we pointed out earlier, while some of the observed heterogeneity is associated with the age and number of children and, to a lesser extent, with economic variables such as wage rates and non-labour income, much of it remains unexplained after controlling for these variables.

This transformation in work choices has created challenges to the formulation of public policy. Most immediately, it poses the question of how to tax two-earner couples. Different countries have found different solutions, with, for example, the USA and Germany taxing incomes jointly, while many other countries, including the UK, Canada and Australia, tax them separately. The large falls in fertility associated with growth in female labour force participation have been largely responsible for the changes in the age structure of the populations of these countries, and so for the associated problem of funding pay-as-you-go social security and pension systems. In the area of family income support, withdrawal of benefits on the basis of total household income leads to very high marginal tax rates on individual incomes and reduced incentives to work.

Discussion of policy towards child support raises issues of whether this is best done by direct lump-sum payments or by the provision of child care facilities outside the home, and, in the case of the former, whether it matters to which parent the payment is actually made. The basic issues of equity and efficiency that underlie the formulation of these policies become more complex in the context of specialisation in household production and female labour supply heterogeneity. For example, it is no longer self-evident that total household income is an adequate measure of a household's standard of living.

These policy issues, together with discussion of the models we need to analyse them, are the focus of the chapters to follow. We do not address policy solutions to poverty due to long-term unemployment or disability, where specific retraining programmes combined with wage subsidies may offer more effective solutions. In our view, these problems, which the data show affect a relatively small proportion of the population of prime working age in most OECD countries, need to be considered in the context of the economic forces that drive them, and the specific moral hazard problems that may be associated with them. The central question we address in the analysis of policy in this book can be posed as follows: how should the system of income taxation be designed to redistribute income within and across the vast majority of households consisting of couples with children, where at least one partner is fully employed?

In the remainder of this chapter we present detailed empirical evidence on the labour supply, hours of domestic work and child care and the earnings of couples to support our view, first, that the household should be modelled as a small economy, and, second, that we need to define the life cycle in terms of the presence and ages of children, taking account of the demands they create – demands that can be met either by work at home or by the market.

1.2 Labour supply and household production

We draw on data for four comparable OECD countries: Australia, Germany, the United Kingdom (UK) and the United States (US). We use data from the following household and time-use surveys for these countries:

- Australia: The Household, Income and Labour Dynamics in Australia Survey, Wave 5, 2005 (HILDA).[1] Australian Bureau of Statistics (ABS) 2006 Time Use Survey (AU TUS).
- Germany: German Socio-Economic Panel, Wave 22, 2005 (GSOEP).[2]
- UK: Expenditure and Food Survey, 2005–6, National Statistics (EFS). Time Use Survey 2000, National Statistics, United Kingdom (UK TUS).
- US: Panel Study of Income Dynamics, 2005 Public Release (PSID).[3] American Time Use Survey 2005 (ATUS).[4]

[1] See Goode and Watson (2007). [2] See DIW Berlin website: www.diw.de/gsoep/.
[3] See PSID website: http://psidonline.isr.umich.edu/.
[4] See US Bureau of Labor Statistics and US Census Bureau website www.bls.gov/tus/.

Table 1.1 *All individuals aged 25–59 years: employment rates by gender*

	AUSTRALIA		GERMANY		UK		US	
	M	F	M	F	M	F	M	F
% Employed	88.1	72.1	84.2	67.5	84.3	70.7	88.4	72.8
% Full-time	80.2	38.4	81.1	44.4	77.8	37.9	84.0	56.0

When male and female employment rates for countries such as Australia, Germany, the UK and the US are compared over time, we observe a large measure of convergence since the 1950s, especially in the latter two countries, due not only to growing female employment, but also to declining male employment. Table 1.1 shows for each of the countries the percentage of males and females of prime working age who are employed, based on data for all individuals aged 25 to 59 in HILDA, GSOEP, EFS and the ATUS.[5] In each country the gap between male and female employment rates is less than 17 percentage points. The male rate ranges from 84.2 per cent in Germany to 88.4 per cent in the US, and the female rate from 67.5 per cent in Germany to 72.8 per cent in the US.

Broad comparisons of employment rates of this kind are sometimes assumed to show that the labour supplies of males and female are converging in the same way as employment rates. This is a mistake. Employment and participation rates can give a misleading picture of the true relation between the characteristics of labour supply of men and women, and this is in fact the case for these four countries. This is because there is a large gap between rates of full-time employment of men and women, as shown by the figures in the second row of table 1.1.[6]

When we select the data for couples, we find an even larger gap between the male and female full-time rate. This is because most singles work full-time, irrespective of gender, and therefore excluding them has this effect.

In the remainder of this chapter we focus on couples.[7] Table 1.2 presents couples' employment rates and average hours of work by gender.[8] In the US, Germany and Australia over 85 per cent of married men are employed full-time, and in the UK 81.7 per cent. Apart from the US, only 34–36 per cent of married women are employed full-time. As a result, there is a large gender difference in average hours of market work, even in the US where the female full-time rate is 50.4 per cent.

[5] The numbers of male and female records in this age category drawn from each of the data files are as follows: HILDA: 3,652 males and 4,064 females; GSOEP: 4,478 males and 4,928 females; EFS: 3,532 males and 3,942 females; and ATUS: 3,743 males and 4,841 females.

[6] The full-time employment rate is computed as the percentage of records showing 35 hours of work or more per week.

[7] The samples include couples in single-family households with no others except children less than 15, dependent students or non-dependent children older than 15.

[8] The numbers of male and female records drawn from the couples samples are: HILDA, 2,084 male and 2,140 female; GSOEP, 3,544 male and 3,910 female; EFS, 2,675 male and 2,808 female; and ATUS, 2,604 male and 3,023 female.

Table 1.2 *Couples: employment rates and labour supplies by gender*

	AUSTRALIA		GERMANY		UK		US	
	M	F	M	F	M	F	M	F
% Employed	91.8	72.8	87.5	64.7	87.9	72.8	91.5	70.0
% Full-time	85.3	34.2	85.3	36.0	81.7	36.1	87.9	50.4
Hours per week	42.2	22.4	46.6	23.7	35.4	21.0	43.2	25.6

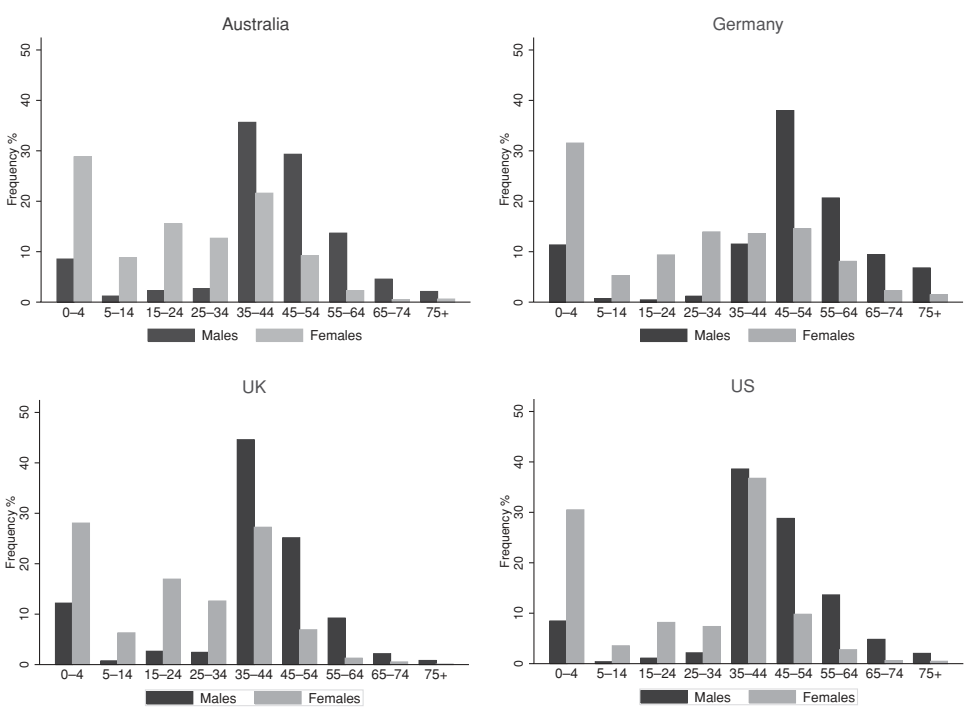

Figure 1.1 All couples: labour supply by gender

The employment rates in table 1.2 also show the very high degree of heterogeneity in female labour supply. While male labour supply shows very little variation, with almost all men working full-time, females are distributed more evenly between zero hours and full-time work.

We illustrate this heterogeneity graphically in figure 1.1 for all couples. The figure presents histograms of 'usual weekly hours of work' for partners aged 25 to 59 years, with the first band representing 0 to 4 hours, and subsequent bands increments of 10 hours. The vast majority of men in all four countries work between 35 and 54 hours a week, with a second much smaller modal frequency at 0–4 hours. The distribution for women is tri-modal in Australia and the UK, with roughly equal frequencies at 0–4 and 35–44 hours per week, and a small peak at 15–24 hours per week. In the US, female

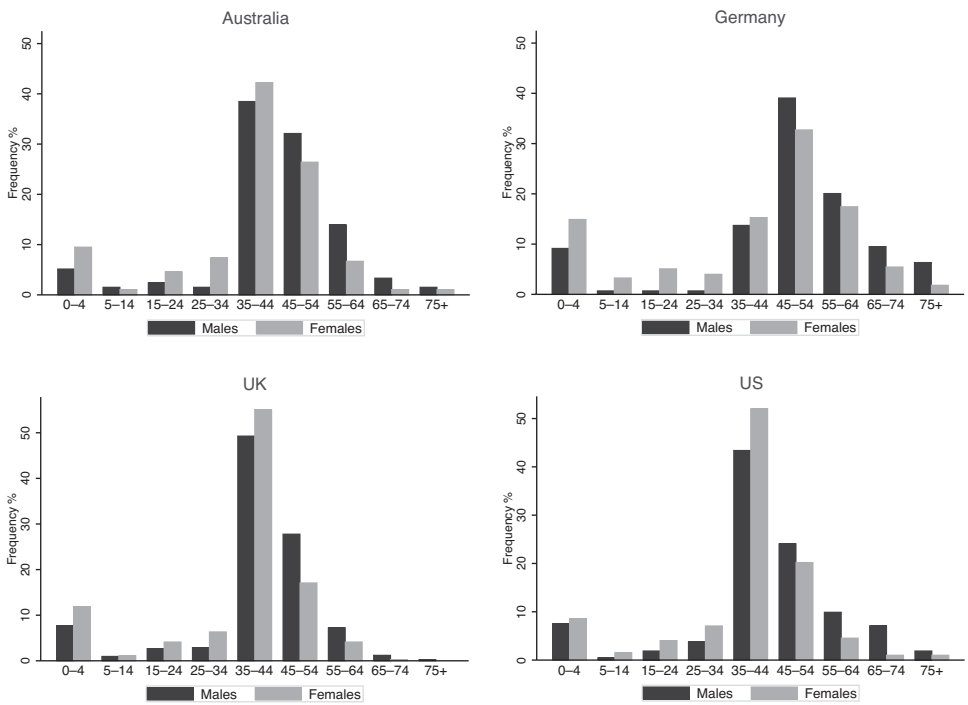

Figure 1.2 Couples: pre-children

labour supply is more nearly bimodal, although there is a significant proportion of married women working part-time. Germany has the highest percentage of prime-aged married women in the 0–4 hours band, and those who are employed tend to be spread evenly across a very wide band of part-time and full-time hours.

Much of the observed heterogeneity in female labour supply is strongly associated with children, as we would expect. However, exactly how and to what extent children play a role, and precisely what that role is, is not well explained by the existing empirical literature. After controlling for demographics, as well as for wage rates and non-labour incomes, female labour supply heterogeneity remains high, suggesting that there is a great deal of room for additional explanatory variables.

We find it useful to explore these points further by taking a very broad-brush life cycle approach. We classify households into three phases:[9] a pre-children phase; a phase in which children aged under 18 years are present in the household; and a post-children phase, in which there are no longer children under 18 years present. The usual weekly hours of work of partners in these phases are depicted graphically across the four countries in figures 1.2 to 1.4.[10]

[9] In chapter 5 we refine this classification considerably. We also deal there with two obvious questions about this approach, the first stemming from the endogeneity of the decision to have children, the second from the possibility of cohort effects (since these data are taken from cross-sections).

[10] The numbers of male and female records in phases 1 to 3 are as follows:

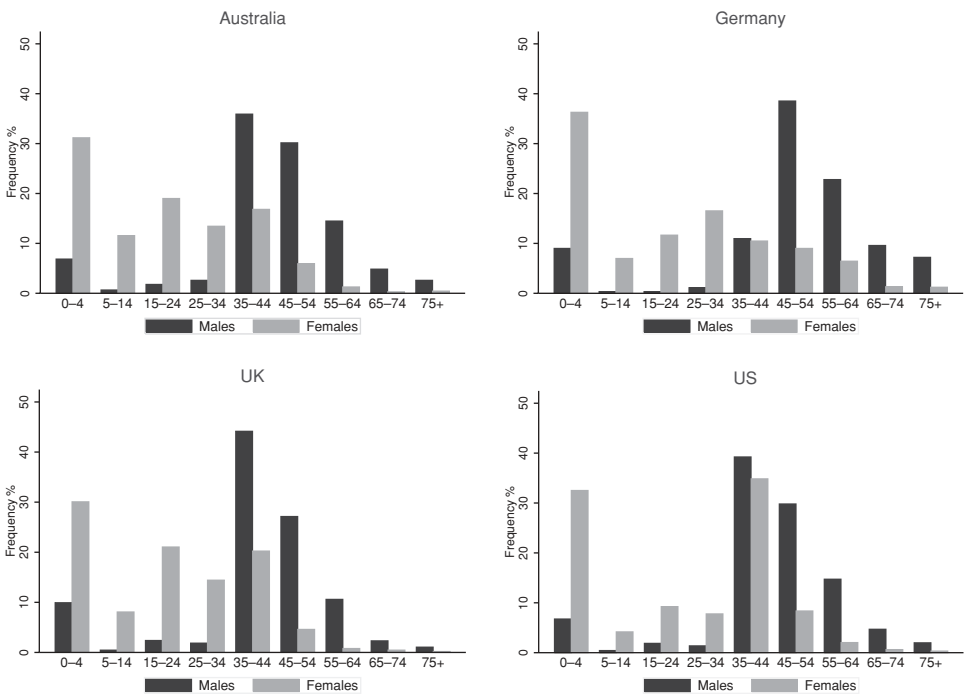

Figure 1.3 Couples: children 0–17 years

In the first phase, shown in figure 1.2, the profiles closely match – partners of prime working age tend to work full-time, and for the same hours, in all four countries.[11] In the with-children phase, the proportion of men working full-time remains about the same, while that of women falls dramatically. At the same time, the heterogeneity in female labour supply emerges. Around a third of women still work full-time, and over a third work part-time. The association between children and female labour supply is also highly persistent.[12] In the post-children phase female labour supply remains well below its pre-children level. In fact, there is little change in most of the countries, and much of what there is would seem to involve a shift from part-time to full-time work.

- HILDA: phase 1: 330 male and 284 female; phase 2: 1,339 male and 1,331 female; phase 3: 415 male and 525 female.
- GSOEP: phase 1: 266 male and 274 female; phase 2: 2,021 male and 2,023 female; phase 3: 1,257 male and 1,613 female.
- EFS: phase 1: 414 male and 363 female; phase 2: 1,499 male and 1,496 female; phase 3: 762 male and 949 female.
- ATUS: phase 1: 201 male and 202 female; phase 2: 1,840 male and 2,099 female; phase 3: 563 male and 722 female.

[11] There are relatively few records in this phase because there are relatively few young married couples without children. However, when we include singles who have not yet had children, and who are therefore essentially in the same life cycle phase, we obtain similar results from much larger samples for all four countries. Almost all men and all women not in higher education work full-time prior to having children.

[12] This is consistent with the results of panel data studies for the US. See, for example, Shaw (1989, 1994).

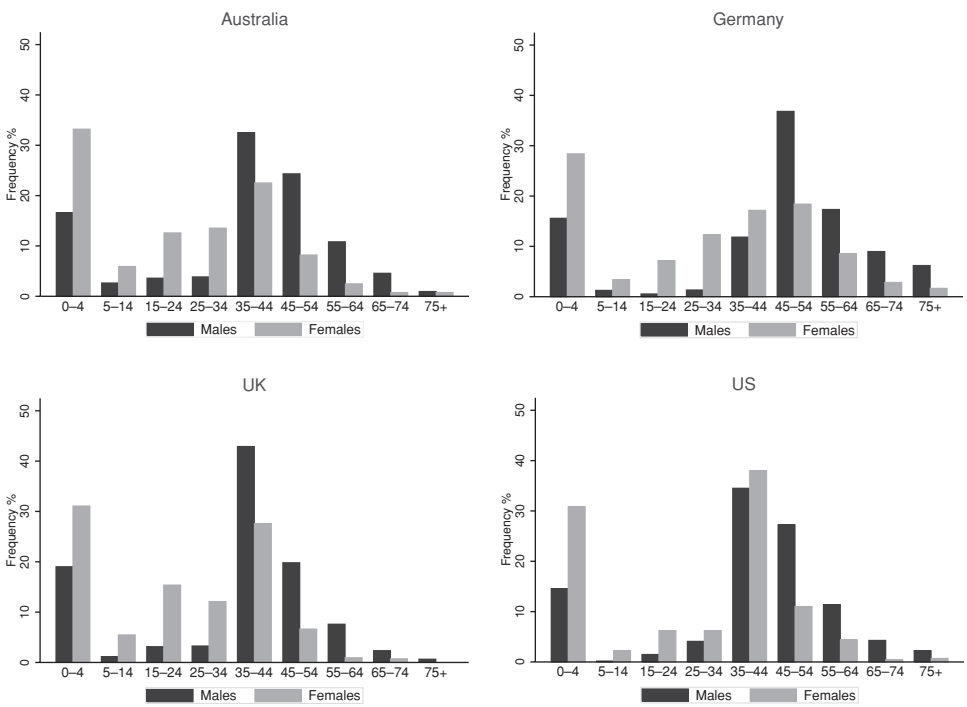

Figure 1.4 Couples: post-children

While the dramatic change in the profile of female hours from phase 1 to phase 2 indicates the strong association between the decision to have children and the labour supply decisions of married women, we would argue that the relationship is far more subtle and complex than the simple view of unidirectional causation that says: children ('demographics') cause the reduction in female labour supply. Given the decision to have children, the observed changes in female labour supply are driven by the economics of investment in the care and education of children in phase 2, much of which is directly influenced by government policy. The effects of the change in female labour supply on current and future income and employment possibilities will, however, then feed back on the household's decision on the timing and number of births. Thus there is a relationship of simultaneity between the fertility and female labour supply decisions, while both are jointly conditioned by underlying factors such as the cost, quality and availability of market child care, on the one hand, and the opportunity costs of parental time, on the other. These latter depend on: net of tax wage rates; productivity in household production; the flexibility of working arrangements in relation to the time demands of child care; the rate of depreciation of work-related human capital when out of the labour force;[13] and the probability of future re-employment, all of which are or can be strongly influenced by public policy.

[13] An extensive literature on work-related human capital accumulation includes the contributions of Eckstein and Wolpin (1989), Altug and Miller (1998) and, more recently, Imai and Keane (2004) and Olivetti (2006), among others.

The argument is straightforward. In phase 1 of the life cycle there is a low demand for home-produced goods and services because there are few of the kinds of goods and services couples in this phase consume for which there are not good, affordable market substitutes, and so there is a low demand for domestic labour in this phase. Put simply, there is nothing much to do in the home, and so it would make no sense for either partner to specialise in household production, or for singles who have not yet had children to do so. This explains why almost all males and females who have not yet had children, whether single or married, work full-time and have close to the same average weekly hours.

The arrival of children creates a very large demand for their care and for investment in their education. Governments of the four countries we are highlighting have taken over much of the role of investing in the education of children once they reach school age. However, all four countries tend to have child care and education sectors for those under school age that are expensive, poorly developed, and frequently difficult to access.

Child care can be provided by some combination of parental time and services bought in from the market. The opportunity cost of parental child care is determined by the present value of the current and future market income forgone, which depends on the factors listed above, and may well differ between the parents because of differences in wage rates and career types. The more costly and difficult it is to access market child care, the more of it will be provided at home, other things being equal. The demand for child care then implies a large demand for household production and introduces a fundamental change in the work choices of couples, which will reflect the relative costs of each partner's time.

This can of course reinforce labour market discrimination, since a lower wage for women in the expectation that they will leave the labour force to look after the children is self-fulfilling. The loss of human capital and career possibilities resulting from leaving full-time work offers an explanation for the strong persistence of female labour supply decisions made in phase 2 into the post-children phase.[14]

Time-use data provide evidence of close substitution between market labour supply and household production, consisting mostly of child care. Table 1.3 reports data means for average weekly hours allocated to these activities by couples in each of the phases. The time input to household production in phases 1 and 3 is computed as time spent on domestic work (washing, cooking, cleaning, etc.) and shopping. In phase 2, the time input to household production is split into domestic work and child care.[15]

In the pre-children phase, both partners typically work full-time in the market, and close to the same average hours. Even in Germany, which has the largest gender gap

[14] See also Attanasio *et al.* (2003).

[15] Market hours are computed from the data on 'usual weekly hours of work' in HILDA, the GSOEP, EFS and PSID. Domestic and child care hours for Germany are based on the time-use data available in Wave 5 of the GSOEP. For Australia, the UK and US we merge time-use data on domestic and child care hours in the AU TUS, UK TUS and ATUS with the HILDA, EFS and PSID samples, respectively, using regression models with the dependent variable of the domestic hours equation specified as the ratio of domestic hours to leisure hours, and of child care hours as the ratio of child care hours to domestic hours. For further detail, see chapter 5.