

COPYRIGHT NOTICE:

John F. Ermisch: An Economic Analysis of the Family

is published by Princeton University Press and copyrighted, © 2003, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

Introduction

THE FAMILY has been undergoing dramatic changes during recent years. In richer countries, marriage and childbearing are occurring much later in people's lives, and they are having fewer children. There is more childbearing outside marriage, more divorce and more one parent families. In some poorer countries, fertility has fallen sharply, while in others there has been little change. Associated with these developments, there have been changes in the ways in which family members interact with one another, including support for elderly parents or children (e.g. payments after divorce), and with markets.

The analysis in this book aims to improve our understanding of how families and markets interact, why important aspects of families have been changing in recent decades and how public policy affects them. It is built on the idea that the standard analytical methods of microeconomics, including the techniques of constrained optimization, can help us to understand resource allocation and the distribution of welfare within the family, intergenerational transfers and transmission, family formation and dissolution and household formation. It also aims to show how economic theories of the family can help to guide and structure empirical analyses of demographic and related phenomena (e.g. labour supply and child support).

The book is intended for research students, social scientists and policy makers who wish to learn how economists analyse family issues. The analysis is relevant to family behaviour in rich and poor countries. Examples of studies that apply the theory are provided throughout the book. This chapter outlines the main arguments of the book.

1.1 INTRA-HOUSEHOLD ALLOCATION

Analysis of the impact of many public policies and technological developments on the welfare of *individuals* requires that we take seriously the view that individualism is the foundation of microeconomic theory. The family is an important institution in the determination of an individual's welfare, and so we must try to understand behaviour within the family in order to assess the welfare consequences of policies and social developments. The

analysis allows for individuals within a family to have different preferences.

When putting a social institution like the family under analytical scrutiny, it is helpful to assume that individuals understand their environment and act rationally to maximize their own welfare. This does not mean that people are perfect in these respects, but to focus on analysis of the institution, we abstract from idiosyncratic aspects of individual behaviour. A fruitful starting point is to assume that people act to maximize their welfare as they evaluate it, given the predicted behaviour of others. It provides a foundation for modelling cooperative behaviour within a family. Family members must obtain welfare from cooperation that is at least as high as they would achieve from this non-cooperative outcome.

Chapter 2 focuses on the behaviour of couples with children. Decisions about when to have children and how many to have are considered later. Benefits from expenditure on children are assumed to be a “public good” for the parents, in the sense that an individual parent’s welfare from total expenditures on children is not affected by the presence of the other parent. Suppose initially that the parents do not cooperate in making decisions in the sense that each parent chooses his(her) contribution to child expenditures to maximize his(her) welfare, taking the contribution of their partner as given. There are two types of outcome from this behaviour. When one parent’s share of total income is not “sufficiently different” from the other’s, both contribute to child expenditures, and only joint family income matters for expenditures on children and each parent’s expenditure on him(her)self. How much is “sufficiently different” depends on each parent’s preferences. If, however, one parent has a relatively small share of family income, then that parent will not contribute to child expenditures. In contrast to the first type of outcome, redistribution of income between parents affects expenditures on children, private expenditure and individual welfare.

The non-cooperative outcome is inefficient (i.e. one parent could be made better off while not making the other worse off), because it encourages “free riding” on the other parent. The best strategy for one parent is to reduce his(her) contribution to expenditure on children when the other parent increases hers(his), and this usually produces too little expenditure on children relative to the efficient level. This non-cooperative model can indicate what the “fallback position” would be if communication and bargaining within the family break down, and how individual preferences and incomes affect this fallback position.

Cooperation between parents is usually a better representation of family behaviour. It achieves an efficient allocation between parents’ private consumption and child expenditure. For the types of individual

preferences usually assumed in economic analysis, the outcome is equivalent to giving each parent a share of joint family income and letting each choose his(her) consumption and his(her) contribution to child expenditure according to his(her) own preferences. In other words, it is like there is an *income sharing rule*, which in general depends on individual incomes and prices and possibly other factors such as marriage market conditions and divorce laws. One interpretation of it is that it reflects bargaining within the family.

Each parent has the alternative of not cooperating, providing an alternative level of welfare, which is called their *threat point*. Corresponding to these threat points are minimum and maximum shares of income allocated to the mother in the cooperative outcome. Individual incomes can affect the cooperative outcome by affecting these threat points. One possible bargaining rule is to maximize the welfare of a “dominant partner”. For example, if the husband were dominant, he would offer his wife just enough to accept this arrangement, which would be her threat point. Another rule is so-called “Nash bargaining”, which maximizes the product of the parents’ *gains from cooperation* (i.e. welfare in the cooperative outcome minus the threat point).

There are two prime candidates for the threat points: welfare if the parents divorce and welfare from a non-cooperative marriage, considered above. It is shown in Chapter 2 that divorce is often not a credible threat, even when welfare in the divorced state exceeds that from a non-cooperative marriage for both partners. This is because bargaining based on the threat points from a non-cooperative marriage produces a better welfare outcome for both partners than divorce. In this case, small changes in the welfare if divorced, say because of changes in welfare benefits to divorced mothers, have no impact on the cooperative outcome from bargaining. There are, however, situations when divorce is a credible threat, but then the outcome is not the one produced by Nash bargaining with the welfare if divorced as the threat points, but rather one partner is indifferent between divorce and marriage. In this case, the opportunities available to each parent if the relationship dissolved would affect allocation and distribution when the couple are together.

Bargaining within the family makes it possible that, for example, an increase in the mother’s income has two effects. It increases family income, which increases expenditure on children and herself. It also may increase the bargaining power of the mother, which could reinforce or offset the income effects, depending on each parent’s preferences for child expenditure. If, as many believe, mothers’ preferences put more weight on children than fathers’ preferences do, then an increase in her bargaining power would also increase expenditure on children. But note that if the threat points are determined by the outcome of a

non-cooperative marriage in which both parents contribute to expenditure on children, then an increase in the mother's income would not affect her threat point or her bargaining power.

Traditional consumer theory usually assumes that the household behaves "as if" it is a single agent, allowing an application of the tools of consumer theory at the household level. This assumption, which is often called the "unitary model", or "consensus model", amounts to assuming that the income sharing rule does not vary with individual incomes. Thus, one important implication of it is that expenditure on children and each parent's private consumption depend only on total family income—the so-called "income pooling" hypothesis. Suppose, for example, that we were comparing two possible cash transfer policies, one which paid the transfer to the mother and the other which paid it to the father. Under the unitary model, expenditure patterns would be invariant to the policy chosen. When the sharing rule is affected by individual incomes, expenditures on children and private expenditure would depend on who received the transfer. An important real-world example of such a policy change in the United Kingdom during the late 1970s soundly rejects the unitary model, as do many other studies. This suggests that children do better when mothers control more of the family resources, that developments which improve women's earning opportunities affect the distribution of welfare within families and that it is possible to target policies on individuals within families.

1.2 ALTRUISM IN THE FAMILY

In economic analysis, a person is said to be *altruistic* toward someone if his(her) welfare depends on the welfare of that person. Altruism, or "caring", is usually defined such that the altruist's welfare depends on the "private utilities" of the altruist and his(her) beneficiary, each of which represents their "private preferences" defined over the person's private consumption and consumption of public goods, such as child expenditures. That is, the altruist's welfare does not depend on how the beneficiary's welfare is obtained. Chapter 3 focuses on the implications of a family decision making rule that maximizes an *effective* altruist's welfare. An altruist is *effective* if he(he) makes financial transfers to his(her) beneficiary, and this happens when he(he) is sufficiently rich relative to his(her) beneficiary.

Maximizing the welfare of an *effective* altruist has some important implications. First, redistribution of income between the altruist and his(her) beneficiary has no effect on outcomes, provided that he(he) remains an effective altruist. Shifting income from him(her) to his(her)

beneficiary would produce an offsetting reduction in transfers to him(her). This means that the income sharing rule is independent of individual incomes.

Altruistic behaviour also provides partial insurance. Suppose that the beneficiary lost his(her) job, causing a fall in his(her) income. Both parties would suffer a decline in welfare, but part of his(her) welfare loss would be offset by higher transfers from the altruist. For the same reason, it also partially insulates them from targeted changes in taxes and benefits.

If private preferences take a particular form, effective altruism also has powerful effects on incentives. A selfish beneficiary has the incentive to choose the efficient level of a public good consumed by both, or alternatively is perfectly content to let the altruist choose it, even though their preferences differ. More generally, both parties would take actions that raise their *joint* income and avoid actions that lower it. This is what Gary Becker has called the “Rotten Kid Theorem”. But the existence of altruistic preferences per se does not eliminate conflict and generate efficient outcomes. When individual incomes are similar, an *effective* altruist may not emerge.

Unfortunately, it is not difficult to find preferences for which even effective altruism does not automatically align the interest of the beneficiary with those of the altruist, in contrast to what the Rotten Kid Theorem would suggest. Except for a very special case of altruistic preferences, a *necessary* condition for such conflict to be avoided is that private preferences take a particular form analysed in Chapter 3.

The Rotten Kid Theorem suggests that parents should delay transfers to their children until late in their lifetime or indeed until after their death, because this provides children with a long-run incentive to consider the interests of the entire family and maximize joint family income. Thus, it suggests that altruistic parents should use bequests rather than gifts. But if beneficiaries suffer disutility of effort in earning their income, bequests discourage effort by the child, because effort is costly and parents compensate for exerting less effort through larger bequests. Because bequests are inefficient, the Rotten Kid Theorem does not hold. Gifts (pre-committed fixed transfers) are preferable in the sense that they are efficient for the family because they only have an income effect. The child is worse off than if he(she) received bequests, because he(she) could work less and obtain higher transfers with bequests, but the parents are better off.

This is an example of the general phenomenon called the “Samaritan’s dilemma.” It arises when a benefactor’s generosity encourages beneficiaries to be less self-sufficient. In the context of saving decisions, parents’ bequests to their child encourage him(her) to over-consume early in life in order to be more impoverished and receive larger bequests later. He(she)

does, however, have an incentive to maximize joint family income with bequests. If the parents pre-commit to gifts, he(she) allocates his(her) lifetime income efficiently between periods, but once he(she) has received the gift, he(she) would wish to take actions that maximize his(her) own income, even if it reduced joint family income. Thus, bequests would produce an inefficient outcome because of the Samaritan's dilemma, while gifts would be inefficient because of the failure of the Rotten Kid Theorem. Even effective altruism fails to produce efficient outcomes in this situation, irrespective of whether transfers are given "early" or "late".

1.3 HOME PRODUCTION AND INVESTMENT

Many goods important to the family, such as investment in children, are "produced" by the family themselves through the combination of parents' time and purchased goods and services. To take a trivial example, the production of meals and the nutrition of family members require someone's time and food purchased on the market. The division of parents' labour and the implications for the costs of home-produced goods, such as the child's human capital, depend on these "home production" relationships. In many respects this is a straightforward application of production and cost analysis from the theory of the firm, but it is helpful to put it in the family context.

Parents can make investments in their own human capital that improve their earning power or their productivity in home production of goods such as their children's human capital. There is substantial evidence, for instance, that more experience in paid employment increases a person's wage. Suppose that this takes the form of simple learning-by-doing. The more a parent works in paid employment, the higher the wage and therefore the cost of his(her) time in home production. If learning-by-doing is sufficiently strong, it leads to complete specialization in market production by one parent. Small differences in wages could tip the balance in favour of complete specialization in paid work by one parent, even though both have exactly the same ability and each parent's time is equally productive in home production. If, for example, sex discrimination in the labour market makes men's wages higher than women's for a given level of human capital, then specialized human capital investments could result in women doing all of the home production and men specializing in market production.

Similar results emerge if home productivity increases with time spent in home production. If, for example, the woman's role in childbearing gives her a comparative advantage in home production, this could tip the

balance toward her contributing all of the home production time when there is learning-by-doing in home production.

Household production theory is primarily used in the book to represent family investments in children's human capital that require parents' time and goods. But it also aids in estimating and interpreting relationships that summarize more specific family activities, such as those that promote the health of family members and those that control fertility. Chapter 4 considers the production of "healthy infants", as indicated by birth weight, the health of family members through nutritional intake, and the "reproduction function", which relates contraceptive use, natural fecundity and luck to the number of births.

1.4 INVESTMENTS IN AND FINANCIAL TRANSFERS TO CHILDREN

Generations are linked by parents' gifts and bequests to their children and by investment in their children's human capital, which affect their earnings and income when they become adults. Whether and how these two types of intergenerational transfer depend on parents' resources and other aspects of family background such as parents' education are studied in Chapter 5. It is assumed that parents care about their children's incomes as adults.

Suppose first that parents have only one child. If parents are rich enough to make financial transfers to their child, then investment in the child's human capital (e.g. his educational level) does not depend on parents' incomes. Parents invest in their child's human capital up to the point that its marginal return equals its marginal cost. Thus, parents make an efficient investment in their child's human capital and then make financial transfers to their child according to their incomes and preferences.

If parents are too poor to make transfers, in the sense that the marginal utility of their own consumption exceeds the marginal utility of transfers, then parents invest less than the efficient amount, and human capital investment depends on their incomes. This suggests three separate effects of, for example, a mother's education on the education of her child. First, there is an *income effect*, which is positive because more educated women earn more. Second, there is a *bargaining effect*, which is positive if mothers' preferences put more weight on the child's income than fathers' and higher education and income increases her bargaining power. Thirdly, there is a *substitution effect*, which depends on any impact of mother's education on the cost of human capital investment in children.

Investments in children may be riskier than many financial investments that the parents could make. Parents' preferences and incomes would

then also matter for the human capital investment among parents who are rich enough to make financial transfers to their child, because of variation in parents' risk aversion with the child's income as an adult. For instance, parents would increase the level of (risky) human capital investment when their own income is larger if they are less averse to risk when their child's future income is larger.

With more than one child, how might parents treat their different children if they are equally concerned about each? Suppose parents' preferences are defined over the total income of each child as an adult. If the parents are sufficiently wealthy to make financial transfers to each of their children, then parents invest in the human capital of each child up to the point that the marginal return equals its marginal cost. Parents invest more in the human capital of a more able child, who then ends up with higher earnings, but a less able child is *fully* compensated by higher monetary transfers in the form of gifts and bequests. Financial transfers would generally differ substantially among children in the family.

Alternatively, parents' preferences may weight children's earnings differently from income derived from parents' gifts and bequests. In this case, financial transfers are the same for each child, and parents' human capital investments may *reinforce* or *compensate* for differences in children's "earnings endowments", depending upon their aversion to inequality between children's earnings. With reinforcement, the ratio of human capital investment (e.g. education levels) between children is larger than the ratio of endowments, and with compensation the opposite is the case. If parents are extremely averse to inequality between children, they invest in the human capital of their children so as to eliminate differences in earnings between children; that is, only equity considerations matter. When there is no inequality aversion, only efficiency matters. In general, both equity and efficiency considerations play a role in parents' human capital investment decisions. Chapter 5 considers the implications of these two models for estimating the returns from education, how we might choose between the two and how we can estimate whether parents compensate for or reinforce differences between their children in innate endowments.

1.5 ECONOMIC THEORIES OF FERTILITY

An important idea in the modern theory of fertility is that the psychic satisfaction parents receive from their children is likely to depend on the amount that parents spend on them as well as the number of children that they have. Gary Becker calls children who have more spent on them

“higher quality” children, the basic idea being that if parents voluntarily spend more on a child, it is because they obtain additional satisfaction from the additional expenditure. It is this additional satisfaction that is called “higher quality”. “Child quality” is now usually identified with the lifetime well-being of the child, which can be increased by investing more in the child’s human capital or by the direct transfer of wealth to the child. An increase in parents’ income may increase the amount spent on children substantially, but this would mainly take the form of higher quality rather than more children. In other words, the income elasticity of the number of children (“quantity”) is probably small compared to the income elasticity of child quality.

It is often assumed that parents view child quantity and quality as substitutes and that they treat all their children equally, in the sense that child quality is the same for each of their children. In this case, their budget constraint contains the product between the number of children and child quality, which implies that the cost (or “shadow price”) of an additional child is proportional to the level of child quality, and the cost (“shadow price”) of raising child quality is proportional to the number of children the parents have. As a consequence, there is an important interaction between family size and child quality. Suppose, for example, that there is a decline in the cost of averting births, say because of the introduction of the oral contraceptive pill, which increases the net marginal cost of a birth without affecting the marginal cost of child quality, thereby reducing family size. This lowers the shadow price of child quality, which in turn raises child quality, which raises the shadow price of children, which lowers family size further, and so on. Lower contraception costs can, therefore, produce *large* increases in child quality and further large declines in fertility. Family size can be highly responsive to changes in prices and incomes, even though children have no close substitutes.

A higher return to human capital increases desired child quality, and through a similar cumulative process reduces fertility and raises human capital investment. Thus, the increases in the returns to human capital investment associated with technical change lead to simultaneous reductions in fertility and increases in human capital investment in children, thereby accounting for important stylized facts of economic development.

Now suppose that there is an increase in parents’ income. If the quality income elasticity exceeds the one for quantity, then the ratio of quality to the number of children rises, thereby increasing the shadow price of an additional child relative to the shadow price of child quality. The substitution effect induced by this increase may be sufficiently large to produce a decline in fertility when income increases. It may, therefore, appear that the income elasticity of fertility is negative, even though children are

“normal goods”, in the sense that parents want more of them when parental income increases.

The factors affecting the cost of children are closely associated with the key role of parental time in the rearing of and investment in children. Parental time in the production of child quality is primarily the mother’s time, and the rearing of children is assumed to be *time intensive* relative to other home production activities. Thus, the cost of children relative to the cost of the parents’ living standard is directly related to the mother’s cost of time. If she has ever been in paid employment, her cost of time is the wage she could earn in employment (i.e. her foregone earnings). The higher her wage, the higher the cost of an additional child and of additional quality per child relative to the cost of improving the parents’ living standard. The relative cost of children also depends on the father’s wage, but probably weakly.

Thus, there are two channels through which men’s and women’s wages affect fertility and child quality. Higher wages for either parent means higher family income, encouraging parents to have more children and to invest more in the human capital of each child or to make larger monetary transfers to them (i.e. higher quality). Higher women’s wages also raise the opportunity cost of a child. If the opportunity cost effect on family size (child quality) dominates the income effect of women’s wages, higher women’s wages reduce family size (child quality). Higher men’s wages mainly affect childbearing through their effect on the couple’s income.

The possibility of purchasing child care, an imperfect substitute for the mother’s time in child rearing, weakens the link between a woman’s wage and the cost of an additional child. Mothers with high wages tend to purchase a much larger proportion of child care time. For them, higher wages have little effect on the cost of children, making it more likely that they increase fertility by raising family income. Similarly, in countries with heavily subsidized child care, mothers contribute much less to child care themselves, making it more likely that women earning higher pay have larger families. At low to moderate levels of wages, a higher mother’s wage tends to reduce fertility, but its negative impact attenuates as her wage rises, or the price of child care falls, because mothers purchase a larger proportion of child care time. The impact of the price of child care on fertility displays a similar interaction, becoming more negative as the mother’s wage rises.

Nevertheless, when examining changes over time, the cost of purchased child care and women’s wages tend to move together, because women’s labour is such an important input to the provision of child care services. Thus, over time we may still expect women’s pay relative to men’s and fertility to be negatively related, because higher women’s pay raises the cost of children.

The ultimate manifestation of low child quality is a child not surviving to adulthood. In light of the *demographic transition* (i.e. the change from a high fertility–high child mortality environment to a low fertility–low mortality one), an interesting question is how fertility responds to changes in the “risk” of child mortality. An autonomous increase in the probability of child survival (e.g. from better water supply or public health) has conflicting impacts. On the one hand, it reduces the price of a surviving birth, thereby encouraging higher fertility. But if parents can influence the chances that their own children survive to become adults by spending more on each child, then it is *possible* that better chances of child survival reduce fertility, provided that exogenous factors affecting child survival substitute for parents’ expenditure.

If we wish to consider decisions about the timing of births, imperfect fertility control, or the consequences of unexpected outcomes like birth control failure or child mortality, a dynamic model is needed. Chapter 6 surveys some of these.

1.6 MATCHING IN THE MARRIAGE MARKET

The process of finding a spouse is one in which information is scarce, and it takes time to gather it. These market frictions affect who marries whom, the gains from each marriage and the distribution of gains between spouses. From an individual woman’s (or man’s) point of view, higher welfare when single, faster arrival of marriage offers and a higher maximum attainable offer allow her to be choosier when selecting a husband. A higher discount rate makes her less choosy, and a higher divorce rate has the same effect because it reduces the perceived benefits from waiting for a better match by making it more likely that a woman will return to the single state.

The behaviour of both sexes is integrated in a marriage market equilibrium in Chapter 7. Suppose first that people’s utility from a marriage depends on their “type” of partner, which can be characterized by various attributes associated with their “attractiveness” as a husband or wife, and that there is no way to “transfer utility” between spouses. The latter assumption means that an individual who would obtain large gains from a match with a particular partner cannot compensate that potential partner to ensure the match is made. Then if marriage market frictions are not too large, positive assortative mating by attractive attributes emerges.

Alternatively, if we assume away frictions, but allow “transferable utility”, there is, in effect, a price mechanism that ensures that jointly efficient matches are made and that each match can be characterized by the “total utility” it generates. Suppose that each person is endowed with

a single attribute (e.g. education), which has a positive effect on total utility from the marriage. Positive (negative) assortative mating with respect to the attribute occurs when attributes are complements (substitutes) in the production of total utility in the marriage, in the sense that the marginal product of one person's attribute is increasing (decreasing) in the attribute of the spouse. With frictions in searching for a partner, it is no longer the case that complementary inputs necessarily generate positive assortative mating.

In the absence of search frictions, the equilibrium outcome is socially efficient. But search frictions produce "sorting externalities", which lead to an inefficient equilibrium. When a man and woman meet, they only match if it is jointly efficient to do so, but by leaving the marriage market they change the composition of types in the market, which affects the expected returns to search for single persons in the market. Their failure to take into account the impact of their match on the welfare of singles in the market produces the inefficiency.

Marriage market frictions also open the possibility of childbearing outside marriage. When a man and women meet, the man can choose to marry the woman, or not, if she will have him. While a woman faces the same choice when she meets a man, she can also choose to have a child by the man and then raise it without the father. Depending on the social welfare system she faces, and whether the father is willing to contribute resources, a woman's welfare when raising a child by herself may be greater than what she obtains when single and childless. But there are also costs in terms of marriage market prospects associated with raising a child alone. A single woman with child may find it more difficult to contact potential husbands while looking after a child. A woman who contacts a man she does not wish to marry, or who will not marry her, would choose to have a child by the man if the short-run gain exceeds the long-term costs in terms of her marriage prospects. Those women who expect to obtain a significant increase in welfare when they marry suffer a greater long-term cost by having a child while single than women whose marriage prospects are such that they expect to gain little from marriage. Thus, women with poorer marriage prospects are more likely to have children outside marriage.

1.7 DIVORCE AND CHILD SUPPORT

As mentioned earlier, expenditure on children, such as investment in their human capital, is considered to be a public good to the parents. When living together, they choose the efficient level of this public good. But after breaking up, the mother usually obtains custody of the children and she

decides the level of expenditure on children. Her former husband can only influence it by making transfers to his former wife. This is plausible because the father cannot usually monitor the division of his transfer between expenditure on children and the mother's consumption, particularly expenditure on young children. The allocation of resources to child expenditure implied by this "contract" is not efficient, because the mother does not take into account the effect of her choices on the welfare of the father. The inefficiency can be interpreted as an *agency* problem—the father can only indirectly affect child expenditure through his ex-wife's choices.

Any such transfers from the father to the mother are voluntary on his part. He makes them because more expenditure on children increases his welfare. He will only make such transfers if his income is "high enough" relative to the mother's income, and this threshold depends on his preferences for child expenditure relative to hers. He will transfer more to the mother the higher his income and the lower is hers. A key feature of the relationship between transfers and child expenditure is that he must transfer more than \$1 to obtain \$1 more expenditure on children, because the mother spends part of the transfer on herself. In other words, he faces a higher effective price for child expenditure when divorced than when he was married, encouraging him to spend less on children after divorce (perhaps nothing), resulting in a lower level of expenditure on children overall.

The probability that a couple divorce is inversely related to the efficiency loss associated with divorce. It is smaller the higher is either spouse's income, because the efficiency loss is larger for a higher desired level of expenditure on children, which increases with income.

Courts or government agencies often stipulate a minimum level of child support payments. Enforcement is not, however, likely to be perfectly effective. A policy that provided better enforcement of child support orders would either increase child support transfers or have no effect on them (because the father already paid more than the ordered amount). Such a policy would also raise expenditure on children among families in which child support increased thereby improving some children's welfare and having no effect on others. But there is another view of the role of child support orders.

Divorced parents and their children could be better off if they could come to a cooperative agreement on resource allocation. There are, of course, an infinite number of such efficient allocations, each involving different amounts of transfers from the father to the mother and entailing different levels of expenditure on children. The court can resolve this indeterminacy by in effect "suggesting" a given cooperative allocation indirectly through the child support order. If the cooperative allocation

implied by that order gives each of the parents higher welfare than they would obtain in the non-cooperative contract just described, then they cooperate. There is a maximum level of child support that the father would pay in the cooperative equilibrium, and a minimum transfer that the mother would accept. If the court sets an order in this range, the divorced couple cooperate and the father pays the ordered amount of child support, even though it may be above the transfer that he would pay in the non-cooperative equilibrium. This is because the efficiency gains from cooperation are sufficient for the father to agree to cooperate and transfer more. If the support order is below the bottom of this range, the mother would not agree to cooperate, despite the fact that the order exceeds what she receives from the father in the non-cooperative equilibrium.

This model has implications for the impact of better enforcement of child support orders. If some parents were induced by the child support order to cooperate when there was no enforcement, perfect enforcement would reduce expenditure on children among this group of parents. The reason is that perfect enforcement changes the order from being a suggested efficient outcome to being the starting point of a bargaining situation between the parents in which their income distribution shifts from the father to the mother by the amount of the child support order. The result is a non-cooperative outcome, which produces lower expenditures on children. In other words, with perfect enforcement, the court becomes an agent for income redistribution rather than an arbitrator who leads some couples to an efficient allocation.

Chapter 8 shows that there are more efficient marriage contracts that specify transfers if the couple divorce. The divorce settlement in these tends to prevent a large discrepancy in each party's welfare between marriage and divorce, thereby providing partial insurance. But these contracts are not likely to be enforceable.

In light of the efficiency losses associated with divorce, behaviour within marriage is likely to be affected by its possibility. If, for example, more participation in paid employment raises future wages, it is likely that the risk of divorce encourages more paid employment by the mother during marriage and, by raising the cost of child quality, lower expenditure on children and lower fertility. These "defensive investments" are undertaken to increase utility later, when utility outcomes are uncertain because of the possibility of divorce.

Divorce law confers certain rights concerning marital dissolution. They define, for each spouse, an *outside option*. For a law that allows unilateral divorce, the outside option is divorce. Either spouse can, without consent of the other, force the marriage to dissolve. With a law requiring mutual consent for divorce, the outside option is marriage. Either party can

refuse to divorce, and without consent of the other, force the marriage to continue.

There are situations in which the law does not matter. Even if, for example, the husband would gain from divorce, his wife may be able to compensate him for staying in the marriage by some change in the way the marriage is conducted. Thus, even if there were unilateral divorce, he would not seek a divorce. In effect, he sells his right to divorce under a unilateral divorce law. A husband who gains from divorce may be able to compensate a wife who loses from divorce by a suitable divorce settlement so that she would also be better off from divorce. Thus, even under a mutual consent law, they would divorce. In effect, she sells her right to the marriage under a mutual consent law.

There are, however, also situations in which divorce would occur under a unilateral divorce law, but not with a divorce law requiring mutual consent, and vice versa. In these situations, a change in the law, including laws relating to marital property and divorce settlements, could affect whether couples divorce, and the outcome is efficient.

1.8 NON-ALTRUISTIC FAMILY TRANSFERS

Can inter vivos transfers between generations within a family be explained by pure self-interest? In the usual model of consumer theory resources are transferred through time by borrowing and lending in the capital market. But suppose there is no such market, or that the person does not have access to it, say because of difficulties in monitoring loans or very large transaction costs. An extended family network including three generations at different stages of life could substitute for a capital market by arranging “loans” to its young members from its middle-aged ones and enforcing repayment later when the young borrowers have become middle-aged and the middle-aged lenders have become old. In this situation, selfish people only have children because they are needed to transfer resources through time. Chapter 9 shows how such an intra-family transfer system can work.

The family transfer rules are set so that it is not possible to devise a different set of family rules that makes any generation better off without making another generation worse off. Once established, these rules would persist over generations until there is a change of circumstances outside the family. Each generation would, of course, prefer that transfers to it when a child be as large as possible and that its transfers to aged parents when middle-aged be as small as possible. But since everyone needs the family transfer rules to survive in old age, the prospect of receiving no support from the family in old age deters any member from disobeying

the rules. These rules also determine fertility. Larger transfers to the elderly relative to those to children increase the fertility required to sustain the family system.

Now suppose that the family has access to a capital market, but that nobody outside the family would lend to a child. The middle-aged now have a choice between providing for their old age by lending in the capital market or by staying within the family system and “lending” to their children, awaiting transfers from them in old age. If the market interest rate is high enough, the middle-aged would be better off by lending to the market than remaining in the family system. A threat of no support from the family in old age is no longer a deterrent, because they can make their own provision for old age through the market.

The opening of a capital market offering a sufficiently high interest rate, or an unexpected rise in the interest rate to such a level in an existing market, toll the death knell for this family system of transfers. Childbearing would also cease in this model of selfish persons. In broad terms, the prediction of this model is consistent with the observation that the growth of the financial sector (or introduction of a state pension system) tends to coincide with a sharp fall in fertility and a decline in private transfers from the middle-aged to their elderly parents. The fact that fertility does not fall to zero, even for couples who make no contribution to the consumption of their own parents (and expect their children to do the same), suggests that the demand for children is not entirely derived from the need for transfers from them to finance consumption in old age.

Adult children can also provide “services” to their parents that do not have clear market substitutes, such as companionship, attention and conforming their behaviour to their parents’ wishes. An increase in such services tends to reduce a selfish child’s well-being because it undermines his(her) independence and may use scarce non-working time. Parents make transfers to their child in exchange for these services. Higher parents’ income always increases their demand for services and therefore transfers. Higher child’s income has an ambiguous effect on transfers. On the one hand, transfers tend to decrease because higher child’s income increases joint family income, which increases the parents’ consumption. On the other hand, the parents must compensate the child by more to achieve the same level of services because the welfare the child requires to participate in the service arrangement increases with his(her) income. If the latter effect dominates, then the positive impact of child’s income on transfers from parents is the opposite to that when altruistic motives dominate parents’ transfer decisions.

Parents’ inter vivos transfers to their children could also reflect a situation in which parents have access to a capital market, but their adult children are not able to borrow against their future income. This can

arise because young adults have not yet established their reputation with lenders. Even though many are a good credit risk, financial intermediaries may not lend to them because they do not know this. The parents of these young adults have an informational advantage on other lenders. They have better information about whether their child will repay a loan. Selfish parents would exploit their bargaining power and charge their child an interest rate in excess of the market rate.

Higher income for the child when he/she is a young adult (for given future income) raises the family wealth. In this intra-family lending arrangement, the parents wish to share in this higher family wealth in terms of higher first period consumption, which encourages them to make smaller transfers. But higher child's income also increases the bargaining power of the child because his/her welfare outside the family lending arrangement increases, and this improves the terms of the family loan from the child's point of view, leading to larger transfers. The net effect is unclear, but it is possible that a higher child's income could increase transfers to him/her, in contrast to what altruistic motives would suggest. Furthermore, higher parents' income does not affect transfers, because these are determined by the child's demand for loans. Also, the probability of receiving transfers is directly related to the child's income later in life, another prediction that contrasts with what altruistic motives would suggest (i.e. no relation).

Chapter 9 also considers the interaction between transfers from parents to their children and the labour market effort of children (e.g. their labour supply). A child's earnings are determined in part by the effort that he/she expends, but also by luck. His/her effort may not, however, be observed by his/her parents (i.e. it is private information), and while parents want to help their children financially when they need it, they also want them to behave responsibly in the sense of expending sufficient effort to support themselves. Transfers from parents may decline or increase with higher child's income depending on the balance of altruistic motives and the aim to provide an incentive for high effort.

1.9 HOUSEHOLD FORMATION

Sharing housing and other consumer durable goods is an implicit transfer that can fully or partially substitute for financial transfers motivated by either exchange or altruistic motives. It is a key factor in household formation decisions (for reasons other than marriage). As a child who becomes an adult starts out in the parental home, there is likely to be asymmetry in bargaining power between them. If the parents are selfish, they could use their bargaining power to extract the child's gain from the

joint consumption economies that arise when they live together. As the predictions of such a model, derived in Chapter 10, are opposite to most empirical evidence, the analysis focuses on parents with altruistic preferences.

It suggests that higher parental income should reduce the probability that the child lives apart from his(her) parents, while higher child's income should increase this probability. When parents do not make financial transfers, higher parental income increases the chances of co-residence because it increases the amount of (joint) housing consumption in the parental home relative to that when living apart. Without financial transfers, higher child's income means that he(her) can more easily afford to purchase his(her) own housing. If financial transfers are made when living apart, but not when living together, higher parental income increases the chances of co-residence because parents would like to provide more help to their child when their income is higher and it is cheaper to do so when living together because of the public good aspect of housing. Conversely, when the child's income is higher, parents choose to provide less help to their child, thereby reducing the need for co-residence to provide support.

The impact of the price of housing on the probability of living apart is intimately related to the price elasticity of parents' housing demand when parents do not make financial transfers to their child. When it is less than a critical value, a higher price of housing reduces the probability that the young adult lives apart from his(her) parents, but the opposite is true if it is above the critical value. These predictions reflect the fact that a higher housing price reduces the child's welfare in the parental home as well as when he(her) lives away from home. If parents did not adjust their housing consumption (zero price elasticity of housing demand), then a young adult's housing and welfare in the parental household would not change, while his(her) utility when living apart would fall; thus, the probability of living apart would fall. When the parents' housing response is relatively small (inelastic housing demand), this fall in the probability continues to hold. If, however, parents' housing response is relatively elastic, a higher housing price entails that the young adult's utility falls more in the parental household as a consequence of the large decline in the public good available when living with parents.

1.10 SOCIAL INTERACTION

An individual's preferences, and therefore behaviour, may depend on what others in society are perceived to be doing. This may take the form of "learning" or "social influence". Certain decisions, particularly

concerning the use of new technology, such as modern contraception, are subject to substantial uncertainty. Learning about other people's experiences through social interaction may reduce this uncertainty and make it more likely that a person adopts a new technology. Social influence captures the possibility that a person's preferences may be altered by those with whom the person interacts. For instance, childbearing outside marriage may be discouraged by social stigma when non-marital births are rare, but this stigma may be eroded as more childbearing is outside marriage.

If this type of social influence is large enough, it is possible that there is more than one stable equilibrium: one in which the phenomenon (e.g. non-marital childbearing) is rare and another in which it is common. If so, "history matters" for the selection of the low-level or high-level equilibrium. Furthermore, *temporary* changes in the socio-economic environment that alter, for example, non-marital childbearing behaviour and/or expectations, can produce dramatic changes in the proportion who become single mothers.

If social influence is relatively strong, but not large enough to produce multiple equilibria, small changes in the socio-economic environment, such as higher state benefits for single mothers, can still produce large changes in the proportion becoming single mothers and other social phenomena. There is a "multiplier effect" of such changes in the "fundamental" determinants of differences in utility between two actions. In other words, each person's actions change not only because of the direct change in some fundamental determinant, but also because of the change in the behaviour of their peers. If this social multiplier is large, populations with slightly different distributions of attributes or preferences could, for example, exhibit very different proportions of women who become single mothers. Social influence or social learning may also explain large and rapid fertility decline in one country or region while there is little change in another, which is similar in terms of socio-economic conditions.

Chapter 11 explores how information on the density of social networks can be used to distinguish between the dominance of social influence or social learning on fertility behaviour. If learning dominates, then both dense and sparse networks containing a larger proportion of women using modern contraceptive methods should increase the chances that a woman will adopt these methods. Because sparse networks are more efficient sources of information, the impact of density on these chances is either zero or negative. By integrating a woman into a larger group, dense networks are more likely to constrain a woman's ability to deviate from prevailing behaviour; that is, they exert a stronger normative influence than sparse networks. A woman's only alternative to agreeing with

other members of this group may be to leave the group. If social influence dominates, then the proportion of women in a woman's social network using modern methods would have a weak effect on her adoption of these methods in sparse networks, but a strong effect in dense networks. That is, the impact of the proportion increases with the density of the network.

Social multipliers and multiple equilibria can also arise through market interactions. For instance, the expected gain from divorce depends on the prospects of remarriage. These prospects depend on the decisions of others to divorce and remarry. If many couples are expected to divorce, then the prospects of remarriage are high because there are more people in the remarriage market. Divorce is then less costly and each particular couple is more likely to decide to divorce. If instead the divorce rate is expected to be low, then divorce is more costly and is less likely to occur. Either a high-divorce or a low-divorce equilibrium may be supported with the same set of fundamental factors affecting divorce decisions. Much of this book focuses on decisions at the individual level and the impacts of prices and resources on them, but chapter 11 indicates that there are often these important feedback effects of the choices of one's peers on one's own choices.

This introductory chapter has suggested how economic analysis can derive predictions about family behaviour, particularly its response to developments in markets, to technological developments and to public policy. The remaining chapters provide details of these analyses, which hopefully provide the foundation for readers to apply these methods to a range of issues related to family behaviour.