This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Demography and the Economy

Volume Author/Editor: John B. Shoven, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-75472-3
ISBN13: 978-0-226-75472-7

Volume URL: http://www.nber.org/books/shov08-1

Conference Date: April 11-12, 2008

Publication Date: November 2010

Chapter Title: Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?

Chapter Authors: Larry E. Jones, Alice Schoonbroodt, Michèle Tertilt

Chapter URL: http://www.nber.org/chapters/c8406

Chapter pages in book: (43 - 100)

2

Fertility Theories
Can They Explain the Negative Fertility-Income Relationship?

Larry E. Jones, Alice Schoonbroodt, and Michèle Tertilt

2.1 Introduction

Empirical studies find a clear negative relationship between income, or wages, and fertility. This finding has been confirmed across time and for different countries. For example, Jones and Tertilt (2008) document a negative cross-sectional relationship between income and fertility in the United States and find that the relationship has been surprisingly stable over time. In particular, the paper shows a negative relationship for thirty birth cohorts between 1830 and 1960, with the income elasticity of fertility remaining roughly constant at about –0.30.[1]

Why do richer people have fewer children, and what explains the relatively time-invariant nature of the relationship? The negative correlation is particularly puzzling if one thinks about children as a consumption good, unless one believes that children are an inferior good. An early discussion of

Larry E. Jones is chair and professor of economics at the University of Minnesota, a visiting scholar at the Federal Reserve Bank of Minneapolis, and a research associate of the National Bureau of Economic Research. Alice Schoonbroodt is a lecturer in economics at the University of Southampton. Michèle Tertilt is an assistant professor of economics at Stanford University, a research affiliate of the Center for Economic and Policy Research, and a faculty research fellow of the National Bureau of Economic Research.

We thank Todd Schoellman, John Knowles, and the participants at the NBER preconference in Boston, the Stanford Junior Faculty Bag Lunch, and the Economics and Demography conference in Napa California for helpful suggestions. We thank Amalia Miller in particular for a thoughtful discussion. Financial support by the NSF (grants SES-0519324 and SES-0452473) and the Stanford Institute for Economic Policy Research (SIEPR) is greatly appreciated. William G. Woolston provided excellent research assistance. Part of this research was completed while Michèle Tertilt was a National Fellow at the Hoover Institution at Stanford.

1. We discuss the empirical evidence in more detail in section 2.2.

this fact appears in the seminal article on fertility choice by Becker (1960). Indeed, this puzzling correlation was the main impetus behind Becker's early work.[2] The ensuing literature can be roughly divided into two strands. One attacks the question from a theoretical point of view and finds that, properly interpreted or with the appropriate additions in choice variables, economic theory says that fertility *should* be negatively related to income. The basic idea is that the price of children is largely time, and because of this, children are more expensive for parents with higher wages. Another argument is that higher-wage people have a higher demand for child quality, making quantity more costly, and hence those parents want fewer children. The other strand of literature attacks the question from an empirical point of view, arguing that the negative relationship is mainly a statistical fluke—due to a missing variables problem. This literature focuses on identifying those crucial missing variables, such as female earnings potential. Once those missing variables are controlled for, fertility and income—so the argument goes—are actually positively related.[3]

In this chapter, we revisit these theories of the cross-sectional relationship between income and fertility. They are largely based on ability or wage heterogeneity. We also formalize a new theory, based on heterogeneity in the taste for children, in which wages are also endogenous. For each of the theories, we catalogue whether they basically never work (i.e., never produce the negative income-fertility relation), whether they work only with specific additional assumptions, or whether they are relatively robust to changes in assumptions. We also often compare the results to the conditional correlations found in the statistical strand of the literature. For those theories that work sometimes, we try to be as explicit as possible about what kinds of conditions are needed (e.g., curvature and/or functional form restrictions) to generate a negative relationship between income and fertility. We also show what goes wrong by giving examples about how they fail. Finally, of the theories that work and appear robust, we ask for more. Can the theory also match the time series properties of fertility? If so, what exactly does it take? If not, why not? Finally, we want to know whether such a theory is consistent with a recursive formulation of dynastic altruism.

Our main findings can be summarized as follows:

2. Quoting from Becker (1960, 217): "Having set out the formal analysis and framework suggested by economic theory, we now investigate its usefulness in the study of fertility patterns. It suggests that a rise in income would increase both the quality and quantity of children desired; the increase in quality being large and the increase in quantity small. The difficulties in separating expenditures on children from general family expenditures notwithstanding, it is evident that wealthier families and countries spend much more per child than do poorer families and countries. The implication with respect to quantity is not so readily confirmed by the raw data. Indeed, most data tend to show a negative relationship between income and fertility." See also the discussion in Hotz, Klerman, and Willis (1993).

3. See Hotz, Klerman, and Willis (1993) for a survey. An early literature review on fertility choice is Bagozzi and Van Loo (1978).

1. (Almost) all theories depend on the assumption that raising children takes time and that this time must be incurred by the parents.

2. Theories based on exogenous wage heterogeneity crucially depend on the assumption of a high elasticity of substitution between consumption and children.

3. Adding a quality choice by itself does not generate a negative fertility-income relationship. The quantity-quality trade-off works only in conjunction with assumptions similar to those needed in list entry (2).

4. Theories based on heterogeneity in tastes for children are able to generate a negative fertility-income relationship without requiring a high elasticity of substitution between consumption and children.

5. Theories that explicitly distinguish between fathers and mothers are very similar to one-parent theories. However, to get fertility to be decreasing in men's income, one needs to assume that there is positive assortative matching of spouses.

6. Several of the theories that match the cross-sectional patterns of fertility also match, at least loosely, some of the broad time series trends in fertility. Theories based on wage heterogeneity produce this relationship more naturally.

7. Extending the models that are successful at matching the cross-sectional properties of fertility choice to fully dynamic models based on parental altruism is very challenging. Basic theories with wage heterogeneity do not appear to be robust to this extension. Theories based on heterogeneity in tastes are more promising, but leave many open questions.

Our findings may be relevant in several different contexts. First, there has been a recent increase in research relating the demographic transition and economic development among macroeconomists.[4] Similarly, several recent contributions try to understand why fertility is higher in poor countries than in rich ones.[5] Further, there is a recent literature that uses dynamic macro-style models to analyze the interplay between fertility, labor force participation, marriage, and inequality[6]—including studies of gender wage gap[7] and the baby boom following World War II.[8] Often dynamic macro-style models are used to analyze the impacts of various policy changes—for example, parental leave policies, the impact of tax reform, welfare reform, and social

4. See, for example, Becker, Murphy, and Tamura (1990); Galor and Weil (1996, 1999, 2000), Greenwood and Seshadri (2002); Hansen and Prescott (2002); Boldrin and Jones (2002); Doepke (2004, 2005); Greenwood, Seshadri, and Vandenbroucke (2005); Moav (2005); Tertilt (2005); Jones and Schoonbroodt (Forthcoming), Murtin (2007); and Bar and Leukhina (Forthcoming). See Galor (2005a, 2005b) for an extensive analysis and a critical survey of theories of the demographic transition.

5. See Manuelli and Seshadri (2009).

6. See Alvarez (1999); Caucutt, Guner, and Knowles (2002); and Falcão and Soares (2008).

7. See Erosa, Fuster, and Restuccia (2005).

8. See Greenwood, Seshadri, and Vandenbroucke (2005); Doepke, Hazan, and Maoz (2007); and Jones and Schoonbroodt (2007).

security.[9] Typically, they use an "off-the-shelf" fertility model as one of their building blocks, and need to make a careful decision about which one to use. What may help guide this choice is an informed understanding of the implications of the models for the fertility-income relationship in the cross section. Because of this, it is natural to use successful models of the cross sectional properties of fertility as a way to inform that choice.

This is easier said than done, however. Economists have been developing and testing theories of fertility ever since Gary Becker's seminal paper, but still there is no full consensus on the motivations behind fertility choices. Here, we provide a systematic comparison of the properties of various fertility theories. We hope that this catalogue may be a useful step toward finding a consensus.

This chapter is organized as follows. In the next section, we summarize the empirical evidence on the fertility-income relationship. Section 2.3 describes a basic model with wage heterogeneity. Section 2.4 develops a new theory based on preference heterogeneity in the desire to have children, which generates endogenous wage heterogeneity. Section 2.5 adds quality to the basic model. In section 2.6 we depart from the simplest framework and analyze more realistic theories with two parents. We investigate whether theories are robust to allowing parents to hire nannies in section 2.7. Section 2.8 pushes several of the working theories to also address the secular decline in fertility, while section 2.9 concludes. The appendix analyzes the extent to which our results apply to a dynastic formulation of fertility.

2.2 Data on Fertility and Income

A robust fact about fertility is that it is decreasing in income. This fact has been documented from a time-series point of view, across countries, and across individuals. Quoting from Becker (1960, 217): "Indeed, most data tend to show a negative relationship between income and fertility. This is true of the Census data for 1910, 1940 and 1950, where income is represented by father's occupation, mother's education or monthly rental; the data from the Indianapolis survey, the data for nineteenth century Providence families, and several other studies as well."[10]

In a recent study, Jones and Tertilt (2008) use U.S. Census Data on lifetime fertility and occupations to document this negative cross-sectional relationship in the United States.[11] They find a robust negative cross-sectional rela-

9. Recent contributions include Aiyagari, Greenwood, and Guner (2000); Erosa, Fuster, and Restuccia (Forthcoming), Fernandez, Guner, and Knowles (2005); Greenwood, Guner, and Knowles (2003); Sylvester (2007); and Zhao (2008).

10. The studies Becker is referring to are U.S. Census (1945, 1955); Whelpton and Kiser (1951); and Jaffe (1940).

11. Income is based on the median annual income for a given occupation in 1950 and adjusted for TFP growth. A measure of income based on occupation is a better measure of

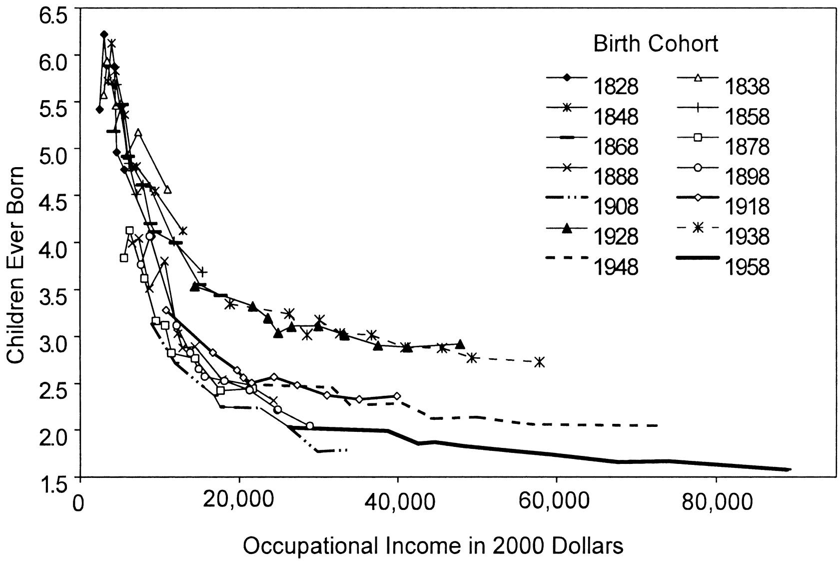

Fig. 2.1 Fertility by occupational income in 2000 dollars
Source: Jones and Tertilt (2008).

tionship between husband's income[12] and fertility for all cohorts for which data is available; that is, for women born between 1826 and 1960.[13] Not only are the correlations always negative, but also they are surprisingly similar in magnitude over time. Figure 2.1, reproduced from their paper, shows this very clearly. While the relationship is not perfect, it seems that most of the fertility decline over time can be "explained" by rising incomes alone, at least in a statistical sense.

To give a sense of the magnitudes, table 2.1 reproduces some of the most relevant numbers from Jones and Tertilt (2008). For a selected number of

lifetime income than income in any particular year. See Ruggles et al. (2004) for a description of how occupational income scores (OIS) are constructed as well as its robustness as a proxy for income. See Jones and Tertilt (2008) for a description of how the OIS was converted into 2000 dollars.

12. The focus on husband's income allows a consistent analysis over time. In particular, it allows the analysis of periods for which data on wife's income is practically nonexistent.

13. Fertility is measured as children ever born (CEB) to the current wife. Of course, this measure could differ from male completed fertility if men had children with different women. Unfortunately not much data on male completed fertility are available. We are aware of two exceptions. First, the 2002 National Survey of Family Growth asked men and women independently about their fertility. Preston and Sten (2008) use this data to construct a measure of the elasticity of male fertility to male education and also find a negative coefficient. Given that divorce was rare for most of the period under consideration, we believe that the wife's fertility is a good proxy. Second, Shiue (2008) compiled Chinese data from 1300 to 1850. She finds a weak positive relationship between male fertility and social status, but since richer men also had more women on average, fertility per wife is actually decreasing.

Table 2.1 **Fertility-income relationship for 14 U.S. cross sections**

Birth cohort	Income elasticity	Top/bottom fertility gap	Fertility	Annual income in 2000 dollars	Number of observations
1826–1830	–0.33	0.95	5.59	4,154	452
1836–1840	–0.20	0.74	5.49	5,064	1,960
1846–1850	–0.32	1.26	5.36	6,173	4,520
1856–1860	–0.35	1.24	4.90	7,525	7,241
1866–1870	–0.34	1.27	4.50	9,173	7,347
1876–1880	–0.42	1.06	3.25	11,182	3,203
1886–1890	–0.45	1.05	3.15	13,631	6,644
1896–1900	–0.50	0.93	2.82	16,616	8,462
1906–1910	–0.42	0.57	2.30	20,255	11,812
1916–1920	–0.25	0.34	2.59	24,690	46,908
1926–1930	–0.17	0.27	3.11	30,097	97,143
1936–1940	–0.19	0.31	3.01	36,688	44,428
1946–1950	–0.20	0.26	2.22	44,723	62,210
1956–1960	–0.22	0.23	1.80	54,517	71,517

Source: Jones and Tertilt (2008).

birth cohorts, the table displays average husband's income and average fertility.[14] To quantify the fertility-income relationship, two different empirical measures were constructed: the income elasticity of fertility, and the fertility gap between the top and bottom 50 percent of the income distribution. The income elasticity roughly hovers around minus one-third, meaning that for a family with an income that is 10 percent higher than another family, the number of children is about 3 percent lower. This is a large difference. For example, for women born during the nineteenth century, those in the bottom half of the income distribution had easily one child more on average than those in the top half. Today, the difference is much smaller in absolute numbers, with a fertility gap of roughly a quarter of a child. But since fertility is significantly lower for all women, the income elasticity has declined only very mildly over time, to about –0.20 for the most recent cohorts.

Note that the income measure used in figure 2.1 and table 2.1 is based on occupations, and can also be viewed as a proxy for wages. Therefore, the findings can be interpreted as showing a negative fertility-wage relationship.

Many other studies have documented this kind of relationship, typically for a specific geographic area at a particular point in time. For example, Borg (1989) finds a negative relationship using panel data from South Korea in 1976, and Docquier (2004) documents a similar relationship for the United States using data from the Panel Study of Income Dynamics (PSID) in 1994.

14. The definitions of fertility and income in the table are identical to those used in figure 2.1.

Westoff (1954) finds a negative relationship between fertility and occupational status for the years 1900 to 1952 using U.S. Census data.

Part of the literature argues that a negative income-fertility relationship is primarily a statistical fluke—that is, it is due to a problem of missing variables. The idea is that once enough variables are controlled for, one would actually find a positive income-fertility relation. Indeed, this was Becker's original view on the topic. He went into great detail focusing on knowledge of the proper use of contraceptives as the important missing variable.[15] Similarly, many authors have argued that a distinction between male and female income is crucial and that the relationship between male income and fertility is indeed (weakly) positive once one correctly controls for female income.[16] Authors of studies that find a positive relationship after controlling for women's wages often interpret such finding as having resolved the "puzzle." This is, however, not necessarily the case. The reason is that even though the finding reconciles the conditional correlations in the data with the simplest model of fertility, the question remains of what kind of theories would explain the unconditional negative correlation of men's wages and fertility. At the very least it requires some assumptions about matching.[17] In this chapter we take a somewhat different approach: rather than controlling for important factors (such as wives' wages) in the data, we try to add such important factors into the model and then ask whether the augmented model delivers the same qualitative facts as the data does.

It is sometimes argued that early on in the development process, a positive relationship between income and fertility existed.[18] Most of the studies that document such a positive relationship are set in agrarian economies, and often income is proxied by farm size. Examples include Simon (1977, chapter 16), who documents a positive relationship between farm size in hectares and the average numbers of children born for rural areas in Poland in 1948, and Clark and Hamilton (2006), who document a positive relationship between occupational status and the number of surviving children in England in the late sixteenth and early seventeenth century (see also Clark

15. He showed that, in his sample, in those households that were actively engaged in family planning, fertility and income were positively related, while the opposite was true for families not engaged in family planning. Other early papers along this line are cited by Becker in his original piece. They include Edin and Hutchinson (1935) and Banks (1955).

16. Empirical studies distinguishing explicitly between husbands and wives include Cho (1968); Fleischer and Rhodes (1979); Freedman and Thorton (1982); Schultz (1986); Heckman and Walker (1990); Merrigan and Pierre (1998); Blau and van der Klaauw (2007); and Jones and Tertilt (2008). The findings are mixed.

17. We discuss this in detail in section 2.6.

18. A more recent version of such a positive relationship is that U.S. fertility is higher than most other countries in the Organization for Economic Cooperation and Development (OECD) even though U.S. income is higher. This does not hold for a larger set of countries, however. See Ahn and Mira (2002) and Manuelli and Seshadri (2009) for a discussion of related points. Bongaarts (2003) finds a slight U-shaped fertility-education relationship in Portugal and Greece using three education levels of women. The other eight countries concur with previous findings of a strictly negative relationship.

[2005] and Clark [2007]). Weir (1995) finds a weakly positive relationship between economic status and fertility in eighteenth century France, while Wrigley (1961) and Haines (1976) document higher fertility in the coal mining areas of France and Prussia than in surrounding agricultural areas during the end of the nineteenth century. Also, Lee (1987) documents a similar finding using data from the United States and Canada.[19] This body of work suggests that the fundamental forces determining the demand for children might be different in areas where agriculture is the primary economic activity.

Of course, there is no reason why the fertility-income relationship should not change over time or vary in different cross sections. It may be that in some subgroups of the population, fertility increases in income once all other relevant correlates are controlled for, while in other subgroups the primary change across the income distribution is in the price of a child and, because of this, that fertility is lower at higher income levels. And in fact, it is plausible that fertility and wealth were indeed positively related in early agrarian economies, but that this relationship was reversed after industrialization.[20]

To sum up, the fact that people with higher lifetime earnings have fewer children seems very robust, at least during the last century and a half in the United States. Other countries and other episodes display a similar relationship. Inspired by these facts, this chapter analyzes which theories of fertility are consistent with this relationship.

2.3 Basic Framework and Results

In this section we introduce notation and explore some basic models of fertility choice. The basic examples that we discuss here focus on the roles played by the nature of the cost of children, the sources of family income, and the formulation of preferences. We find that the simplest versions of these ideas do not generate a negative relationship between fertility and income. Special assumptions on the nature of costs of children, the utility function, the sources of income, and/or the child quality production function are needed. This is not to say that these theories are wrong. Rather, by making explicit the assumptions behind the ideas we hope to facilitate the testing of the theories and, ultimately, to improve our understanding of fertility decision-making.

19. See also the papers cited in Lee (1987).

20. For example, Skirbekk (2008) (using a large data set including various world regions over time) finds that as fertility declines, there is a general shift from a positive to a negative or neutral status-fertility relation. Those with high income/wealth or high occupation/social class switch from having relatively many to fewer or the same number of children as others. Education, however, depresses fertility for as long as this relation is observed (early twentieth century).

To keep the analysis tractable, we focus on a static, monoparental setup. This approach allows for closed form solutions and lets us focus on the basic mechanics behind the results. Obviously, there are many dynamic elements in real world fertility decision-making; for example, choices about the timing of births, and so forth. We see our basic examples as a way to gain insights into modeling ingredients of more complex dynamic models. Clearly, many important features are left out in the simplest example we start with. Some of these features are particularly important and we come back to those in later sections of this chapter. One such element is that any child necessarily has a father and a mother. In fact, many authors have emphasized that it may be female time rather than male time that is important to generate the negative relationship between fertility and income. We get back to this in section 2.6. In later sections of the chapter we extend the model to include more dynamic elements, including limited forms of human capital/child quality (sections 2.4 and 2.5) and parental altruism (appendix).

Two more caveats are in order. First, throughout the chapter we analyze only rational theories of fertility.[21] Behavioral concerns might be relevant, especially for teenage childbearing, but are not considered here. Second, we focus on theories in which children provide direct utility benefits; that is, children are a consumption good. Note that children are sometimes also viewed as an investment, providing old-age security.[22] While the investment motive may have important implications for the fertility-income relationship, this analysis is beyond the scope of this chapter and is left for future research.

2.3.1 The Basic Model

The general static model of fertility choice that we consider is as follows. People maximize utility subject to a budget constraint, a time constraint, and a child quality production function. People (potentially) derive utility from four different goods: consumption, c, number of children, n, the average quality of children, q, and leisure, ℓ. Producing children takes b_0 units of goods and b_1 units of time (per child). We let l_w denote the time spent working and normalize the total time endowment to one. The wage per unit of time is denoted by w. In addition to labor income, we also allow for nonlabor income, y. Finally, child quality is a function of educational child inputs, s (we abstract from direct parental time inputs into child quality). Thus, the choice problem is as follows:

21. We also abstract from costs and technologies to prevent births or to inseminate artificially. Several authors have given these issues more thought, and we refer the reader to them (see, e.g., Hotz and Miller (1988); Goldin and Katz (2002); Bailey (2006); and Greenwood and Guner, (Forthcoming)).

22. Examples include Ehrlich and Lui (1991); Boldrin and Jones (2002); and Boldrin, De Nardi, and Jones (2005). Zhao (2008) uses the Boldrin-Jones framework to jointly address the fertility decline and the narrowing of fertility differentials by income in response to changes in social security.

$$\text{(1)} \qquad \max_{c,n,q,e,l_w} \quad U(c, n, q, \ell)$$
$$\text{s. t.} \quad l_w + b_1 n + \ell \leq 1$$
$$c + (b_0 + s)n \leq y + wl_w$$
$$q = f(s).$$

In order to highlight the crucial ingredients to generate a negative income (or wage) to fertility relationship, we distinguish between various combinations of utility specifications, concept of wealth/income/earnings used, costs of children, and quality production functions. We now briefly discuss each of these components.

Utility: We focus on separable utilities. That is:

$$U(c, n, q, \ell) = u_c(c) + u_n(n) + u_q(q) + u_\ell(\ell).$$

We consider the CES utility case, $u_x(x) = \alpha_x(x^{1-\sigma_x} - 1)/(1 - \sigma_x)$ for values of $\sigma_x > 0$. We will often distinguish three cases: (a) $\sigma_x > 1$ (high curvature, low elasticity of substitution); (b) $\sigma_x < 1$ (low curvature, high elasticity of substitution); and (c) $\sigma_x = 1$ corresponding to log utility.[23]

Income/Wealth: We use the following (standard) language: w is the wage, $W = w + y$ is total wealth, and $I = wl_w$ is earned income (often also called labor earnings). In most of our examples, there are only two uses of time (working and child-rearing), in which case earned income is equal to $w(1 - b_1 n)$. An interesting special case is where all income is labor income, $y = 0$ and $W = w$. In several examples, we focus on the fertility-earnings (rather than wage) relationship. In these examples, there is no wage heterogeneity. However, the logic underlying those examples can easily be generalized to (endogenous) wage heterogeneity. We do so in section 2.4. In this context, the wage will be equal to human capital, H, and human capital is a function of schooling inputs. For simplicity, we will omit H and say that the wage w is a function of schooling inputs.

Costs of Children: We allow for both goods and time costs, denoted by b_0 and b_1, respectively. To get starker results, we sometimes shut down one of the two types of costs. It turns out that a time cost appears to be essential to almost all the theories and examples we present here. To see this, note that with separable utility, no time cost ($b_1 = 0$) and no quality in utility ($\alpha_q = 0$), n is a normal good, and hence, it follows that n is increasing in both

23. This utility function has the added advantage that, in some cases, it can be interpreted as the problem in Bellman's equation for a Barro-Becker style dynasty with parental altruism. There, the term $u_n(n)$ is the value function for continuations. This interpretation is only valid for certain choices of the α_n's however. See appendix for details.

y and w.[24] Thus, we will typically require that $b_1 > 0$. While it seems fairly obvious that it takes time to raise a child, it is less clear whether the time spent must be the parent's time rather than a nanny or a day care center. We analyze the implications of allowing for nannies in section 2.7.[25]

Quality Production Function: One important feature for the quantity-quality trade-off to generate the desired relationship is the specification of the quality production function, $f(\cdot)$. We experiment with various specifications. Note that making special assumptions on $f(\cdot)$ is technically equivalent to making special assumptions on $u_q(\cdot)$. That is, let $v_q(\cdot) = u_q(f(\cdot))$ and make assumptions about this function. The interpretation, however, can be quite different. With homothetic preferences to start with, unless $f(s)$ is of the form $f(s) = s^\kappa$, this introduces nonhomotheticity into the overall problem (1). We will analyze quality production functions in some detail in section 2.5.

Leisure: For some of the examples in sections 2.6 and 2.7, we need leisure as an alternative use of time in order to reproduce the negative fertility-income relationship. For most examples, this is not necessary, and hence we will typically assume that $\alpha_\ell = 0$.

2.3.2 The Price of Time Theory

To highlight the necessary ingredients, we start by discussing a simple example that does not generate the desired negative relationship between fertility and income. We then show what special assumptions are needed to obtain the desired result.

Starting from the general formulation (1), we assume log utility ($u_x(x) = \alpha_x \log(x)$), no utility from child quality ($\alpha_q = 0$) or leisure ($\alpha_\ell = 0$), and no nonlabor income ($y = 0$). Then the problem reduces to

$$\max_{c,n} \quad \alpha_c \log(c) + \alpha_n \log(n) \quad \text{s. t.} \quad c + b_0 n \leq w(1 - b_1 n). \tag{2}$$

The solution for fertility is:

$$n^* = \frac{\alpha_n w}{(\alpha_c + \alpha_n)(b_0 + wb_1)}.$$

24. When $\alpha_q > 0$, the constraint becomes nonlinear, which complicates matters. In certain cases, the problem can be written in aggregate quality $Q = nq$. In this case, if $b_1 = 0$, both n and Q are normal goods and hence increase in both y and w.

25. We restrict attention to linear child costs. Analyzing the robustness of our results to other child cost specifications would be of interest. There seems to be little consensus in the empirical literature on the shape of the child cost function, however. Empirical papers that estimate the costs of children and economies of scale in the household include Hotz and Miller (1988), Bernal (2008), Lazear and Michael (1980), and Espenshade (1984). Taking maternal health and maternal mortality risk into account, one might also want to argue that a convex cost function is the most reasonable formulation (e.g., Tertilt 2005).

As is apparent from this example, as long as the goods cost of children is positive ($b_0 > 0$) higher-wage households (higher w) will have strictly *more* children in this setup. This is the *opposite* prediction from what we observe in the data. Setting the goods cost to zero with just a time cost results in fertility choice being independent of w—still, not a negative relationship. Adding leisure or child quality (say, with $q = f(e) = e$) will not reverse this result (see section 2.5).

To give the price of time theory a chance, it seems fairly obvious that a deviation from log utility is needed; that is, a specification where income and substitution effects do not cancel out. Thus, we turn now to general Constant Elasticity of Substitution (CES) utility functions. Also, since a time cost is essential here and a goods cost does not really add anything, we set $b_0 = 0$ and assume $b_1 > 0$, but reintroduce nonlabor income, $y \geq 0$. Thus, our next example takes the form

$$\text{(3)} \qquad \max_{c,n} \quad \alpha_c \frac{c^{1-\sigma} - 1}{1 - \sigma} + \alpha_n \frac{n^{1-\sigma} - 1}{1 - \sigma}$$

$$\text{s. t.} \quad c \leq y + w(1 - b_1 n).$$

It is easy to solve for a closed form solution of this specification. Optimal fertility is given by:

$$n^* = \frac{y/w + 1}{(\alpha_c b_1/\alpha_n)^{1/\sigma} w^{(1-\sigma)/\sigma} + b_1}.$$

Elasticity of substitution: In problem (3) wage heterogeneity leads indeed to a negative wage-fertility relationship if the right amount of curvature is assumed in the utility function. To see this, assume first that $y = 0$. If the only way in which individuals differ is in their wages, we can see that when $\sigma \geq 1$, fertility is either independent of or increasing in w. However, when $\sigma < 1$, it follows that $n^*(w)$ is decreasing.

The intuition here is simple: when the only cost of children is time, and that time must be the parents' own time, higher wage families face a higher price of children. This induces the usual wealth and substitution effects familiar from demand theory. Certainly it implies that compensated demand for children is decreasing. This is not sufficient, however, to automatically imply that the demand for children is decreasing in income, since those families that face higher prices also have more wealth. Thus, it depends on which of the two forces is stronger. If the elasticity of substitution between children and consumption is high enough (low σ), the substitution effect dominates and $n^*(w)$ is decreasing, as in the data.

Moreover, it can be seen that this relationship is approximately isoelastic when y is small and w is large relative to b_1. In this example, the income elasticity of demand for children is $(\sigma - 1)/\sigma$.

In sum, this theory works, but not without extra restrictions on preferences. An additional requirement could be that the formulation be consistent with dynamic maximization in a setting with parental altruism à la Barro and Becker (1989) (i.e., parents care about number and utility of children multiplicatively). In the first section of the appendix we discuss the relationship between this static problem and a reinterpretation of it as the Bellman equation of a dynamic problem. The difficulty with the dynamic reinterpretation of the current example is that α_n is no longer a parameter but represents children's average level of utility. It therefore becomes a function of the wage. It turns out that once this is taken into account properly, fertility is independent of the wage independently of σ. Moreover, Jones and Schoonbroodt (Forthcoming) show that in this kind of model, $\sigma > 1$ is needed to generate the decreases in fertility observed over the past 200 years in response to increased productivity growth and decreased mortality. Hence, it seems that this dynamic interpretation of the static model presented here is at an impasse to get both the cross-sectional and trend features of fertility at the same time. In the first section of the appendix, we show that with preference heterogeneity, both the cross section as well as the trend observations can be generated.

Nonlabor Income: An alternative specification that also works is to assume log utility but positive nonlabor income. Assume $\sigma \to 1$ and $y > 0$, then the solution to (3) becomes

$$n^* = \frac{\alpha_n(y/w + 1)}{(\alpha_c + \alpha_n)b_1.}$$

Note that for $y > 0$, fertility is indeed decreasing in the wage.[26] Note that the slope of the relationship depends on the size of the nonlabor income. That is, for small amounts of nonlabor income fertility is decreasing in the wage only very mildly, and in the limit, when nonlabor income is zero, fertility does not depend on the wage at all.

Note, however, that the only income that would really qualify as nonlabor income here are gifts, lottery income, bequests, and the like.[27] Since most families have no or very little such nonlabor income, it is questionable whether this should be the main mechanism by which fertility and income are connected. Yet variations of this formulation are used a lot in the literature. For example, the refinement that it is *female* time that determines the opportunity cost falls into this category. In particular, sometimes y is interpreted as the husband's income and w as the wife's wage. Then fertility

26. Adding nonlabor income effectively changes the curvature of the utility function, and hence the technical reason that makes this example succeed is similar to the $\sigma < 1$ case shown previously. The interpretation, of course, is very different.

27. Any interest income from assets that are accumulated labor earnings would be proportional to labor income, and hence would not generate the result outlined here.

is decreasing in the latter. We will turn our attention to two-parent fertility models in section 2.6.

Nonhomothetic preferences: Another way to generate the desired relationship is to move away from homothetic utility.[28] Assume, for example, that $\sigma_c = 0$. Then the problem to solve is

$$\max_{c,n} \quad \alpha_c c + \alpha_n \frac{n^{1-\sigma} - 1}{1 - \sigma} \tag{4}$$

$$\text{s.t.} \quad c \leq (1 - b_1 n)w.$$

And the solution is:

$$n^* = \left[\frac{\alpha_n}{\alpha_c b_1}\right]^{1/\sigma} w^{-1/\sigma},$$

which is clearly decreasing in w for any value of σ.[29] We are not emphasizing nonhomothetic utilities any further, because one broader aim of the proposed research agenda here is to develop a theory that encompasses cross-sectional, trend, and cyclical features of fertility choice. Embedding this example into a fully dynamic growth model has the unfortunate property that income shares to consumption tend to one. Because of this these models would be of limited use.

2.4 Endogenous Wage Differences

In the previous section we focused on theories of the cross-sectional relationship between fertility and wages in which the fundamental difference was exogenous variation in ability (wages). In this section, we explore an alternative view with an alternative causation. Suppose that the basic source of heterogeneity is in tastes for children versus material goods—some people want large families and others want to travel the world, go to fancy restaurants, and drive a sports car. This basic difference in taste for either "lifestyle" affects the investment in human capital and hence, wages. That is, parents who want large families will allocate less time to developing market-based skills in anticipation of having many children, and will therefore have lower wages and lower earned income.

Rather than assuming people differ in their taste for children, one could simply assume that people differ exogenously in fertility and choose human capital investments accordingly. This kind of model also gets the basic relationship right, and is useful for understanding the basic mechanism. We

28. See, for example, Greenwood, Guner, and Knowles (2003).

29. This specification (with $\sigma \rightarrow 1$) is used in Fernandez, Guner, and Knowles (2005); Erosa, Fuster, and Restuccia (Forthcoming); and Erosa, Fuster, and Restuccia (2005). Note that the income elasticity of demand for children here is $-1/\sigma$, which is close to the data for $\sigma = 3.0$.

start with this simple version, even though the interpretation of exogenous fertility is not straightforward. We then move to a more general case that has a more plausible interpretation: deterministic heterogeneity in the taste for children versus consumption goods. Here schooling is chosen in anticipation of fertility decisions.

Finally, as long as raising children takes time, a simpler mechanism can be considered. Again assuming taste heterogeneity, parents who choose large families will have less time available to work and hence will have lower earned income, even if wages are exogenous. This simplification will be helpful in subsequent sections. Note that whenever the simple mechanism works and one can generate a negative fertility-*income* relationship, it is straightforward to also generate a negative fertility-*wage* relationship by adding endogenous human capital investments to the model.

2.4.1 Exogenous Fertility and Endogenous Wages

The simplest version illustrating the mechanism we want to focus on is one where fertility is exogenously different across people. Let $\bar{n}_i$ be the number of children that are attached to adult i. Each child requires b_1 units of parental time. The parent solves one lifetime maximization problem by choosing how much time (net of child-rearing time) to allocate to schooling versus earning wages. Even though we write this as a one-period problem, the decisions are best interpreted in a sequential fashion: time is first spent on schooling, l_s, which determines future human capital al_s. Normalizing the wage per unit of human capital to one, al_s is also the wage, so that total lifetime income simply becomes $wl_w = al_s l_w$. The problem then is:

$$
\begin{aligned}
(5) \qquad & \max_{c, l_w, l_s} \quad \alpha_c \frac{c^{1-\sigma}}{1-\sigma} + \alpha_n \frac{\bar{n}_i^{1-\sigma}}{1-\sigma} \\
& \text{s. t.} \quad l_s + l_w \leq 1 - b_1 \bar{n}_i \\
& \qquad\quad w = al_s \\
& \qquad\quad c \leq wl_w.
\end{aligned}
$$

The solution is

$$l_s^i = l_w^i = \frac{1 - b_1 \bar{n}_i}{2}.$$

It follows immediately that the wage is decreasing in fertility.

$$w_i = al_s^i = \frac{a}{2}(1 - b_1 \bar{n}_i).$$

Note that the derived negative relationship is quite robust; that is, it does not depend on specific functional forms or parameter restrictions. The only crucial assumption is that it takes time to raise children.

One interpretation of this example is that people are ex ante identical, but are exposed to stochastic fertility shocks (e.g., birth control failures). Then, ex post, people will have different fertility realizations, which leads them to optimally invest different amounts into human capital. However, for such shocks to be the main driving force behind the negative fertility-income relationship, it would need to be the case that most people know their fertility realizations *before* they make their human capital accumulation decisions. While this seems implausible for schooling decisions, it is more plausible for human capital that is accumulated on the job through experience. Exogenous fertility shocks may also be important for some margins, such as drop-out decisions for girls who become pregnant in high school.

2.4.2 Endogenous Fertility and Endogenous Wages

Next, we extend the basic intuition given before to allow for both the choice of fertility and the endogenous determination of wages. Assume now that parents differ in their preferences for children; that is, some people value children more than others. To do this, we add a fertility choice to problem (5) and allow for preference heterogeneity. We also generalize the model along two other dimensions, which will turn out to be useful later on. First, following Ben-Porath (1976) and Heckman (1976), we allow for decreasing returns in the human capital accumulation process: $w = al_s^{\nu_s}$, $\nu_s \in (0, 1]$. Second, we allow for decreasing returns when working. That is, an individual working l_w units (hours/weeks/years) will earn a total income of $wl_w^{\nu_w}$, $\nu_w \in (0, 1]$. While this formulation is nonstandard (i.e., most of the literature assumes that income is linear in hours worked), we find it quite plausible since many jobs pay a premium for full-time work. Note also that setting $\nu_w = 1$ gives the standard model in which income is the product of an hourly wage and hours worked. The modified problem then is

$$
\begin{aligned}
(6) \qquad & \max_{c,n,l_w,l_s} \quad \alpha_c \frac{c^{1-\sigma}}{1-\sigma} + \alpha_n \frac{n^{1-\sigma}}{1-\sigma} \\
& \text{s. t.} \quad l_s + l_w \leq 1 - b_1 n \\
& \qquad\quad w = al_s^{\nu_s} \\
& \qquad\quad c \leq wl_w^{\nu_w}.
\end{aligned}
$$

The first-order conditions are:

$$l_s: \qquad \alpha c(al_s^{\nu_s}l_w^{\nu_w})^{-\sigma} a\nu_s l_s^{\nu_s-1} l_w^{\nu_w} = \alpha_n \left(\frac{1-l_s-l_w}{b_1}\right)^{-\sigma} \frac{1}{b_1}$$

$$l_w: \qquad \alpha c(al_s^{\nu_s}l_w^{\nu_w})^{-\sigma} a\nu_w l_s^{\nu_s} l_w^{\nu_w-1} = \alpha_n \left(\frac{1-l_s-l_w}{b_1}\right)^{-\sigma} \frac{1}{b_1}.$$

It follows immediately that $l_s = (\nu_s/\nu_w)l_w$. Using this, the optimal amount of work solves the following equation

$$\alpha_c a^{1-\sigma}\, \nu_s\left(\frac{\nu_s}{\nu_w}\right)^{\nu_s-1-\nu_s\sigma} l_w^{-(\nu_s+\nu_w)\sigma+\nu_s+\nu_w-1} = \alpha_n\left(\frac{1}{b_1}\right)^{1-\sigma}\left(1-\frac{\nu_s+\nu_w}{\nu_w}l_w\right)^{-\sigma}.$$

It is easy to derive closed form solutions for two special cases: (a) constant returns to scale ($\nu_w + \nu_s = 1$) and a general σ; and (b) general production function, but assuming log utility $\sigma = 1$.[30] The solution for case (b) is

$$l_w^* = \frac{\alpha_c\nu_w}{\alpha_n + (\nu_s + \nu_w)\alpha_c}$$

$$l_s^* = \frac{\alpha_c\nu_s}{\alpha_n + (\nu_s + \nu_w)\alpha_c}$$

$$n^* = \frac{1}{b_1}\left(\frac{\alpha_n}{\alpha_n + (\nu_s + \nu_w)\alpha_c}\right).$$

Note that the wage rate is

$$w^* = a(l_s^*)^{\nu_s},$$

which increases monotonically in time spent at school. Taking derivatives with respect to the child preference parameters, α_n, gives

$$\frac{\partial n^*}{\partial \alpha_n} = \frac{(\nu_s + \nu_w)\alpha_c}{b_1[\alpha_n + (\nu_s + \nu_w)\alpha_c]^2} > 0$$

$$\frac{\partial l_s^*}{\partial \alpha_n} = \frac{-\alpha_c\nu_s}{[\alpha_n + (\nu_s + \nu_w)\alpha_c]^2} < 0.$$

Thus, clearly, people who have a higher preference for children will have both—more children and a lower wage.

As can be seen from these expressions, fertility is independent of the raw learning ability, a. That is, without differences in preferences, parents will all have the same fertility.[31]

There are a couple of special cases where the implicit relationship between fertility and wages can be solved for explicitly.

In addition to $\sigma = 1$, now assume that $\nu_w = \nu_s = 1$: human capital is linear in years of schooling, and total income is simply the wage multiplied time spent working. For this case, we can substitute out all preference parameters to derive an equilibrium relationship between wage and fertility that will hold across all consumers (i.e., independent of their individual α_n and α_c):

$$n^* = \frac{1}{b_1}\left(1 - \frac{2}{a}w^*\right).$$

In this case, it follows that fertility is linearly decreasing in wages.

30. We analyze case (a) with dynastic altruism in the second section of the appendix.

31. Of course, if in addition one assumes that $\sigma < 1$, then fertility decreases in a for the same reasons as in section 2.3.2.

A second case that admits a straightforward closed form solution is when $\nu_s = \nu_w$. Then, the relationship can be written as:

$$n^* = \frac{1}{b_1}\left(1 - 2\left(\frac{w^*}{a}\right)^{1/\nu_s}\right).$$

In this case the relationship between the wage and fertility is nonlinear, with its curvature determined by the parameter ν_s.

In sum, this direction of causation generates the negative income-fertility and wage-fertility relationships under fairly general assumptions. In the second section of the appendix, we add parental altruism to this model. Similar results go through.

2.4.3 An Aside on Wages vs. Income

Here we have focused on the cross-sectional relationship between wages and fertility when the basic heterogeneity is differences, across people, in preferences for children vis-à-vis consumption goods. To do this we need a model in which wages themselves are endogenous. An alternative, weaker version of a similar property can be derived without explicitly including human capital formation in the model. This involves the relationship between fertility and income. For simplicity, assume that all households have the same w. Recall the solution to problem (3).

$$n^* = \frac{(y/w) + 1}{(\alpha_c b_1/\alpha_n)^{1/\sigma} w^{(1-\sigma)/\sigma} + b_1},$$

and consider two families that differ only in their values of α_n and/or α_c. As we can see, the family with the higher α_n will have more children for any value of σ and y. It also follows that this family will have lower earned income, $I = [1 - b_1 n^*(\alpha_n, \alpha_c)]w$, simply because it will spend more time raising children and less time working. Thus, preference heterogeneity of this type will also generate a negative correlation between fertility and earned income, without further assumptions on elasticities, or the formation of human capital, as long as children take parental time.

2.4.4 Empirical Evidence and Related Work

Empirical papers have confirmed the mechanism emerging from section 2.4.1 in the data, though most research (with the exception of Angrist and Evans [1998]) focuses on its importance for female wages, or income, and has little to say about the relationship between male income and fertility as shown in figure 2.1.[32] Similarly, the structural microeconomics literature, as well as some authors in the macroeconomics literature, also primarily focuses

32. Nor do they say much about most of the time period we are discussing, in which few women were earning market wages. In addition, good data for IV estimation (on twins, for example) has only become available recently.

on female wages. These papers address the mechanism emerging from section 2.4.2, though not in isolation. We review these results following.

Empirical Evidence

There is a large statistical literature that tries to assess the effect of (exogenous) fertility variation on labor supply, experience accumulation, and wages and/or earned income (see Browning [1992] for an early review). Mincer and Polachek (1974) find that work interruptions for childbearing have led to large human capital depreciations. Mincer and Ofek (1982) find that longer interruptions cause larger human capital losses. While there is a large and rapid increase in wages upon reentry, full earnings potential is not regained after interruption and reentry. These findings suggest that children have a lasting effect on income through forgone experience, which is a specific type of human capital accumulation.[33]

These papers view the number of children as exogenous. More recent research has focused on identifying valid instruments for fertility, such as miscarriages and unwanted pregnancies. For example, Miller (Forthcoming) finds that an exogenous *delay* in childbirth leads to a substantial increase in earnings, wage rates, and hours worked. She finds evidence for both fixed wage penalties and lower returns to experience for mothers. Since delay in fertility is typically associated with lower completed fertility, this result suggests that the number of children may have a strong effect on human capital accumulation of various types.

While all the papers mentioned so far focus on female earnings and leave father's and family income aside, Angrist and Evans (1998) use instrumental variable (IV) estimation to look at both parents' labor supply and labor income as well as family income. They look at families with two children and use the gender composition of the existing children as an instrument for the desire to have a third child. The authors find that families with a stronger desire for a third child work less and earn less. This is true for wives alone, husbands alone, and family income.[34] Unfortunately, nothing is said about hourly wages. Note that income is measured *before* the family actually has the third child. The fact that income is already lower prior to childbirth is in line with the aforementioned theory: people who want to have more children (i.e., higher α_n) anticipate working less in the future, and thus have a weaker incentive to accumulate human capital through experience.

33. Mincer and Polachek (1974) go on to answer the question: "Do family size and number of children currently present affect the accumulation of earning power beyond the effect on work experience? The answer is largely negative: when numbers of children and some measures of their age are added to work histories in the [regression] equations, the children variables are negative but usually not significant statistically" (S 95).

34. Their instrument is based on the following observation. Families with two children of the same sex are more likely to have a third child because sex mix is presumably preferred. Since gender of children is exogenous, the willingness to bear a third child—in the hope for the opposite sex—is also largely exogenous.

Related Theory

As for the mechanism in section 2.4.2 with endogenous fertility, the structural microeconomics literature on joint fertility and female labor supply choices also use preference heterogeneity to generate a distribution of fertility and wages as observed in the data. Again, the focus is on *female* labor supply, experience, schooling, and wages or earnings, while our mechanism is meant to address men (see figure 2.1) as well as women (see section 2.6 for details). Furthermore, permanent taste is typically not the only source of heterogeneity in these papers. Fixed and stochastic ability heterogeneity, as well as preference shocks over the life cycle, are additional necessary ingredients to fit the data. Francesconi (2002) estimates such a combined model with part-time and full-time employment. In a similar framework, Del Boca and Sauer (2009) analyze the effects of institutions on fertility, timing, and labor supply decisions. Finally, Keane and Wolpin (2006) add schooling and marriage decisions to estimate the effects of welfare programs on fertility and female labor supply.[35] All these papers use some version of the mechanism described here, though not in isolation. Our aim is to contrast pure taste and pure ability heterogeneity. In reality, of course, both may be relevant.

Finally, this mechanism is also sometimes used in the macroeconomics literature. For example, Erosa, Fuster, and Restuccia (Forthcoming) have stochastic fertility opportunities and stochastic values of children, together with learning-by-doing on the job, so that higher fertility translates into lower wages.[36] Again, male investment decisions are assumed not to be affected by fertility preferences and realizations. A similar mechanism is also at work in Erosa, Fuster, and Restuccia (2005) and Knowles (2007).[37]

2.4.5 Outlook

While the empirical evidence seems to support the idea that heterogeneity in tastes for children is to some extent responsible for the observed negative fertility-income relationship, this mechanism has received far less attention in the theoretical literature. Rather, most research starts with the assumption that exogenous differences in income (or ability) cause fertility to vary systematically across the income distribution. We therefore address

35. This literature is based on a combination of two basic models: Eckstein and Wolpin (1989), who analyze female labor force participation and experience accumulation with exogenous fertility heterogeneity, and Hotz and Miller (1988), who analyze contraceptive effort with taste heterogeneity, thereby endogenizing fertility but abstracting from labor supply and human capital accumulation of any kind.

36. Although, this is not the only channel through which fertility and income are related in their model.

37. Attanasio, Low, and Sanchez-Marcos (2008) analyze a similar model to Eckstein and Wolpin (1989) with exogenous fertility and endogenous experience to account for the increase in female labor force participation across cohorts.

the preference channel in all subsequent sections. Recall from section 2.4.3 that a simpler version of the mechanism can be used to derive a negative fertility-income relationship. For tractability, we use this shortcut when we analyze preference heterogeneity in sections 2.5 and 2.6. However, in all cases, the model can easily be extended to human capital accumulation and wages. We reintroduce endogenous wages in section 2.7, where we present an example in which parental time is not essential and in the appendix, where we build the dynastic analog of problem (6).

2.5 Quantity-Quality Theory

In this section, we revisit the idea that the demand for child quality naturally leads richer parents to want more quality and thus less quantity, what is often called the quantity-quality hypothesis.[38] This idea turns out not to be a very robust theory of the negative fertility-income hypothesis.

In his seminal work, Becker (1960) argued that there is a trade-off between quantity and quality of children. Originally, however, Becker did not propose the quantity-quality trade-off as an explanation for why fertility and income were negatively correlated. Indeed, in the 1960 paper Becker argues, by analogy with other durable goods, that economic theory suggests that fertility and income should be positively related, but perhaps only weakly so, while quality of children and income should be strongly positively correlated. The intuition for Becker's argument is simple. While richer parents do spend more on their children (better schools, better clothes, higher bequests, etc.), richer people spend more on everything. They have higher quality houses and cars as well, yet no one would argue that we should expect rich people to have fewer houses than poor people. As a first cut, the same logic should apply to children: richer people would want more quality, but probably not less quantity, the same way they also would not want better but fewer cars.

So what makes children different? Hotz, Klerman, and Willis (1993),

38. Empirical evidence about the quantity-quality trade-off is mixed (see Schultz [2005] for a useful summary). While the negative relationship between family size and various measures of child quality—in terms of investments or outcomes—is clearly negative, it is controversial whether this is a causal relationship. In particular, when using twin births as exogenous fertility variations, researchers have not always found a negative effect on these quality investments or outcomes of children. One regularity seems to carry through most studies, however: the negative relationship between number and quality of children is more strongly negative in developing countries (e.g., Rosenzweig and Wolpin (1980) for India; Li, Zhang, and Zhu (2008) for rural China) than it is in more advanced societies (e.g., Angrist, Lavy, and Schlosser (forthcoming) for Israel; Black, Devereux, and Salvanes (2005) for Norway). Yet Cáceres-Delpiano (2006) finds that a twin on a later birth reduces the likelihood that older children attend private school in U.S. Census data from 1980. One reason for the discrepancies between rich and poor countries might be the availability of high quality public schools in developed countries. For example, De la Croix and Doepke (2009) find that the effect of income on household choices, in terms of fertility and private schooling, diminishes as the quality of public schooling goes up.

reviewing Becker's arguments, seem to emphasize that what might be the case is that not children per se are normal goods, but that expenditures on children are: "If children are normal goods in the sense that total expenditures on children are an increasing function of income, then the sum of the income elasticities of the number and quality of children must be positive [. . .], but it is still possible that the income elasticity of demand for the number of children is negative [. . .] if the income elasticity of quality is large enough" (295). This is not our reading of the paper. Our reading is that, by analogy, quantity should be slightly increasing in income and quality should be greatly increasing in income. Becker's argument is, then, that the observation of a negative relationship is a missing variables problem, namely knowledge about contraceptives. Becker and Lewis (1973) and Becker and Tomes (1976) were important follow-ups on Becker (1960). Becker and Lewis (1973) argue that, once income is measured correctly, the true fertility-income elasticity is positive, even if the observed one is negative. Becker and Tomes (1976) argue that the quality production function has an endowment component that generates a negative correlation between fertility and income.

Following, we derive conditions under which simple examples including child quality can generate this negative correlation without making children inferior goods. We start with the simplest specification of the example in section 2.3 with log utility and a linear quality production function. In this example, it becomes apparent that even with quality choice and ability heterogeneity, we need a positive time cost and zero goods costs for fertility to be nonincreasing in income. Next, we derive the requirements on the quality production function for fertility to be strictly decreasing in wages—under both wage and taste heterogeneity. One example that generates the desired relation is an affine production function with a positive constant, as in Becker and Tomes (1976), together with the assumption that children take time while child quality requires purchased inputs as in Moav (2005). Various interpretations of this specification can be used to accommodate the cross section of fertility with respect to income and the trend in fertility over time. Finally, under preference heterogeneity, none of these requirements on the quality production function are needed.

2.5.1 A Simple Example

First, we show by example that including a quality choice in and of itself does not necessarily lead to a negative relationship between fertility and income. That is, including quality does *not* necessarily lead richer people to want fewer children. They might want more quality and accordingly, a smaller increase in number of children—as argued in Becker (1960)—but the relationship between fertility and income is still positive.

Suppose $U(c, n, q) = \alpha_c \log c + \alpha_n \log n + \alpha_q \log q$, $\alpha_q > 0$, $q = f(s) = s$ and $y = 0$. Then the problem from section 2.3 is:

$$\begin{aligned} \max_{c,n,q,s,l_w} \quad & \alpha_c \log c + \alpha_n \log n + \alpha_q \log q \\ \text{s. t.} \quad & l_w + b_1 n \leq 1 \\ & c + (b_0 + s)n \leq w l_w \\ & q \leq s. \end{aligned}$$

This is a version of the problem considered in Becker and Lewis (1973), while Becker (1960) assumed $b_0 = b_1 = 0$. The constraint set in this problem is not convex because of the term ns. We therefore rewrite the problem in terms of total quality, $Q = qn$.[39] We also know that the constraints hold with equality. Using this, the problem becomes:

$$\begin{aligned} \max_{c,n,Q} \quad & \alpha_c \log c + (\alpha_n - \alpha_q) \log n + \alpha_q \log Q \\ \text{s. t.} \quad & c + b_0 n + Q \leq w(1 - b_1 n). \end{aligned}$$

This is now a standard problem under the assumption that $\alpha_n > \alpha_q$. The solution is given by:

$$\begin{aligned} n^* &= \frac{\alpha_n - \alpha_q}{(\alpha_c + \alpha_n)(b_0 + b_1 w)} w \\ q^* &= \frac{\alpha_q (b_0 + b_1 w)}{\alpha_n - \alpha_q} \\ c^* &= \frac{\alpha_c}{\alpha_c + \alpha_n} w. \end{aligned}$$

Similar to what we found in the example in section 2.3.2, as long as the goods cost is positive ($b_0 > 0$), fertility is strictly increasing in the wage, w.[40] On the other hand, if $b_0 = 0$, fertility is independent of w, while earned income is $I = w(1 - b_1 n^*)$. Again, this does not give a negative relationship between income and fertility since there is no heterogeneity in fertility choice. Instead, we get an extreme version of Becker's original argument. That is, if there is only a time cost of children, $b_0 = 0$, then we have high income elasticity of quality per child (q is strictly increasing in w and hence I) and low income elasticity of number of children (n is independent of w or I).[41]

39. Rosenzweig and Wolpin (1980) write a model with $b_1 = 0$, but a children-independent price of quality. If this price is strictly positive, our formulation cannot be used.

40. Whether earned income, $I = (1 - b_1 n)w$, increases or decreases depends on the size of the increase in n in response to an increase in w. In the present example, we have:

$$\frac{dI}{dw} = (1 - b_1 n) - b_1 w \frac{dn}{dw} = \frac{(\alpha_c + \alpha_q)(b_0 + b_1 w)^2 + (\alpha_n - \alpha_q) b_0^2}{(\alpha_c + \alpha_n)(b_0 + b_1 w)^2} > 0.$$

Thus, in this case, income and fertility are positively related.

41. It is useful to note that the time intensity in the cost of children matters (the relative size of b_0 and b_1) for the size of these effects. Also, similarly to the cost of time theory, one could vary the elasticity of substitution in the utility function. We leave this part to the reader.

There are at least two ways in which this "negative result" can be overturned. First, keeping wage heterogeneity, the quality production function can be generalized. Second, one can consider preference heterogeneity instead of ability heterogeneity in this simple example. We consider these two avenues in turn following.[42]

2.5.2 The Quality Production Function

The next example is based on the analysis in Moav (2005), who argued that producing children takes time, while educating each child requires goods costs. This assumption makes quality relatively cheaper for higher wage people and one might expect a quantity-quality trade-off to result. However, the comparative advantage alone does not imply that higher wage people have fewer children, as we have seen before. The properties of the human capital production function are also a crucial ingredient, as noted in Moav (2005).

We make the same assumptions as before, except that we let $q = f(s)$ be unspecified for now. The maximization problem is given by:

$$\max_{c,n,q,s} \quad \alpha_c \log c + \alpha_n \log n + \alpha_q \log q \tag{7}$$

$$\text{s. t.} \quad c + b_0 n + sn \leq w(1 - b_1 n)$$

$$q = f(s).$$

The first order conditions give

$$\frac{sf'(s)}{f(s)} = \frac{\alpha_n}{\alpha_q}\left(\frac{s/w}{b_0/w + b_1 + s/w}\right) \tag{8}$$

$$n^* = \left(\frac{\alpha_n}{\alpha_c + \alpha_n}\right)\frac{1}{b_0/w + b_1 + s^*/w}. \tag{9}$$

Let the elasticity on the left-hand side of equation (8) be $\eta(s) \equiv sf'(s)/f(s)$.[43]

42. We have also explored a third channel—nonseparable preferences—to a limited degree (cf. Jones and Schoonbroodt, Forthcoming). For example, assume $q = s$ and solve:

$$\max_{\{c,n,q\}} \quad \alpha_c \log c + \log [[(\alpha_n - \alpha_q)n^\rho + \alpha_q(nq)^\rho]^{1/\rho}]$$

$$\text{s. t.} \quad c + (b_0 + b_1 w)n + nq \leq w.$$

In this case, if $\rho \in (0, 1)$ then n and $Q = nq$ are substitutes in utility and fertility is decreasing in w, while the opposite is true if $\rho < 0$. In the text, we are implicitly assuming the case where $\rho \to 0$. The substitutes case works because number of children is time intensive and hence more costly to high wage parents while the price of quality is the same across people. Another way of generating a negative income-fertility relationship through a quantity-quality trade-off is to assume that the educational choice is indivisible: the choice is between skilled and unskilled children. This mechanism was used in Doepke (2004). In this case, low ability people would choose (some) unskilled children and have more of them than high ability people who have skilled children. Among the latter group, however, fertility will be increasing in ability again.

43. Note that unless $f(s) = s^\lambda$ for some $\lambda > 0$, this formulation is very similar to the nonhomothetic preference example given in section 2.3 since we can rewrite the utility function as $\alpha_c \log c + \alpha_n \log n + \alpha_q \log f(s)$.

Ability Heterogeneity

Suppose that households differ in their abilities, w. In the case where $b_0 = 0$, we can see from equation (9) that for n^* to be a decreasing function in w, s^*/w needs to be increasing in w. But the right-hand side of (8) is increasing in this ratio. Thus, the left-hand side has to be increasing as well. Hence, we need that $\eta'(s) > 0$, which is purely a property of $f(s)$. An example of a human capital production function that satisfies this property was first introduced by Becker and Tomes (1976):[44]

$$f(s) = d_0 + d_1 s, \qquad d_0 > 0, d_1 > 0.$$

In this case, the solution is:

$$s^* = \frac{(\alpha_q/\alpha_n)b_1 w - d_0/d_1}{(1 - \alpha_q/\alpha_n)},$$

which is well-defined as long as $\alpha_q < \alpha_n$ and d_0 is small enough; that is, $d_0 < d_1(\alpha_q/\alpha_n)b_1 w$.[45] Solving for n^* gives

$$n^* = \frac{(\alpha_n - \alpha_q)/(\alpha_c + \alpha_n)}{b_1 - d_0/wd_1}.$$

From this it is clear that $\partial n^*/\partial w < 0$.

Finally, notice that this example still requires a time cost. In fact, in the case with $b_0 > 0$, the solution is given by:

$$s^* = \frac{(\alpha_q/\alpha_n)(b_0 + b_1 w) - d_0/d_1}{(1 - \alpha_q/\alpha_n)},$$

which is well-defined as long as

$$\text{(10)} \qquad \alpha_q < \alpha_n \qquad \text{and} \qquad \frac{\alpha_q}{\alpha_n}(b_0 + b_1 w) > \frac{d_0}{d_1}.$$

Solving for n^* gives

$$n^* = \frac{(\alpha_n - \alpha_q)/(\alpha_c + \alpha_n)}{b_1 + b_0/w - d_0/wd_1}.$$

Hence, fertility is decreasing in w if and only if

$$\text{(11)} \qquad \frac{d_0}{d_1} > b_0.$$

In the case where $b_1 = 0$, conditions (10) and (11) are mutually exclusive.

44. De la Croix and Doepke (2003, 2004) use a more complex production function that allows quality to depend on parental human capital, but overall has similar properties: $f(s, w) = d_1(d_0 + s)^{\gamma} w^{\tau}$, where $\gamma, \tau \in (0, 1)$ are parameters. Examples of production functions that do *not* satisfy the condition include $f(s) = s^a$ and $f(s) = as$, which lead to a constant s^*/w, and $f(s) = \log(s)$ and $f(s) = \exp(as)$, which lead to decreasing s^*/w.

45. Otherwise $s = 0$ is the solution.

Interpretation and Further Predictions of the Model

Becker and Tomes (1976) interpret d_0 as an endowment of child quality, or "innate ability." In this interpretation, one might want to take intergenerational persistence in ability into account. If the child's quality endowment and parent's ability, w, are positively correlated in the sense that $E(d_0) = w$, then fertility is, again, independent of w while quality is still increasing in w. An alternative would be that in those families in which parents have higher market wages, the marginal value of education is higher—d_1 is perfectly positively correlated with w. For example, assume that $d_1 = \kappa w$. Then even if innate ability, d_0, is perfectly correlated with w, fertility is still decreasing while education is increasing in w. This educational investment does not require time per se. Instead, for a given amount of goods, the high ability parent produces more quality.

An alternative interpretation of d_0 is publicly-provided schooling. Since this has increased over time, we see that the predicted response is that fertility will increase, at least holding w fixed. In contrast, holding d_0 fixed, an increase in income over time would cause fertility to decrease. Hence, under this interpretation the example suggests that the increase in income was more important than the increase in publicly-provided schooling.[46]

Preference Heterogeneity

Next, assume that w is the same for all households, but suppose that people differ in their preference for the consumption good, α_c. In all the previous examples, the more people like the consumption good, the fewer children they will have and, as long as $b_1 > 0$, the more income they will earn. However, the quality choice, q, is independent of α_c and hence income, I.

If, on the other hand, we consider heterogeneity in the preference for children, α_n, we see that the more people like children, n (relative to both consumption, c, and quality, q), the more they will have, the less income they will earn, and the less quality investments they make per child. Thus, in this case, fertility and income are still negatively related, while quality per child will be positively related with income.

Note that this does not depend on any particular assumption about goods costs or the quality production function. As usual, however, a positive time cost is required so that earned income, I, is decreasing in number of children, n, which generates the negative correlation.[47]

46. See the conclusion for suggestive simulations of such changes over time.

47. Pushing the idea of preference heterogeneity one step further, Galor and Moav (2002) argue that the forces of natural selection selected individual preferences that are culturally or genetically predisposed toward investment in child quality, bringing about a demographic transition.

2.6 Married Couples and the Female Time Allocation Hypothesis

A refinement of the price of time theory of fertility is to view the decision-making unit as a married couple and to explicitly distinguish between the time of the wife and the husband. In this version, since it is typically the case that most child care responsibility rests with the woman, it is the time of the wife that is critical to the fertility decision.[48] In its simplest form, the idea is that the price of children is higher for high productivity couples, even if only the husband works.[49]

The aim of this section is threefold. First, we test how robust the results derived in previous sections are to introducing women explicitly. In particular, we ask whether the same restrictions on parameters are necessary to generate a negative fertility-relationship when the division of labor within couples is taken into account. Second, we move to more general formulations that model home production explicitly, examining the restrictions needed on the home production technology under log utility (in the spirit of Willis [1973]). Third, we show that specific patterns of assortative mating are needed to match the data. A richer model also necessitates a more nuanced look at the data. The findings in the empirical literature can be summarized as the following three findings:

1. The correlation between fertility and wife's wage (or productivity). Evidence suggests that this correlation is strongly negative whether controlling for the husband's wage or not.
2. The conditional correlation between fertility and husband's wage, holding the wife's wage constant. Evidence here is very mixed (e.g., Blau and van der Klaauw [2007] find it is strongly positive, Jones and Tertilt [2008] find it is negative, and Schultz [1986] finds that it depends on the exact subgroup of the population one considers; see following).
3. The unconditional correlation between fertility and husband's wage. Evidence suggests that this correlation is strongly negative in the data.

48. A related idea was first formalized in Willis (1973), who studied the time allocation problem for a couple in which the time of both the husband and wife are used in raising children while consumption is produced using the time of the wife and market-purchased goods.

49. In the words of Hotz, Klerman, and Willis (1993): "A second major reason for a negative relationship between income and fertility, in addition to quality-quantity interaction, is the hypothesis that higher income is associated with a higher cost of female time, either because of increased female wage rates or because higher household income raises the value of female time in nonmarket activities. Given the assumption that childrearing is a relatively time intensive activity, especially for mothers, the opportunity cost of children tends to increase relative to other sources of satisfaction not related to children, leading to a substitution effect against children. As noted earlier, the cost of time hypothesis was first advanced by Mincer (1963) and, following Becker's (1965) development of the household production model, the relationship between fertility and female labor supply has become a standard feature of models of household behavior" (298–99).

We show that simple examples imply that fertility should be decreasing in the productivity or wage of the wife (1) and (weakly) increasing in the wage of the husband (2). Because of this theoretical result, much of the empirical literature has taken the stand that the negative estimated correlation between income of the husband and fertility (3) is contaminated by a missing variables problem—the productivity of the wife. Since productivities or wages within couples are typically positively correlated, a downward bias (perhaps enough to change the sign) is induced on the true effect of husband's income on fertility. One might think that this effect is large enough, in theory, that any restrictions on the form of preferences, and so forth, are no longer necessary. This is not what we find in the following examples. Rather, we find that specific assumptions on elasticity, the home production function, and assortative mating (either in terms of productivities or preferences) are still required to generate facts (1) and (3).[50] We summarize those combinations of assumptions that successfully generate facts (1) and (3) in table 2A.1 in the appendix.

2.6.1 Empirical Findings

Testing predictions (1) and (2) in the data is complicated because of the difficulty in obtaining direct measures of the value of the wife's time. Until recently many wives did not work and even now, those that do are a "selected" sample. Hence, other proxies must be used, such as inferred productivities based on a Mincer regression or education. The evidence on (1) and (3) are quite robust while evidence on (2) is mixed. Following is a summary of the findings of three recent studies.

Schultz (1986) estimates a reduced-form fertility equation based on his household demand framework:[51]

$$n_i = \beta_0 + \beta_1 \ln w_{fi} + \beta_2 w_{mi} + \beta y_i + \varepsilon_i,$$

where n is the number of children, w_f and w_m are female and male wages, respectively, y is asset income, and ε is an error term. This equation is estimated separately for different age and race groups. The data are from the 1967 Survey of Economic Opportunities, an augmented version of the Current Population Survey. He finds that

> [I]n every age and race regression the wife's wage is negatively associated with fertility. The coefficient on the husband's predicted wages changes sign over the life cycle, adding to the number of children ever born for

50. Given the mixed evidence on fact (2), we do not focus too much on the model prediction for fact (2).

51. Schultz (1986, 91) also says: "Empirical studies of fertility that have sought to estimate the distinctive effects of the wage opportunities for men and women generally find β_1 to be negative, while β_2 tends to be negative in high-income urban populations and frequently positive in low-income agricultural populations (Schultz (1981))."

> younger wives [. . .] but contributing to lower fertility among older wives. [. . .] For white wives over age 35 and for black wives aged 35–54, a higher predicted husband's wage is significantly associated with lower completed fertility. The elasticities of fertility with respect to the wage rates of wives and husbands are of similar magnitude for blacks and whites, although for blacks the level of fertility is higher and wage levels are lower. [. . .] These estimates give credence to the hypothesis that children are time-intensive. In all age and race regressions the sum of the coefficients on the wife's and husband's wage rates is negative and increases generally for older age groups. [. . .] The hypothesis that children are more female than male time-intensive is also consistent with these estimates. (Table 1, 93)

Using National Longitudinal Survey of Youth (NLSY) longitudinal data for women born between 1957 and 1964, Blau and van der Klaauw (2007) find that

> [A] one standard deviation increase in the male wage rate is estimated to have some fairly large effects on white women, but none of the underlying coefficient estimates are significantly different from zero. Several of the black and Hispanic interactions are statistically significant, however, and the simulated effects are in some cases quite large. A higher male wage rate increases the number of children ever born to black women by 0.169. . . . For Hispanic women, a higher male wage rate [also] increases fertility. . . . [A] higher female wage rate generally has effects that are of the opposite sign from those of the male wage rate. As with the male wage rate, the effects are not significantly different from zero for whites, but for blacks and Hispanics a higher female wage rate has negative effects on fertility that are significantly different from zero. Children ever born decline by about 0.1 for blacks and Hispanics. (29–30)

Jones and Tertilt (2008) also experiment with this hypothesis. Since very few women worked in the early cohorts, education is chosen as a measure of potential income. They find that children ever born (CEB) is declining in both the education level of the wife and the husband, and significantly so. Moreover, the coefficients on husband's and wife's education are similar in size (the wife's being slightly larger) and there is no systematic time trend.

2.6.2 Theory

It is convenient to break this variant of the story into two separate parts: one in which the woman does not work in the market, and one in which she can and does. Roughly, we can think of the first version as corresponding to a time in history when very few married women participated in the formal labor market. The second corresponds to more recent history. It is clear that the critical features necessary to reproduce the observations must be different in the two cases. We summarize all models that are consistent with the facts in table 2A.1 in the appendix.

Full Specialization in the Household

In this example, the husband works in the market, l_m, earning wage, w_m, or enjoys leisure, ℓ_m, while the wife works only in the home, l_{hf}, so that her trade-off is between how much time to allocate to producing home goods versus raising children, $b_1 n$, or enjoying leisure, ℓ_f. Her productivity in home production is denoted w_f. This setup may be more relevant to the early period in the data when (married) women's labor force participation was roughly zero.

The gender-specific utility function is given by

$$U_g = \alpha_{cg}\log(c_g) + \alpha_{ng}\log(n) + \alpha_{\ell g}\log(\ell_g) + \alpha_{hg}\log(c_{hg}),$$

where $g = f, m$ indicates gender, c_g is market consumption, n is the number of children, ℓ is leisure, and c_{hg} is the home good. Note that only the husband's leisure is needed for some of the following results. That is, $\alpha_{\ell f}$ could be zero, while the husband needs an alternative use of time to generate any endogenous wage/income heterogeneity for the husband. Given our previous results, we assume that children cost only time (i.e., $b_0 = 0$).

We assume that there is unitary decision making in the household. The family solves the problem:

$$(12) \qquad \max_{\{c_m, c_f, c_{hm}, c_{hf}, n, \ell_m, \ell_f, l_m, l_{hf}\}} \quad \lambda_f U_f + \lambda_m U_m$$

$$\text{s. t.} \quad \begin{aligned} & c_f + c_m \le w_m l_m \\ & l_m + \ell_m \le 1 \\ & c_{hf} + c_{hm} \le w_f l_{hf} \\ & l_{hf} + \ell_f + b_1 n \le 1. \end{aligned}$$

Here ℓ_f and ℓ_m are leisure of the female and male respectively, w_m is the wage of the man, w_f is the productivity of the woman in home production, and c_{hf} and c_{hm} are consumption of home goods by the woman and the man, respectively. Note that it is assumed that the wife spends b_1 hours for each child being raised (and the husband spends none). To keep it simple, assume perfect agreement of couples: assume $\alpha_{xf} = \alpha_{xm} = \alpha_x$ for $x = c, h, n, \ell$. Further, without loss of generality, assume $\lambda_f + \lambda_m = 1$ and $\alpha_c + \alpha_n + \alpha_\ell + \alpha_h = 1$.

This problem separates into two maximization problems, one concerning the allocation of the man's time and one concerning the allocation of the woman's time. The one for the man is straightforward and does not involve fertility. Notice however, that male earnings are increasing in α_c since leisure becomes less desirable relative to consumption. The problem for the woman's time allocation is:

$$\max_{\{c_{hm},c_{hf},n,\ell_f\}} \quad \lambda_f\alpha_\ell \log(\ell_f) + \lambda_f\alpha_h \log(c_{hf}) + \lambda_m\alpha_h \log(c_{hm}) + (\lambda_f + \lambda_m)\alpha_n \log(n)$$

$$\text{s. t.} \quad b_1 w_f n + c_{hf} + c_{hm} + \ell_f \le w_f.$$

The solution is:

$$n^* = \frac{\alpha_n}{\lambda_f\alpha_\ell + \alpha_h + \alpha_n}\frac{1}{b_1}. \tag{13}$$

Ability Heterogeneity, Elasticity, and the Home Production Function Suppose households differ in their productivities, (w_f, w_m). We see that n^* is independent of woman's productivity in the home. If education is a good proxy for female home productivity, then the evidence in Jones and Tertilt (2008) contradicts this model implication. That is, this model is not consistent with fact (1).[52] Fertility is also independent of w_m, holding w_f fixed. Finally, even if the productivity of the husband and wife are positively correlated (or independent), fertility is independent of both productivities. Thus, fact (3) is not predicted here either.[53] Clearly, something is missing in the theory.

As can be seen from the previous, since the couple's problem splits into two separate maximization problems, and the one for the wife's time looks just like those discussed in section 2.3 (additional goods permitting), the natural next step is to analyze a more general version in which utility is given by:

$$U_g = \alpha_c \frac{c_g^{1-\sigma}}{1-\sigma} + \alpha_n \frac{n^{1-\sigma}}{1-\sigma} + \alpha_\ell \frac{\ell_g^{1-\sigma}}{1-\sigma} + \alpha_h \frac{c_{hg}^{1-\sigma}}{1-\sigma}.$$

With $\sigma < 1$, it follows that n^* will be decreasing in the productivity at home of the wife, w_f, fact (1). Holding the wife's productivity fixed, fertility is still independent of the husband's wage—fact (2). Thus, if w_f and w_m are positively correlated, and $\sigma < 1$, the partial correlation between n^* and w_m is negative as well—fact (3). This example is summarized in the first row of table 2A.1.

A second variation that also reproduces the negative correlation in the cross section can be obtained by making the home production technology slightly more complex. Assume that utility is given by

$$U_g = \alpha_n \log(n) + \alpha_h \log(c_{hg}),$$

where the home good, c_{hg}, is produced using market goods, c, and time of the wife, l_{hf}, with productivity w_f; that is, $c_{hf} + c_{hm} = F(c, w_f l_{hf})$. To simplify

52. One should note that though fact (1) is based on evidence from the twentieth century, so a model where fertility is constant across women, conditional on husband's income, could still be a good description of the nineteenth century.

53. It can also be shown that if children have a nonmarket goods cost, $b_0 > 0$, n^* is increasing in w_f. It follows that if w_f is positively correlated with w_m (which is what we might expect), n^* and w_m will also be positively correlated.

the analysis, we now assume that leisure is not valued, $\alpha_{\ell_g} = 0$. Thus the problem is:

$$\max_{\{c_m, c_f, c_{hm}, c_{hf}, n, \ell_m, l_{hf}\}} \quad \lambda_f U_f + \lambda_m U_m$$
$$\text{s. t.} \quad c \leq w_m$$
$$b_1 n + l_{hf} \leq 1$$
$$c_{hf} + c_{hm} \leq F(c, w_f l_{hf}).$$

The first-order conditions can be reduced to one equation involving the amount of time the wife spends making home goods, which directly relates to fertility:

$$(1 - l_{hf}) = \frac{\alpha_n}{\alpha_h} \frac{1}{w_f} \frac{F(w_m, w_f l_{hf})}{F_2(w_m, w_f l_{hf})}$$
$$n^* = \frac{1 - l_{hf}}{b_1}.$$

That is, time spent in child-rearing $(1 - l_{hf})$ is positively related to the relative desirability of children to consumption, α_n/α_h, and negatively related to the productivity of the wife, w_f, all else equal. Thus, so is fertility, n^*. When F is assumed to be CES, $F(c, w_f l_{hf}) = [\delta c^{\rho} + (1 - \delta)(w_f l_{hf})^{\rho}]^{1/\rho}$, this becomes:

$$(14) \qquad n^* = \frac{1 - l_{hf}}{b_1} = \frac{\alpha_n}{\alpha_h}(b_1(1 - \delta))^{-1}\left[\delta\left[\frac{w_m}{w_f}\right]^{\rho} l_{hf}^{1-\rho} + (1 - \delta) l_{hf}\right].$$

We can see from the second equality that in the Cobb-Douglas case $(\rho \to 0)$, $(1 - l_{hf})$ is independent of both w_m and w_f, but does depend on α_n/α_h. Thus, the same must be true of n^* (first equality).

We can also see that for any value of ρ, if w_f and w_m are proportional $(w_f = \phi w_m)$, then l_{hf} is independent of w_m and w_f and hence the same is true for fertility. That is, under perfect assortative mating, fertility and the wage of the husband and the productivity of the wife are independent.

When this correlation is imperfect and $\rho \neq 0$, the analysis is more complicated. We will assume another extreme, that w_m and w_f are independent, in what follows. When $\rho > 0$, market goods and female time are substitutes in the production of consumption. An increase in w_m holding w_f fixed causes l_{hf} to fall. Hence, n^* rises in this case. That is, fertility is an increasing function of husband's wage if w_m and w_f are independent.

On the other hand, when $\rho < 0$, market goods and female time are complements in the production of consumption. An increase in w_m holding w_f fixed causes l_{hf} to rise. Hence, n^* falls in this case. That is, fertility is a decreasing function of husband's wage if time and goods are complements and wages of husbands and wives are independent.

Thus, assuming enough complementarity between time and goods in production, F, and enough independence between productivities of husbands and wives, also gives a model that can reproduce the negative correlation between husbands income and fertility—fact (3). From equation (14) it is also obvious that female and male productivities enter in the opposite ways. Thus, if $\rho < 0$, it follows immediately that a higher female home productivity leads to lower optimal fertility. Of course, home productivity is difficult to measure, and hence, it is not obvious that this implication is counterfactual. Alternatively, assume $w_f = \overline{w}$, that is, women are homogenous in their home productivity (e.g., perhaps because more schooling does not increase productivity in cooking, cleaning, etc.). Then, we still generate fact (3), while the model has nothing to say about women. But again, given that home productivity is difficult to assess empirically, this may well be in line with the facts. This result is summarized as row 2 in table 2A.1.

In sum then, we see that fertility and wages/home productivities are uncorrelated without the same kinds of assumptions over utility function curvature that we have identified in earlier sections. As a substitute, we can generate the observed curvature, even with unitary elasticity in preferences, if we move away from unitary elasticity in the home production technology. But this requires the right correlation between husband's wages and wife's productivity in the home.

Preference Heterogeneity Now assume there is heterogeneity in tastes rather than productivities; that is, households differ in how much they like children, α_n, consumption, α_c, and/or the home good, α_h. Going back to problem (12), the comparative statics of fertility with respect to preference parameters can immediately be derived from equation (13). Similarly, one can solve for labor earnings. Note that since the woman does not work in the market in this version, total household earnings are equal to male earnings and are given by:

$$I_m = w_m(1 - \ell_m^*) = w_m\left(1 - \frac{\alpha_\ell}{\alpha_c[\lambda_f/\lambda_m + 1] + \alpha_\ell}\right).$$

The results are as follows:[54]

1. With heterogeneity in α_c alone, while (male) earnings are increasing in α_c, fertility is the same for all households.
2. With heterogeneity in α_n or α_h alone, (male) earnings are the same for all households while fertility is decreasing in α_h and increasing in α_n.
3. With simultaneous heterogeneity in α_c and α_h and a positive correlation of these preferences within households, fertility will be negatively correlated with husband's earnings, fact (3). This finding hinges on the husband having

54. Using a model along the lines of section 2.4, these findings can be generalized to apply to male wages instead of labor earnings.

an alternative use of time to market work—leisure, in this case. This case is row 3 in table 2A.1.

In sum, only the third case (heterogeneity in tastes for all consumption goods, and positive correlation of these tastes within couples) can generate the negative income-fertility relationship observed for men. Similar results can be derived in the examples with general elasticities or home production functions.

Partial Specialization

To capture better the realities of the twentieth century, we now allow for more gender symmetry. Women and men both work in the market and there is no home production. We still assume that only women can raise children. Also, as before, we add leisure, ℓ_g. Then, husbands have to allocate their time between work and leisure, while women's time is allocated between three activities: working, enjoying leisure, and child-rearing. This example might be more relevant for the more recent experience, when women's labor force participation has been relatively large.

The gender-specific utility function is given by:

$$U_g = \alpha_{cg} \log(c_g) + \alpha_{ng} \log(n) + \alpha_{\ell g} \log(\ell_g),$$

and the couple solves the problem:

$$\max_{\{c_m, c_f, n, \ell_m, \ell_f\}} \quad \lambda_f U_f + \lambda_m U_m$$

$$\text{s. t.} \quad c_f + c_m \leq w_m(1 - \ell_m) + w_f(1 - \ell_f - b_1 n),$$

where ℓ_f and ℓ_m are leisure of the female and male, respectively, and w_f and w_m are the respective wages. Each child takes b_1 units of female time. Without loss of generality, assume that $\lambda_f + \lambda_m = 1$ and $\alpha_c + \alpha_n + \alpha_\ell = 1$. Define $W = w_f + w_m$ as total wealth.

Given the assumption of logarithmic utility, we obtain the standard result that expenditure on each good is a constant fraction of wealth, given by preferences:

$$c_f = \lambda_f \alpha_c W;$$

$$c_m = \lambda_m \alpha_c W;$$

$$w_f \ell_f = \lambda_f \alpha_\ell W;$$

$$w_m \ell_m = \lambda_m \alpha_\ell W;$$

$$b_1 w_f n = (\lambda_m + \lambda_f) \alpha_n W.$$

This immediately implies that:

$$n^* = \frac{(\lambda_m + \lambda_f)\alpha_n}{b_1} \left[1 + \frac{w_m}{w_f} \right]. \tag{15}$$

Comparing equation (15) to the full specialization analogue (13), one can see that the main difference is that the male wage and the husband's weight affect optimal fertility in the partial specialization versions, but not when full specialization is assumed. With partial specialization, the time allocation of husband and wife is more interdependent since they can, to some extent, substitute tasks between them. This is technologically infeasible in the full specialization model and hence, male wages are irrelevant for fertility choices.

Ability Heterogeneity Suppose households differ in their market wages, w_f and w_m. We see that fertility, n^*, is decreasing in the wife's wage, w_f, if the husband's wage, w_m, is held constant. Further, fertility, n^*, is increasing in the husband's wage, w_m, if the wife's wage, w_f, is held constant.

Thus, this model is consistent with fact (1) and in line with some authors' findings on fact (2) (e.g., Blau and van der Klaauw 2007). What remains to be seen are conditions under which fact (3)—that is, the negative correlation between male wages and fertility—can be accommodated as well. From equation (15), we also see that:

$$E\,[n|w_m] = \frac{(\lambda_m + \lambda_f)\alpha_n}{b_1}\left[1 + w_m E\left[\frac{1}{w_f}|w_m\right]\right].$$

Thus, the partial correlation between fertility and husband's income depends on $E[1/w_f|w_m]$. That is, it depends on the correlation between husband's and wife's market wages. Depending on the matching pattern, we can distinguish three cases:

1. Perfectly (positively) correlated wages within couples:

(a) If $w_f = \phi w_m$, then $E[1/w_f|w_m] = 1/\phi w_m$, and so n^* is independent of w_m.

(b) Similarly, if $w_f = \phi w_m^{\nu}$, then $w_m E[1/w_f|w_m]$ is increasing (decreasing) in w_m if $\nu < 1$ ($\nu > 1$). That is, n^* is increasing in w_m for $\nu < 1$ and decreasing in w_m for $\nu > 1$. Note that $\nu > 1$ means that a 1 percent increase in the husband's wage is associated with a more than 1 percent increase in the productivity of his wife.

(c) More generally, assuming matching can be characterized by a deterministic function $w_f(w_m)$, then n^* is decreasing in w_m if and only if $[w'_f(w_m)]/(w_f/w_m) > 1$. In words, the elasticity of female wages with respect to male wages must be larger than one. This seems unlikely. This case is summarized in row 4 in table 2A.1.

2. Independent wages within couples:

Then $E[1/w_f|w_m] = E[1/w_f]$, and so n^* is increasing as a function of w_m.

3. Negatively correlated wages within couples:

Suppose that $w_f = D - \nu w_m$ (where $D > 0$ so that $w_f > 0$). In this case $w_m E[1/w_f|w_m] = w_m/(D - \nu w_m) = 1/(D/w_m - \nu)$. Again this is increasing in w_m.

Thus, this version of the theory is consistent with fact (1)—that the regression coefficient on wife's wage is positive—and with the "debated fact (2)" that the regression coefficient on husband's income is positive (as in Blau and van der Klaauw [2007]). But this version is not consistent with a negative partial correlation between husband's income and fertility (unless the correlation is positive with $\nu > 1$, which seems unlikely). Thus, simply considering couples does not remove the need for special assumptions about the curvature on utility as in the previous simpler examples.

Preference Heterogeneity From equation (15), we can also see the relationship between income and fertility when the basic source of heterogeneity is in preferences. For example, if couples differ in their values of α_c and assuming both α_ℓ and α_n are lower so that $\alpha_c + \alpha_\ell + \alpha_n = 1$ for all households, those with higher desire for consumption choose lower leisure (both ℓ_f and ℓ_m), and also lower fertility, n^*. Because of this, those couples with higher α_c will have both higher incomes, since they work more, and lower fertility (row 5, table 2A.1). Note that we have assumed that couples are matched perfectly in terms of their preferences.

2.7 Nannies

So far, the assumption that children take time has been an essential ingredient for deriving a negative wage-fertility relationship. It is easy to see that with goods costs only, none of the previous examples work. That is, with $b_0 > 0$ and $b_1 = 0$, the negative wage-fertility relationship gets reversed in any of the (working) examples of sections 2.3, 2.4, 2.5, and 2.6.

While it is fairly obvious that children are time-intensive, it is less clear that it is specifically the parent's time that is needed. In fact, outsourcing child care is quite common, and has been throughout history. Examples include nannies, au pairs, relatives, wet nurses, and even orphanages.[55] In short, these kind of arrangements mean that even though children take time to raise, this time, in principle, can be hired. Hence, it is not clear why the price of children should be higher for high wage people.

In this section we first show how, when buying nanny-time is an option, higher wage parents will choose to have *more* children in simple models. We then ask what assumptions would restore the negative wage-fertility relationship, *even when* hiring nannies is possible. We give one example where a specific type of preference heterogeneity gives the desired result.

55. In the nineteenth century, many poor children were sent to orphanages, even when the parents were still alive, but too poor to feed the children. In 1853, Charles Loring Brace founded the Children's Aid Society, which rescued more than 150,000 abandoned, abused, and orphaned children from the streets of New York City and took them by train to start new lives with families on farms across the country between 1853 and 1929.

2.7.1 An Example with Ability Heterogeneity

To see that the assumption of *parental* time is a critical one, consider the following simple example:

$$\max_{c,n,\gamma} \quad \alpha_c u(c) + \alpha_n u(n)$$
$$\text{s.t.} \quad c + w_n(1 - \gamma)b_1 n \leq w(1 - \gamma b_1 n),$$

where $b_1 n$ is the total time requirement for raising n children, as before, but the time cost of children can now be split into parental time, $\gamma b_1 n$, or nanny time, $(1 - \gamma)b_1 n$, where $\gamma \in [0, 1]$. We denote the cost of a nanny by w_n per unit of time.

The optimal use of nannies in this example depends on the relative market wage of nannies versus parents. As long as $w < w_n$, it is never optimal to hire a nanny ($\gamma^* = 1$), and hence, this case is analog to our previous analysis of examples in which children require parental time. On the other hand, when $w > w_n$, parents prefer to hire a nanny, so that $\gamma^* = 0$. This case is equivalent to examples where children are a goods cost only, and there we have seen that $dn^*/dw \geq 0$. So while in this example $dn^*/dw < 0$ is possible, it occurs only in the region where nannies are irrelevant.

Thus, if some people have market wages that are lower than wages of nannies and others have higher wages, this model implies a v-shaped wage-fertility relationship. That is, fertility is downward sloping in wages for people with wages below the nanny wage and upward sloping thereafter. Recall from figure 2.1 however, that the data do not display such a v-shaped relationship.[56]

Going one step further, one may ask: what determines the nannies' wage? Notice that in this model, everyone is equally productive at child care. One unit of time produces $(1/b)$ children. Since this is the case, everyone with a market ability, w, below the nannies' wage would be better off becoming a nanny and raising $(1/b)$ children since leisure is not valued. Everyone with ability above the nannies' wage would hire a nanny. The nannies' wage is then determined through demand and supply and w_n should be the lowest wage observed in the data. That is, we would observe an increasing relationship between wages and fertility throughout the income ladder.

One might rephrase the question as follows: why is fertility decreasing in wages even for those people whose (after-tax) wages are higher than the hourly cost of day care or nannies?

56. Some authors have argued that at the very top of the income distribution, the fertility-income relation might be positive. Due to top coding and small samples at the top of the income distributions, these estimates are often statistically insignificant. Also, if this theory were applied to such a v-shape, it would mean that nannies are so expensive (either due to high wages or high tax wedges) that only the top income group finds it worthwhile hiring nannies. This seems to be at odds with the evidence as well.

There are, of course, several plausible answers to this question, such as the moral hazard problem involved in child care. Even though, in principle, nannies can be hired, if there is some effort involved in raising a high quality child, then the incentives for a nanny might be different from those of a parent. If monitoring is costly, parents might optimally choose to do the child-rearing themselves. In this case, the opportunity cost of a child again is increasing in income. Alternatively, perhaps parents *enjoy spending time* with their children over and above the pure utility effect of having children. If people derive pleasure from, say, spending the weekend with their children, then nannies are a poor substitute for own child-rearing. To the best of our knowledge, these ideas have not been formalized seriously, yet.[57] Also, not everyone is equally productive in raising children; in particular, if nannies are also teachers. While we believe these are interesting and potentially promising channels, they are well beyond the scope of this chapter, and are left for future research. In the next subsection, we pursue yet another possibility, based on preference heterogeneity and endogenous wages along the lines of section 2.4.

2.7.2 A Working Example with Preference Heterogeneity

The idea is that people differ in how much they like "material goods" goods vis-à-vis nonmaterial goods such as children and leisure. That is, some people like a "market-consumption lifestyle" while others like a "family-leisure lifestyle." Because of these different preferences, the former invest more in human capital and therefore have a higher wage, while the latter know they will enjoy leisure, which makes human capital investments less profitable. These are also the people who like large families. As we will see in the next example, one can recover the negative wage-fertility relationship in this setup, even allowing for nannies. However, the result rests on a particular form of preference heterogeneity across households. Therefore, rather than seeing this example as a definite answer to the question raised at the beginning of this section, we view it as a starting point for discussion and further research.

The starting point here is the example of section 2.4, where parents make schooling choices for themselves, which in turn determine their wage. To keep it simple, assume $\nu_s = \nu_w = 1$. We add one additional good to the utility function: leisure, ℓ. As before, each child requires a time input, b_1. Again, this can be a nanny's time, $(1 - \gamma)b_1 n$, or the parent's time, $\gamma b_1 n$, (where $\gamma \in [0, 1]$). In this choice, the parent takes the nanny's wage, w_n, as given.

The choice problem is:

57. Erosa, Fuster, and Restuccia (Forthcoming) have an indirect way of modeling the idea that parents like to spend time with children. That is, the value of staying at home can only be enjoyed if the mother gave birth in the past but has not returned to work since.

$$\max_{c,n,\ell,l_s,l_w,\gamma} \quad \alpha_c \log(c) + \alpha_n \log(n) + \alpha_\ell \log(\ell)$$
$$\text{s.t.} \quad l_s + l_w + \ell + \gamma b_1 n \leq 1$$
$$w = al_s$$
$$c + w_n(1-\gamma)b_1 n \leq wl_w.$$

It is easy to see that $l_s^* = l_w^*$. That is, given the child care choice, γ, and the leisure choice, ℓ, this maximizes market income. In terms of the nanny choice, one can show that an interior choice is never optimal. We therefore solve the problem for $\gamma = 1$ and $\gamma = 0$ and show that, assuming people differ in preferences, fertility and wages are negatively related for both $\gamma = 1$ and $\gamma = 0$.[58] Finally, we compare utilities across the two choices and derive the condition on parameters for which parents optimally hire a nanny.

Suppose the parent cares for the child, $\gamma = 1$. Then the solution is given by:

$$l_s^* = \frac{\alpha_c}{\alpha_\ell + \alpha_n + 2\alpha_c}$$
$$n^* = \frac{\alpha_n}{(\alpha_\ell + \alpha_n + 2\alpha_c)b_1}$$
$$\ell^* = \frac{\alpha_\ell}{\alpha_\ell + \alpha_n + 2\alpha_c}.$$

This is very similar to the solution in section 2.4, except that leisure is an additional choice variable. All the results go through. In particular, if parents take care of their children themselves, those who like the consumption good more; that is, higher α_c relative to α_n and α_l, will invest more in human capital, l_s, and hence have higher wages, $w = al_s$. They will also choose fewer children and less leisure.

In the case where parents choose to outsource child care, $\gamma = 0$, the solution is given by:

$$l_s^* = \frac{\alpha_c + \alpha_n}{\alpha_\ell + 2(\alpha_n + \alpha_c)}$$
$$n^* = \frac{\alpha_n(\alpha_c + \alpha_n)}{[\alpha_\ell + 2(\alpha_c + \alpha_n)]^2} \frac{a}{w_n b_1}$$
$$\ell^* = \frac{\alpha_\ell}{\alpha_\ell + 2(\alpha_c + \alpha_n)}.$$

Again, suppose that people differ in their preference for the consumption good α_c. Then, time in school, and hence wages, are strictly increasing in α_c

58. Formally, when $\gamma = 0$, the problem reduces to a pure goods cost example with $b_0 \equiv w_n b_1$.

and fertility is strictly decreasing in α_c as long as leisure is not too important (the exact condition is: $2(\alpha_c + \alpha_n) > \alpha_\ell$). Hence, we obtain the negative fertility-wage relationship even if nannies are hired.

Finally, the condition for using a nanny is given by:

$$U|_{\gamma=0} > U|_{\gamma=1}$$

iff

$$\frac{a}{w_n} > \left[\frac{\alpha_c^{\alpha_c}(\alpha_\ell + 2(\alpha_n + \alpha_c))^{(\alpha_\ell+2(\alpha_n+\alpha_c))}}{(\alpha_\ell + \alpha_n + 2\alpha_c)^{(\alpha_\ell+\alpha_n+2\alpha_c)}(\alpha_n + \alpha_c)^{(\alpha_n+\alpha_c)}}\right]^{1/\alpha_n}.$$

The higher one's ability, a, relative to nanny wages, w_n, the more likely it is that the parent will hire a nanny. This is similar to the logic in the previous example with the v-shaped (or increasing) fertility-wage relationship. What is different here is that, assuming households differ in α_c, fertility and wages will be negatively related even among those parents who do use nannies; that is, those who choose a goods cost rather than a time cost.

Figure 2.2 illustrates the model graphically. In this example, all households have the same ability, a, but differ in their preferences, α_c. The figure then plots optimal choices as a function of α_c, both conditional on using a nanny

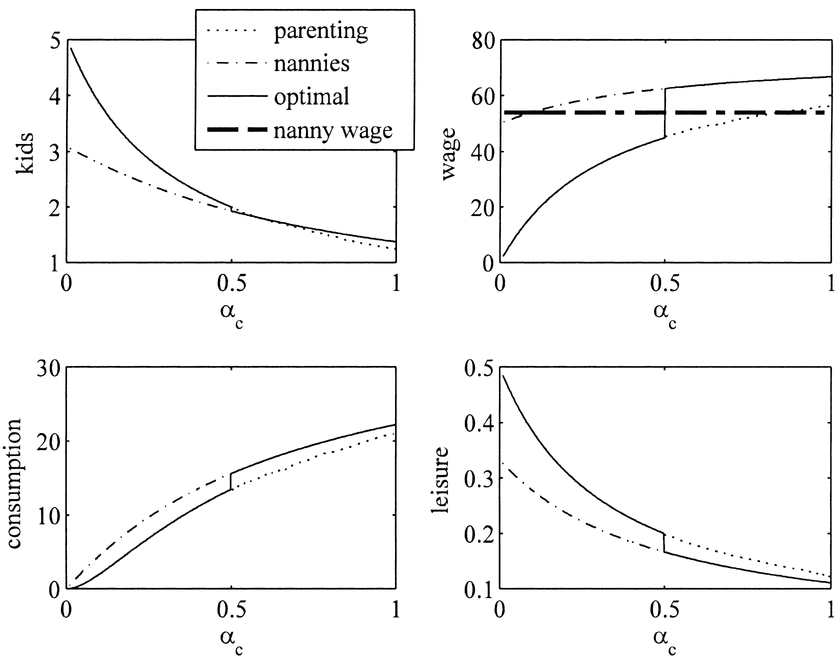

Fig. 2.2 Example with nanny choice

or parenting one's own child. The solid line depicts the solution under the optimal nanny choice. The figure shows clearly how fertility decreases and wages increase in the desire to consume (α_c). Once consumption becomes important enough, people optimally will use a nanny. At this point, the wage jumps up discretely: the decision to use a nanny frees up time, which will be used partly for schooling, which directly translates into the wage. At this point, consumption jumps up and leisure jumps down. Fertility falls somewhat, but note that for high α_c types, parents who use nannies have higher fertility than they would have had if nannies did not exist.

The mechanism behind this example is essentially the same as in section 2.4. People who put a higher weight on consumption goods will invest more in schooling, and hence have higher wages. At the same time, they care less about children and hence have fewer. Note that having leisure in this example is crucial, because once nannies become an option, parents allocate their time only between investing in (own) human capital and working. Given our functional forms, without leisure ($\alpha_\ell = 0$), the optimal allocation would be $l_s^* = l_w^* = 0.5$. But then wages would no longer differ across people, since independent of the preference parameters, everyone would make the same schooling choice. Adding leisure allows for an alternative use of time so that optimal schooling, and hence wages, actually differ across people with different preferences.[59] People who value consumption goods more choose more schooling and less leisure, and therefore have higher wages. These same people also have fewer children. This logic holds even when child care time can be outsourced to nannies, since it is ultimately the relative dislike of children that drives the low fertility of high wage people, and not the high time cost of children. Because of this logic, heterogeneity in *preferences*, rather than in exogenous ability, is essential for this result. Starting from exogenous ability heterogeneity would lead to very different conclusions, as is obvious from the previous solution (and recalling $w = al_s^*$): higher a people have both higher wages *and* more children.

Of course, the mechanism in this example is probably not the only (or even the main) reason for why higher wage people choose lower fertility, even when nannies are an option. Our goal here is to raise an important question and propose a first attempt to answer it. One limitation of the present example is that nanny quality is not a choice. When nanny quality is an input into child quality, specific functional form assumptions are needed to preserve the desired result. This relates back to the quantity-quality trade-off analyzed in section 2.5.

59. This is similar to the preference heterogeneity examples in the couples section, in which the leisure of the husband generated the desired correlation even if his time was not needed to raise children.

2.8 Time Series Implications

Throughout most of this chapter, we have focused on what kind of theories of fertility can match the downward sloping fertility-wage relationship observed in cross-sectional data. We have seen that special assumptions are needed, such as a high elasticity of substitution between fertility and (parent's) consumption. One might want to ask more of such theories. For example, one might want to know the conditions under which such models could also match the decline in average fertility over the last century and a half. In other words, which of these theories can also get the time series facts right, or, how must they be modified to do so?[60] Our static examples are too stylized to empirically test them in any serious fashion. Yet, from section 2.2 there emerged several stylized facts and one way to tackle this question is to see which of the theories can produce a picture that looks qualitatively like figure 2.1. The stylized facts that emerge from this figure can be summarized as:

1. Fertility is very high at low wages (about 6).
2. Fertility is very low at high wages (about 2).
3. Fertility is decreasing (and convex) in wages for each cross section.
4. Fertility falls over time, as consecutive cross sections move to the right.

In terms of forcing variables, it is not obvious which exogenous changes over time to consider. One obvious change over this time period are increases in wages driven by Total Factor Productivity (TFP) growth. Another potentially important change is the development of education, both through technological change that made human capital production more efficient and changes in government policies through the (free) public provision of schooling. Sometimes it is argued that children have become more costly over time, and so we look at this change as well. The interpretation of this change, however, is not straightforward.

Next, we show four numerical examples, each based on a different theory analyzed in the text. Each graph displays four cross-sectional relationships between income and fertility. Depending on the example, the difference between people within a cross section (i.e., on one line) is either wages or preferences, while the difference between different cross sections (i.e., between the four different lines) is either wages, schooling technology, and/or child-rearing costs.

The first two figures are based on two different examples from section 2.3. Figure 2.3 is based on problem (3) while figure 2.4 is based on problem (4),

60. One could also ask the opposite question: which of the existing theories of the demographic transition can generate the cross-sectional fertility facts? Such an analysis is beyond the scope of this chapter.

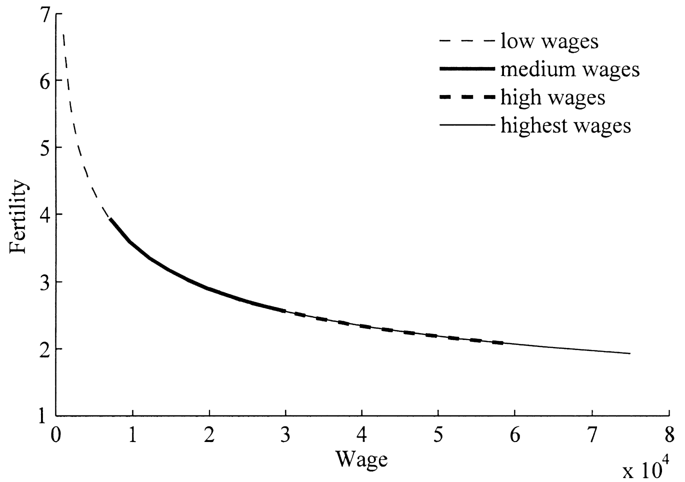

Fig. 2.3 Time series based on price of time example, $\sigma \subset 1$, increasing wages

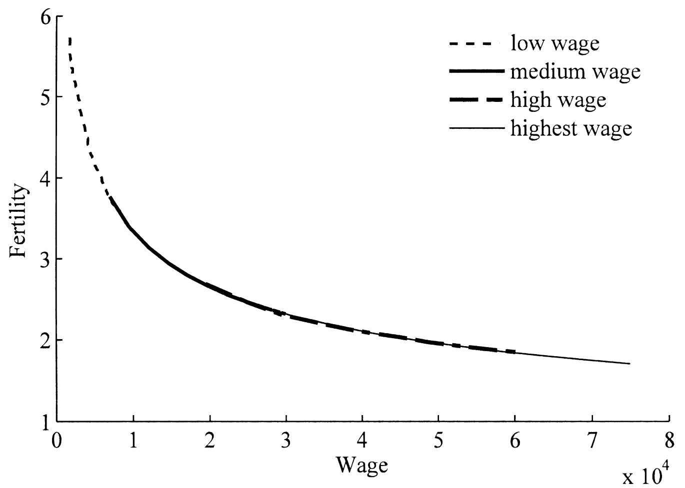

Fig. 2.4 Time series in example with nonhomothetic utility, increasing wages

both variants of the simplest "price of time theory."[61] In each case, the only difference across people (both in the cross section and over time) is wages. Both examples match the stylized facts described before fairly well. Thus, as long as one is willing to assume a high elasticity of substitution between parent's consumption and fertility, the basic theory seems to work well—at

61. The main qualitative difference between the two examples is that the income elasticity is constant in figure 2.4, while it is increasing in absolute value in figure 2.3. Recall also that the empirical elasticity appears to slightly decrease over this time horizon (as shown in table 2.1).

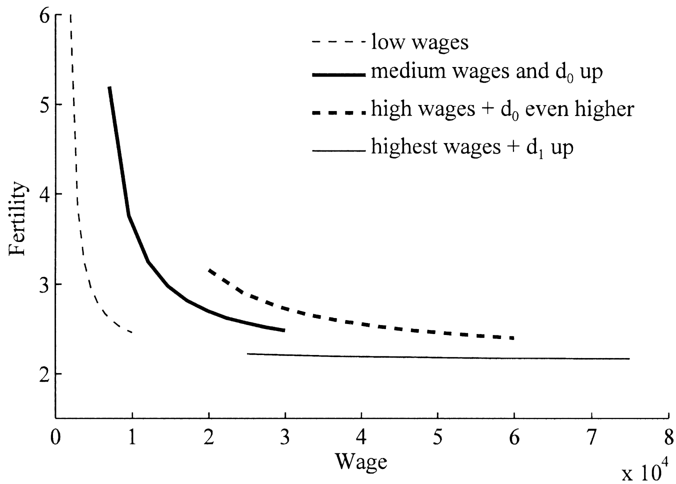

Fig. 2.5 Time series based on quantity-quality example

least in this simple formulation. Once one moves to a truly dynamic formulation, where parents have preferences over their children's utility, the same logic no longer holds, as we discuss in the first section of the appendix. The intuition is simple: when wages go up, both parents' and children's wages are affected. Thus, while the opportunity cost of having a child is higher for richer parents, the benefit of having a child also increases (because the wage of a child of a rich parent is also high). Thus, even though these results seem like strong successes for the theory at first glance, there are other reasonable, but more stringent, requirements for which their success is more limited.

Figure 2.5 considers the quantity-quality trade-off example from problem (7) with $f(s) = d_0 + d_1 s$. Note that to distinguish this example from the first two pictures, this assumes log-utility, and all curvature comes in through the child quality production function only. In this example, fertility is essentially hyperbolic in wages, and hence the shape of the curve does not match figure 2.1 very well.[62] However, this example lends itself to think about potential changes in the education sector. In addition to increasing wages, consecutive cross sections in figure 2.5 face different quality production functions. In particular, the second cross section has a higher d_0, which one could interpret as the introduction of elementary public education. The third cross section has an even higher d_0, which might represent a further expansion of the public education system. The last cross section has a higher d_1, which is a parameter that determines the returns to parental education inputs. This could be interpreted as improvements in education technology. Alterna-

62. One way of stating the qualitative difference between figure 2.5 and the data is that the income elasticity of fertility in the example converges to zero very fast as wages increase, while in the data, the elasticity is roughly constant.

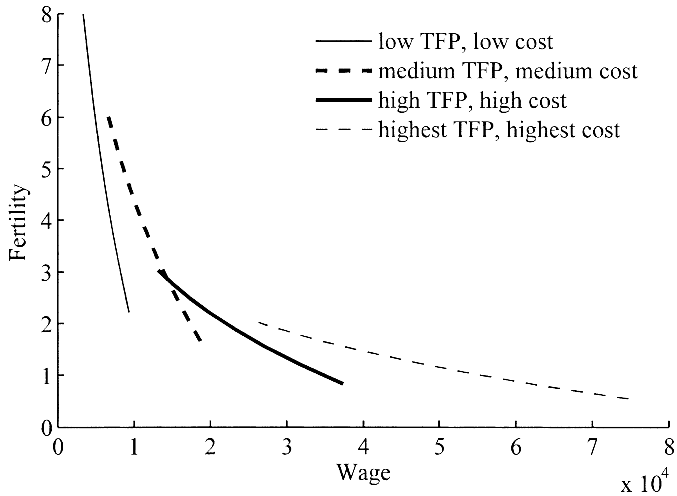

Fig. 2.6 Time series based on increasing TFP and increasing cost of children, cross section due to preference heterogeneity

tively, without this last change in the child quality production function, the last cross section would simply be a continuation of the third cross section, converging to 2.14 children (in this example) as wages go to infinity. So while this picture matches figure 2.1 qualitatively, more work on the underlying changes in education technology (i.e., their historical analogues) would be required before one could call this theory a success.

Finally, figure 2.6 is based on the preference heterogeneity example from section 2.4. In this figure the cross section and time series both slope downward, but the mechanisms behind the two are different. The cross section is based on preference heterogeneity. That is, people who like children invest less in market-specific human capital and therefore have lower wages, while those who put a higher weight on consumption goods do the opposite and therefore have higher wages. Over time, as in the previous examples, we assume that average productivity, a, goes up. However, in this example, increases in productivity do not affect fertility decisions. Hence, without more bells and whistles (e.g., changing the curvature to the utility function), this example will not lead to falling fertility for consecutive cross sections. Thus, we have added a second channel to the time series in the figure: increases in child costs—that is, the units of time required per child increase exogenously over time. This picture looks roughly like the data, but its interpretation is not clear; that is, what is the real-world analogue of an increase in child-rearing costs (measured in units of time)?[63]

63. One rationale for this change may be the progressive introduction of child labor laws. That is, while the time cost remained the same, the time that children contribute to the household's income decreased. Hence, this would be equivalent to a net increase in the time cost.

These simple examples are only meant to spur thinking about the possibilities of the models examined in this chapter. Much more work in carefully calibrating/estimating the relevant parameters and documenting the needed changes in the forcing variables, is necessary before any final conclusions can be drawn. In the end, we cannot offer a clear answer to our own question, but we hope that the ideas here will stimulate further research, leading to a better understanding of fertility decision-making.

2.9 Conclusion

We have investigated the ability of fertility theories to match the cross-sectional relationship between fertility and income. The main focus has been on comparing two sets of theories, one in which ability heterogeneity causes fertility differences and another in which heterogeneity in the taste for children causes income differences. Several interesting findings emerge and are summarized in table 2.2. In particular, we find that low incomes *cause* high fertility only if the elasticity of substitution between consumption and the number of children is high. Empirical research estimating this elasticity would be desirable.

Theories based on taste heterogeneity, on the other hand, do not require any elasticity assumptions. The mechanism causing the negative income-fertility relationship is a very different one, and does not depend on the relative sizes of income and substitution effects. Thus, one may conclude that taste-based theories are more robust. Another advance of taste-based theories is that the assumption of parental time as a critical input into child production is not necessarily needed.

One may also require theories to generate simultaneously a negative income and child quality relationship. While this follows immediately from ability-driven stories, the result is somewhat harder to generate within the class of taste-driven stories. Whether two-parent versions of these theories can generate male wages to be negatively correlated with fertility depends on the details of the models. Generally speaking, with additional assumptions, both classes of theories can do so. However, these both require specific assumptions about how spouses are matched, or about how male and female inputs are combined in family production. In particular, taste-based stories require assortative matching along *preference* lines, while ability-driven stories require assortative matching (or complementarities in production) in *abilities.* Finally, one may ask whether the same driving force that explains the cross-section can also generate the time trend. This is a relatively easy task to accomplish for ability-based stories, because literally the same force that causes richer people to have fewer children in the cross-section also operates as incomes go up for everyone, and thereby mechanically causes a demographic transition. It seems clear that the same mechanism will *not* be able to generate a demographic transition in taste-based theories, unless one believes that tastes for children declined systematically over time.

Table 2.2 **Comparison**

Assumptions and robustness	Ability heterogeneity	Taste heterogeneity
Elasticity (c,n)	Elasticity < 1	Elasticity irrelevant
Parental time	Crucial	Not necessary
Can also get child quality to increase in income?	Yes, plus may help relax elasticity assumption	Depends on details of preference heterogeneity
Can get fertility to decrease in male income, when women do child-rearing?	Need positive assortative matching in ability or complementarities in home production	Need matching along preference lines
Can model also match time series?	Yes	No

In some ways, the analysis in this chapter raises more questions than it answers. It points to several directions for further research, both theoretical as well as empirical. On the empirical side, estimates of the elasticity of substitution between own consumption and children (and child quality versus quantity) would be useful. More generally, clever ways of empirically estimating the contribution of taste-based versus ability-based theories in explaining the negative fertility-income correlation would be valuable. One such attempt is provided in Amialchuk (2006), who uses PSID data and finds that in response to income shocks (specifically, job displacements), couples do not change their lifetime fertility in a significant way. Angrist and Evans (1998), on the other hand, estimate the impact of exogenous variation in fertility (due to twins) on parents' labor supply and find little effect. To the extent that human capital is accumulated on the job, this finding can be interpreted as showing a negligible causal effect from fertility shocks to income. It does not, however, invalidate theories based on preference heterogeneity for consumption goods vis-à-vis children. Clearly, further empirical research to test the various theories is needed.

In addition, a better empirical understanding of the spousal matching process would be helpful. While assortative mating in education has long been documented in the data (e.g., Pencavel 1998), assortative mating in preferences has received less attention. Recent research estimating preferences for marriage markets (e.g., Ariely, Hitsch, and Hortacscu 2006; Lee 2008) may prove useful for understanding better why higher income men have fewer children even though, typically, their wives do most of the child-rearing. Is it because high ability men tend to marry high ability women? Or is it because men with a preference for consumption goods tend to marry women with similar preferences, leading them to spend most of their income on material goods and less on children accordingly?

New research should also develop models of fertility that allow parents to outsource child care. All successful theories of fertility rely on the assumption that it takes the parents' time to raise children. Alternative child care options exist, yet as soon as child care can be bought in the market, the time cost becomes a goods cost for the parents. However, models with only goods cost cannot generate a negative income fertility relationship (with one very specific exception). More theoretical research would be of interest here. For example, modeling explicitly that nannies require monitoring, which in itself may be time-intensive, could be a promising avenue to pursue.

Finally, we found that expanding the successful models to full dynamic versions based on parental altruism is very challenging. Dynamic models are very important for understanding the connection between cross-sectional fertility differences and the demographic transition. More research in this area is needed.

Appendix

Adding Parental Altruism

So far, our focus has been on examining simple models of fertility choice that give rise to the observed pattern in the cross section with respect to income. As we have seen, there are several examples that are capable of this, though they differ in their details. One property that is missing from all of the examples in the main text, however, is altruism of parents toward their children. That is, parents are made happy by things that increase the utility of their children. Altruism introduces an additional dynamic aspect to the fertility choice automatically: when choosing their own fertility levels, parents must forecast the utility levels of their own children. Following this logic, the utility of the children will depend on the utility levels of their own children—that is, the grandchildren—and so forth. Thus, the utility of the current period decision maker depends on the entire future evolution of the path of consumption and fertility, not just the levels chosen this period.

Although this task sounds complex, models of fertility choice based on parental altruism of this form have been worked out in detail in Becker and Barro (1988) and Barro and Becker (1989). Here we develop a simple version of the Barro-Becker model (B-B henceforth) and discuss its relationship with the examples developed in the main text. We show that the simple example discussed in section 2.3 can be interpreted as the problem solved by the typical parent under a setting with dynastic altruism, but that this requires some extra assumptions and has some additional implications. In particular, the simple, static problem with homothetic preferences can be interpreted as the problem from the Bellman's equation for the fully dynamic model where the term relating to fertility choice corresponds to the value function for continuation payoffs. However, this interpretation has the additional implication that the value function also depends on the wage, and because of this, has the property that families with different base wage rates all make the same fertility choices. Thus, although the high elasticity homothetic example has the correct cross-sectional property in the static example, this property does not extend to the fully dynamic version of the model.

In the simplest version of the B-B model, the time t parent solves:

$$\max_{c_t, n_t} \quad u(c_t) + \beta g(n_t) U_{t+1},$$
$$\text{subject to:} \quad c_t + \theta_t n_t \leq w_t,$$

where c_t is current period consumption, n_t is the fertility choice, and U_{t+1} is the utility level of the typical child. Assuming that $g(n) = n^{\eta}$, $u(c) = c^{1-\sigma}/(1 - \sigma)$, successively substituting and changing to aggregate variables for all of the descendants of a given time 0 household, the equilibrium sequence of choices can be represented as the solution to the following time 0 maximization problem:

$$\max_{\{C_t, N_t\}} \quad \sum_{t=0}^{\infty} \frac{\beta^t N_t^{\eta+\sigma-1} C_t^{1-\sigma}}{1-\sigma}.$$

Subject to:

$$C_t + \theta_t N_{t+1} \leq w_t N_t,$$
$$N_0 \text{ given},$$

where C_t is aggregate consumption in period t, N_t is the number of adults in period t, θ_t is the cost of producing a child, and w_t is the wage rate. Implicit in this formulation is the assumption that each adult has the same level of consumption $C_t/N_t = c_t$ in any period.

For this problem to satisfy the typical monotonicity and concavity restrictions, some restrictions on σ and η must be satisfied. There are two sets of parameter choices that satisfy these requirements. The first is the original assumption in Becker and Barro (1988) and Barro and Becker (1989): $0 \leq \eta + \sigma - 1 < 1$, $0 < 1 - \sigma < 1$ and $0 < \eta = \eta + \sigma - 1 + 1 - \sigma < 1$. In this case, $U > 0$ for all $(N, C) \in R_{+}^{2}$. The second possibility is one that allows for intertemporal elasticities of substitution in line with the standard growth and business cycle literature: $\sigma > 1$, $\eta + \sigma - 1 \leq 0$. In this case, utility is negative and $\eta < 0$. When $\eta = 1 - \sigma$ (allowed under both configurations), utility becomes a function of aggregate consumption only.[64]

There are two types of situations under which this maximization problem becomes a stationary dynamic program (where the state variable is N).

64. This formulation for the dynasty utility flow gives rise to some very useful simplifications that we will exploit later. One disadvantage of it, however, is that it is not equivalent to logarithmic utility when $\sigma = 1$. However, when $\eta = 1 - \sigma$ and $\sigma \to 1$, the preferences will converge to those given by the utility function $\Sigma\beta^t \log(C_t)$. See Bar and Leukhina (forthcoming) for an explicit derivation of Barro-Becker preferences with an intertemporal elasticity of substitution (IES) equal to one.

Both cases require constant growth in wages: $w_t = \gamma_w^t w_0$. The first is when the cost of children is in terms of goods, and this cost grows at the same rate as wages: $\theta_t = a\gamma_w^t$. The second case is when the cost of having a child is in terms of time only, $\theta_t = b_{1t}w_t$, where b_{1t} is the amount of time it takes to raise one surviving child.

In either of these cases, the problem of the dynasty overall has a homogeneous of degree one constraint set and an objective function that is homogeneous of degree η. Because of this structure, it follows that the solution to the sequence problem has several useful properties that we will exploit later.

Following the discussion in section 2.3, it follows that only the time cost case is capable of matching the facts from the cross section and hence, we will limit our attention to this case.

Under the special case that $\eta = 1 - \sigma$, it follows that the value function for this problem, $V(N)$, is homogeneous of degree $1 - \sigma$ in N – $V(N) = V(1)N^{1-\sigma}$. Because of this fact, it follows that, after detrending, Bellman's equation for this problem can be written as:

$$V(N) = \sup_{\{C,N'\}} \quad \frac{c^{1-\sigma}}{(1-\sigma)} + \hat{\beta}V(1)N'^{(1-\sigma)}$$

$$\text{s.t.} \quad C + \theta N' \leq wN,$$

where $\hat{\beta} = \beta\gamma_w^{\eta}$. Variable $V(1)$ can be found explicitly. It is given by:

$$V(1) = \frac{(w + \theta(\pi - \gamma_N))^{1-\sigma}}{(1-\sigma)(1 - \beta\gamma_N^{\eta}\gamma^{1-\sigma})}.$$

It follows that the solution to the dynastic problem has a representation in which each date t adult chooses his own consumption and fertility level so as to solve:

$$\max_{\{c,n\}} \quad \frac{c_t^{1-\sigma}}{1-\sigma}\hat{\beta}\, V(1)n_t^{1-\sigma}$$

$$\text{s.t.} \quad c_t + \theta_t n_t \leq w_t.$$

Note that this problem is similar to the CES utility function problem laid out in section 2.3.2. However, there is one important difference. The coefficient on fertility cannot be chosen freely. In particular, it is easy to see that $V(1)$ depends on the wage. Indeed, it follows directly that it is increasing in the wage. Because of this, it follows that the results from the comparative statics concerning the dependence of fertility on the wage are not necessarily valid. In the dynamic version of the problem both the objective function (i.e., Bellman's Equation) and the constraints depend on the wage.

In fact, it can be shown that the equilibrium choice of fertility is given by:

$$\text{(A1)} \qquad n_t = \frac{N_{t+1}}{N_t} = \gamma_N = \left(\beta\gamma_w^{1-\sigma}\left[\frac{w_0}{\theta_0} + \pi\right]\right)^{1/\sigma} = \left(\beta\gamma_w^{1-\sigma}\left[\frac{1}{b_1} + \pi\right]\right)^{1/\sigma},$$

where the last equality follows from assuming that all costs of children are in terms of time, $\theta_0 = b_1 w_0$.

It follows that fertility choices are independent of the level of wages of the family. Thus, although it seems as if the time cost case can reproduce the cross-sectional properties of fertility choice (when $\sigma < 1$ is assumed), this is not true once one restricts attention to static problems that have a dynamic rationalization.[65]

We can also use this framework to get some idea about the implications for differences in fertility across families when preferences for children are the basic source of heterogeneity. For example, we can see that if families differ in their levels of patience, β, differences in the cross section are preserved in the time series. Thus, for example, if for two families, i and i', we have that $\beta_i > \beta_{i'}$, it follows that $n_{it} > n_{i't}$ for all t. Thus, the cross-sectional variation in fertility choice is preserved in the time series.[66] It should be noted however, that this will also have the implication that families with higher fertility also have higher savings rates. This probably does not hold in the cross section.

A Dynamic Version of the Endogenous Wage Example

Next, we develop a version of the endogenous wage model in section 2.4 that is consistent with parental altruism, as in the B-B model.

Assume that the resource constraints are given by those of problem (6), but assume that $\nu_s + \nu_w = 1$. (To simplify notation, write $\nu_s = \nu$ and $\nu_w = 1 - \nu$.) Using capital letters to denote aggregate quantities (i.e., defining $L_t \equiv N_t l_t$, etc.), the planner's problem can be rewritten as:

$$\text{(A2)} \qquad \max \quad \sum_{t=0}^{\infty} \frac{\beta^t N_t^{\eta+\sigma-1} C_t^{1-\sigma}}{1-\sigma}$$

$$\text{s.t.} \quad L_{st} + L_{wt} + L_{nt} \le N_t$$

$$C_t \le a L_{st}^{\nu} L_{wt}^{1-\nu}$$

$$bN_{t+1} \le L_{nt}.$$

As just shown,the constraint correspondence is homogeneous of degree 1 and the utility function is homogeneous of degree η in initial condition N_0.

65. Here we have assumed that wage differences across families are permanent—that is, if i and i' represent two distinct families then we are assuming that $w_{it+1}/w_{i't+1} = w_{it}/w_{i't} = \gamma w$. An interesting question is whether this result will be overtuned when one moves away from this assumption. Jones and Schoonbroodt (forthcoming) find that a high growth rate lowers fertility if $\sigma > 1$ and vice-versa (see also equation [16]). This suggests that with intergenerational mean reversion in income, poor households expect a high income growth rate and would have more children than rich ones as long as $\sigma < 1$. In this context, Zhao (2008) uses a model with filial altruism as in Boldrin and Jones (2002), where mean reversion is crucial, both in the cross section and over time (when social security crowds out fertility). We leave the analysis of intermediate cases (i.e., partially correlated dynastic incomes) to future research.

66. As mentioned before, this assumes that the differences across families is permanent: $\beta_{it} > \beta_{i't}$ for all t.

Assuming that $\eta = 1 - \sigma$ as just shown, the value function is of the form $V(N) = V(1)N^{1-\sigma}$. It follows that the Bellman Equation is:

$$V(N) = \sup_{C,N'} \frac{C^{1-\sigma}}{1-\sigma} + \beta V(1) N'^{(1-\sigma)}$$

$$\text{s.t.} \quad L_s + L_w + bN' \leq N$$

$$C \leq aL_s^{\nu} L_w^{1-\nu}.$$

So for the appropriate choice of α_n and α_c, the solution to problem (6) can be interpreted as the solution to the dynamic problem (A2) with $N_0 = 1$ in some cases. Here, normalizing $\alpha_c = 1$, it follows that $\alpha_n = \beta V(1)$.

It is not clear in this framework exactly which comparative statics exercise corresponds to the one in section 2.4, where α_n is increased. In principle, it could correspond either to an increase in β, or to any increase that makes $V(1)$ larger. In what follows, we consider only the implications of increases, across dynasties, of increases in β's.

Using the first order conditions to the problem in sequence form and simplifying, we obtain a characterization of the balanced growth path dynamics. The system is determined by the division of time between schooling and working and the intertemporal choice of family size involving fertility. It is given by:

$$\frac{L_{wt}}{L_{st}} = \frac{1-\nu}{\nu}, \text{ and}$$

$$n_t^{\sigma} = \gamma_N^{\sigma} = \frac{\beta}{b_1}.$$

That is, fertility is increasing in β. Because of this fact, it follows that both L_{st}/N_t and L_{wt}/N_t are decreasing in β, and hence, fertility and income (or wages) are negatively related as desired.

Thus, for the endogenous wage example, an explicit dynastic form can be provided that is still consistent with the cross-sectional facts. There are still some issues here, however. Foremost, when discount factors differ across agents, strong forces for borrowing and lending are typically present. The analysis here ignores these considerations. It is not certain that the results will be robust to this extension.[67]

Summary of Findings for Couples' Models

In table 2A.1 we summarize the sets of assumptions that are able to generate both a negative correlation between husband's as well as wife's income and fertility.

67. Another issue not considered here is variants of intergenerational persistence in preferences.

Table 2A.1 **Couples: Model versions that work**

Data					< 0	Mixed	< 0
Specialization in production	Exogenous heterogeneity	Curvature in utility	Spousal matching	Other	(1) $\delta n/\delta w_f$	(2) $(\delta n/\delta w_f)\vert w_f$	(3) $\delta n/\delta w_m$
1 Full[c]	Ability	$\sigma < 1$	corr $(w_f, w_m) > 0$		< 0	$= 0$	< 0
2 Full[c]	Ability	$\sigma = 1$	$w_f = \overline{w}$	$\rho < 0$[d]	n.a.[b]	< 0	< 0
3 Full[c]	Taste	Any σ	Prefs matching	Leisure	< 0[a]	n.a.[e]	< 0[a]
4 Partial	Ability	$\sigma = 1$	$w'_f(w_m)\frac{w_m}{w_f} > 1$		< 0	> 0	< 0
5 Partial	Taste	Any σ	Prefs matching	Leisure	< 0[a]	n.a.[e]	< 0[a]

Note: n.a. = not applicable.

[a]The correlation here is for earnings not wages, but including schooling, as in section 2.4, will extend the results to wages.

[b]All women have the same productivity in this example, hence the correlation $\delta n/\delta w_f$ is not well-defined.

[c]In the full specialization version, w_f refers to female productivity at home and cannot be compared directly to data on female wages.

[d]Market goods and female time are complements in home production.

[e]The conditional correlation will depend on the details of the matching process. With perfectly aligned preferences, there is no residual variation in the husband's wage/earnings, conditional on the wife's wage/earnings. Hence the conditional correlation is not well-defined in the model analyzed in the chapter.

References

Ahn, N., and P. Mira. 2002. A note on the changing relationship between fertility and female employment rates in developed countries. *Journal of Population Economics* 15 (4): 667–82.
Aiyagari, S. R., J. Greenwood, and N. Guner. 2000. On the state of the union. *Journal of Political Economy* 108 (2): 213–44.
Alvarez, F. E. 1999. Social mobility: The Barro-Becker children meet the Laitner-Loury dynasties. *Review of Economic Dynamics* 2 (1): 65–103.
Amialchuk, A. 2006. The effect of husband's earnings shocks on the timing of fertility. University of Houston. Unpublished manuscript.
Angrist, J., and W. N. Evans. 1998. Children and their parents' labor supply: Evidence from exogenous variation in family size. *The American Economic Review* 88 (3): 450–77.
Angrist, J. D., V. Lavy, and A. Schlosser. Forthcoming. New evidence on the causal link between the quantity and quality of children. *The Journal of Labor Economics*. Cambridge, MA: National Bureau of Economic Research, December.
Ariely, D., G. Hitsch, and A. Hortacscu. 2006. What makes you click? An empirical analysis of online dating. University of Chicago. Unpublished Manuscript.
Attanasio, O., H. Low, and V. Sanchez-Marcos. 2008. Explaining changes in female labour supply in a life-cycle model. *The American Economic Review* 98 (4): 1517–52.
Bagozzi, R. P., and M. F. Van Loo. 1978. Fertility as consumption: Theories from the behavioral sciences. *The Journal of Consumer Research* 4 (4): 199–228.
Bailey, M. J. 2006. More power to the pill: The impact of contraceptive freedom on women's lifecycle labor supply. *Quarterly Journal of Economics* 121 (1): 289–320.
Banks, W. H. 1955. Differential fertility in Madison Country, New York, 1965. *Milbank Memorial Fund Quarterly* 33:161–86.
Bar, M., and O. Leukhina. Forthcoming. Demographic transition and industrial revolution: A macroeconomic investigation. *Review of Economic Dynamics.*
Barro, R., and G. Becker. 1989. Fertility choice in a model of economic growth. *Econometrica* 57 (2): 481–501.
Becker, G. S. 1960. An economic analysis of fertility. In *Demographic and economic change in developed countries,* no. 11 in Universities—National Bureau Conference Series, ed. Universities National Bureau Committee for Economic Research, 225–56. Princeton, NJ: Princeton University Press.
———. 1965. A theory of the allocation of time. *The Economic Journal* 75 (299): 493–517.
Becker, G., and R. Barro. 1988. A reformulation of the theory of fertility. *Quarterly Journal of Economics* 103 (1): 1–25.
Becker, G. S., and H. G. Lewis. 1973. On the interaction between the quantity and quality of children. *Journal of Political Economy* 81 (2): S279–88.
Becker, G., K. Murphy, and R. Tamura. 1990. Human capital, fertility, and economic growth. *Journal of Political Economy* 98:812–37.
Becker, G. S., and N. Tomes. 1976. Child endowments and the quantity and quality of children. *Journal of Political Economy* 84 (4): S143–62.
Ben-Porath, Y. 1976. Fertility response to child mortality: Micro data from Israel. *Journal of Political Economy* 84 (4): S163–78.
Bernal, R. 2008. The effect of maternal employment and child care on children's cognitive development. *International Economic Review* 49 (4): 1173–1209.

Black, S. E., P. J. Devereux, and K. G. Salvanes. 2005. The more the merrier? The effect of family size and birth order on children's education. *The Quarterly Journal of Economics* 120 (2): 669–700.
Blau, D., and W. van der Klaauw. 2007. The impact of social and economic policy on the family structure experiences of children in the United States. Department of Economics and Carolina Population Center, University of North Carolina. Unpublished Manuscript.
Boldrin, M., M. De Nardi, and L. E. Jones. 2005. Fertility and social security. NBER Working Paper no. 11146. Cambridge, MA: National Bureau of Economic Research, February.
Boldrin, M., and L. E. Jones. 2002. Mortality, fertility and saving in a Malthusian economy. *Review of Economic Dynamics* 5:775–814.
Bongaarts, J. 2003. Completing the fertility transition in the developing world: The role of educational differences and fertility preferences. *Population Studies* 57 (3): 321–35.
Borg, M. O. 1989. The income–fertility relationship: Effect of the net price of a child. *Demography* 26 (2): 301–10.
Browning, M. 1992. Children and household economic behavior. *Journal of Economic Literature* 30 (3): 1434–75.
Cáceres-Delpiano, J. 2006. The impacts of family size on investment in child quality. *Journal of Human Resources* 41 (4): 738–54.
Caucutt, E., N. Guner, and J. Knowles. 2002. Why do women wait? Matching, wage inequality, and the incentives for fertility delay. *Review of Economic Dynamics* 5 (4): 815–55.
Cho, L.-J. 1968. Income and differentials in current fertility. *Demography* 5 (1): 198–211.
Clark, G. 2005. Human capital, fertility, and the industrial revolution. *Journal of the European Economic Association* 3 (2-3): 505–15.
———. 2007. *Farewell to alms: A brief economic history of the world.* Princeton, NJ: Princeton University Press.
Clark, G., and G. Hamilton. 2006. Survival of the richest: The Malthusian mechanism in pre-industrial England. *Journal of Economic History* 66 (3): 707–36.
De la Croix, D., and M. Doepke. 2003. Inequality and growth: Why differential fertility matters. *The American Economic Review* 93 (4): 1091–1113.
———. 2004. Public versus private education when differential fertility matters. *Journal of Development Economics* 73 (2): 607–29.
———. 2009. To segregate or to integrate: Education politics and democracy. *Review of Economic Studies* 76:597–628.
Del Boca, D., and R. M. Sauer. 2009. Life cycle employment and fertility across institutional environments. *European Economic Review* 55 (3): 274–92.
Docquier, F. 2004. Income distribution, non-convexities and the fertility-income relationship. *Economica* 71 (282): 261–73.
Doepke, M. 2004. Accounting for fertility decline during the transition to growth. *Journal of Economic Growth* 9 (3): 347–83.
———. 2005. Child mortality and fertility decline: Does the Barro-Becker model fit the facts? *Journal of Population Economics* 18 (2): 337–66.
Doepke, M., M. Hazan, and Y. Maoz. 2007. The baby boom and World War II: A Macroeconomic analysis. NBER Working Paper no. 13707. Cambridge, MA: National Bureau of Economic Research, December.
Eckstein, Z., and K. I. Wolpin. 1989. Dynamic labour force participation of married women and endogenous work experience. *The Review of Economic Studies* 56 (3): 375–90.

Edin, K. A., and E. P. Hutchinson. 1935. *Studies of differential fertility.* London: P.S. King and Son.
Ehrlich, I., and F. T. Lui. 1991. Intergenerational trade, longevity, and economic growth. *Journal of Political Economy* 99 (5): 1029–59.
Erosa, A., L. Fuster, and D. Restuccia. Forthcoming. A general equilibrium analysis of parental leave policies. *Review of Economic Dynamics.*
———. 2005. A quantitative theory of the gender gap in wages. Working Papers tecipa-199, University of Toronto, Department of Economics.
Espenshade, T. J. 1984. *Investing in children: New estimates of parental expenditures.* Washington, DC: Urban Institute Press.
Falcão, B. L. S., and R. R. Soares. 2008. The demographic transition and the sexual division of labor. *Journal of Political Economy* 116 (6): 1058–1104.
Fernandez, R., N. Guner, and J. Knowles. 2005. Love and money: A theoretical and empirical analysis of household sorting and inequality. *The Quarterly Journal of Economics* 120 (1): 273–344.
Fleischer, B. M., and G. Rhodes. 1979. Fertility, women's wage rates, and labor supply. *The American Economic Review* 69 (1): 14–24.
Francesconi, M. 2002. A joint dynamic model of fertility and work of married women. *Journal of Labor Economics* 20 (2): 336–80.
Freedman, D. S., and A. Thorton. 1982. Income and fertility: The elusive relationship. *Demography* 19 (1): 65–78.
Galor, O. 2005a. From stagnation to growth: Unified growth theory. In *Handbook of economic growth,* vol. 1, ed. P. Aghion and S. Durlauf, 171–293. North Holland: Elsevier.
———. 2005b. The demographic transition and the emergence of sustained economic growth. *Journal of the European Economic Association* 3 (2–3): 494–504.
Galor, O., and O. Moav. 2002. Natural selection and the origin of economic growth. *The Quarterly Journal of Economics* 117 (4): 1133–91.
Galor, O., and D. N. Weil. 1996. The gender gap, fertility and growth. *The American Economic Review* 86 (3): 374–87.
———. 1999. From Malthusian stagnation to modern growth. *The American Economic Review, Papers and Proceedings* 89 (2): 150–54.
———. 2000. Population, technology, and growth: From Malthusian stagnation to the demographic transition and beyond. *The American Economic Review* 90 (4): 806–28.
Goldin, C., and L. F. Katz. 2002. The power of the pill: Oral contraceptives and women's career and marriage decisions. *Journal of Political Economy* 110 (4): 730–70.
Greenwood, J., and N. Guner. Forthcoming. Social change: The sexual revolution. *International Economic Review.*
Greenwood, J., N. Guner, and J. A. Knowles. 2003. More on marriage, fertility, and the distribution of income. *International Economic Review* 44 (3): 827–62.
Greenwood, J., and A. Seshadri. 2002. The U.S. demographic transition. *The American Economic Review, Papers and Proceedings* 92 (2): 153–59.
Greenwood, J., A. Seshadri, and G. Vandenbroucke. 2005. The baby boom and baby bust. *The American Economic Review* 95 (1): 183–207.
Haines, M. R. 1976. Population and economic change in nineteenth century eastern Europe: Prussian upper Silesia, 1840–1913. *The Journal of Economic History* 36 (2): 334–58.
Hansen, G. D., and E. C. Prescott. 2002. Malthus to Solow. *The American Economic Review* 92 (4): 1205–17.

Heckman, J. J. 1976. A life cycle model of earnings, learning, and consumption. *Journal of Political Economy* 84 (4): S11–S44.
Heckman, J. J., and J. R. Walker. 1990. The relationship between wages and the timing and spacing of births: Evidence from Swedish longitudinal data. *Econometrica* 58 (6): 1411–41.
Hotz, V. J., J. A. Klerman, and R. J. Willis. 1993. The economics of fertility in developed countries. In *Handbook of population and family economics,* vol. 1, ed. M. Rosenzweig and O. Stark, 275–347. North Holland: Elsevier.
Hotz, J. V., and R. A. Miller. 1988. An empirical analysis of life cycle fertility and female labor supply. *Econometrica* 56 (1): 91–118.
Jaffe, A. J. 1940. Differential fertility in the white population in early America. *Journal of Heredity* 31:407–11.
Jones, L. E., and A. Schoonbroodt. 2007. Baby busts and baby booms: The fertility response to shocks in dynastic models. Discussion Papers in Economics and Econometrics, no. 0706, University of Southampton.
———. Forthcoming. Complements versus substitutes and trends in fertility choice in dynastic models. *International Economic Review.*
Jones, L. E., and M. Tertilt. 2008. An economic history of fertility in the U.S.: 1826–1960. In *Frontiers of family economics,* vol. 1, ed. P. Rupert, 165–230. Bingley, UK: Emerald Press.
Keane, M. P., and K. I. Wolpin. 2006. The role of labor and marriage markets, preference heterogeneity and the welfare system in the life cycle decisions of black, hispanic and white women. PIER Working Paper Archive no. 06-004, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.
Knowles, J. 2007. Why are married men working so much? The macroeconomics of bargaining between spouses. IZA Discussion Papers no. 2909, Institute for the Study of Labor.
Lazear, E. P., and R. T. Michael. 1980. Family size and the distribution of real per capita income. *The American Economic Review* 70 (1): 91–107.
Lee, R. 1987. Population dynamics of humans and other animals. *Demography* 24 (4): 443–66 (Presidential Address to the Population Association of America).
Lee, S. 2008. Preferences and choice constraints in marital sorting: Evidence from Korea. Stanford University. Unpublished Manuscript.
Li, H., J. Zhang, and Y. Zhu. 2008. The quantity-quality tradeoff of children in a developing country: Identification using Chinese twins. *Demography* 45:223–43.
Manuelli, R. E., and A. Seshadri. 2009. Explaining international fertility differences. *Quarterly Journal of Economics* 124 (2): 771–807.
Merrigan, P., and Y. S. Pierre. 1998. An econometric and neoclassical analysis of the timing and spacing of births in Canada from 1950–1990. *Journal of Population Economics* 11 (1): 29–51.
Miller, A. Forthcoming. The effects of motherhood timing on career path. *Journal of Population Economics.*
Mincer, J. 1963. Market prices, opportunity costs, and income effects. In *Measurement in economics: Studies in mathematical economics and econometrics in honor of Yehuda Grunfeld,* ed. C. F. Christ, L. A. Goodman, and D. Patinkin, 67–82. Stanford, CA: Stanford University Press.
Mincer, J., and H. Ofek. 1982. Interrupted work careers: Depreciation and restoration of human capital. *The Journal of Human Resources* 17 (1): 3–24.
Mincer, J., and S. Polachek. 1974. Family investments in human capital: Earnings of women. *The Journal of Political Economy* 82 (2): S76–S108.

Moav, O. 2005. Cheap children and the persistence of poverty. *Economic Journal* 115 (500): 88–110.
Murtin, F. 2007. American economic development or the virtues of education. Stanford Center for the Study of Poverty and Inequality. Unpublished Manuscript.
Pencavel, J. 1998. Assortative mating by schooling and the work behavior of wives and husbands. *The American Economic Review* 88 (2): 326–29.
Preston, S., and C. Sten. 2008. The future of American fertility. University of Pennsylvania, Discussion Paper. Manuscript prepared for NBER Conference on Demography and Economics. April, Napa, California.
Rosenzweig, M. R., and K. I. Wolpin. 1980. Testing the quantity-quality fertility model: The use of twins as a natural experiment. *Econometrica* 48 (1): 227–40.
Ruggles, S., M. Sobek, T. Alexander, C. A. Fitch, R. Goeken, P. K. Hall, M. King, and C. Ronnander. 2004. Integrated public use microdata series: Version 3.0. Minneapolis: Minnesota Population Center.
Schultz, T. 1981. *Economics of population.* Reading, MA: Addison Wesley.
———. 1986. The value and allocation of time in high-income countries: Implications for fertility. *Population and Development Review* 12:87–108.
———. 2005. Effects of fertility decline on family well being: Opportunities for evaluating population programs. Yale University. Working Paper.
Shiue, C. H. 2008. Human capital and fertility in Chinese clans, 1300–1850. University of Colorado at Boulder. Unpublished Manuscript.
Simon, J. L. 1977. *The economics of population growth.* Princeton, NJ: Princeton University Press.
Skirbekk, V. 2008. Fertility trends by social status. *Demographic Research* 18 (5): 145–80.
Sylvester, M. S. 2007. The career and family choices of women: A dynamic analysis of labor force participation, schooling, marriage and fertility decisions. *Review of Economic Dynamics* 10 (3): 367–99.
Tertilt, M. 2005. Polygyny, fertility, and savings. *Journal of Political Economy* 113 (6): 1341–71.
U.S. Census. 1945. Differential fertility 1910 and 1940. U.S. Bureau of the Census, Census of Population. Washington, DC: GPO.
———. 1955. Fertility. U.S. Bureau of the Census, Census of Population. Washington, DC: GPO.
Weir, D. R. 1995. Family income, mortality, and fertility on the eve of the demographic transition: A case study of Rosny-Sous-Bois. *The Journal of Economic History* 55 (1): 1–26.
Westoff, C. F. 1954. Differential fertility in the United States: 1900 to 1952. *American Sociological Review* 19 (5): 549–61.
Whelpton, P., and C. Kiser, eds. 1951. *Social and psychological factors affecting fertility.* Milbank Memorial Fund.
Willis, R. J. 1973. A new approach to the economic theory of fertility behavior. *Journal of Political Economy* 81 (2): S14–S64.
Wrigley, E. A. 1961. *Industrial growth and population change.* Cambridge: Cambridge University Press.
Zhao, K. 2008. Social Security, differential fertility, and the dynamics of the earnings distribution. University of Western Ontario. Unpublished Manuscript.