

PRIVATE INCENTIVES AND SOCIAL INTERACTIONS: FERTILITY PUZZLES IN ISRAEL

Charles F. Manski
Northwestern University

Joram Mayshar
Hebrew University of Jerusalem

Abstract

This paper explores how private and social incentives for fertility may have combined to produce the complex fertility pattern observed in Israel in the past half-century. Fertility has declined within some ethnic-religious groups, moderately increased in others, and parts of the ultra-Orthodox Jewish population have experienced a *reverse fertility transition*, in which childbearing has increased rapidly and substantially. We present a theoretical analysis of the social dynamics of fertility that shows how private preferences, preferences for conformity to social norms in childbearing, and piecewise linear child allowances could have combined to yield such a complex fertility pattern. We then explain the identification problem that makes it so difficult to infer the actual Israeli fertility process from data on completed fertility. (JEL: J13, Z13, H53)

1. Introduction

Demographers have long wavered between two perspectives on fertility behavior: Homo Economicus, who chooses an optimal family size, and Homo Sociologicus, who conforms to social norms [see Andorka (1978) and Hammel (1990)]. These two perspectives dominate the ongoing debate on the underlying sources of the fertility transition—the recent historic sharp decline in fertility toward rates that are below replacement. Summaries of this debate by Cleland and Wilson (1987) and by Bongaarts and Watkins (1996) find fault with the economic approach for failing to explain why the fertility transition has occurred in different places under rather different economic conditions, and why the initial stage of the transition is so rapid, relative to the more gradual pace of economic change. Both articles cite a general need to incorporate sociological elements into economic models of fertility, and a particular need to understand how social norms affect fertility and evolve.

Acknowledgments: This research was supported in part by a grant from the Israel Foundations Trustees (1998–2000), and in part by United States National Science Foundation Grant SES-0001436. The Israel Central Bureau of Statistics prepared for our use the merged Census and Birth Registry data file analyzed here. We are grateful to the Bureau staff, including Elana Jacoby and Eliezer Kamon, for their efforts. Arie Beresteanu, Philip Cross, Shira Klein, Francesca Molinari, and Jora Stixrud provided excellent assistance in our analysis of the data. We thank Eliahu Ben-Moshe, Momi Dahan, Charles Kamen, and Ari Peltial for very constructive comments on an earlier version of the paper.

E-mail addresses: Manski: cfmanski@northwestern.edu; Mayshar: msjoram@mscc.huji.ac.il

Although discussions of fertility sometimes give the impression that economics and sociology offer competing models of fertility, the two approaches are compatible with one another. Women may choose family size to maximize utility functions that recognize both private and social incentives for fertility, and social incentives may themselves evolve as an outcome of these childbearing decisions. Indeed a rapidly growing literature on the economics of social interactions has sought to understand how such processes operate in education, residential choice, labor markets, and many other environments [see Becker and Murphy (2000), Manski (2000), and the references therein]. Kohler (2001) presents a recent attempt to formally incorporate social interactions into the analysis of fertility choice.

This paper explores how private and social incentives for fertility may have combined to produce the rather complex fertility pattern observed in Israel in the past half-century. Table 1 summarizes random sample data on completed fertility (lifetime number of children) among Israeli women alive and married in 1995, whose year of first marriage was either before 1955 or in the period 1970–1980. The table is based on a special file that merged information from the 1983 Census and the 1995 national Population Registry prepared at our request by the Israel Central Bureau of Statistics. It shows that fertility in Israel displays considerable variation both over time and across major ethnic/religious subgroups of the population. (Further details about the data are provided later.)

Comparing women whose first marriage occurred prior to 1955 with those who first married 1970–1980, Table 1 shows that completed fertility declined among Jewish women of Mizrahi origin (whose father was born in Asia or North Africa) and among non-Bedouin Arab women, in accord with trends worldwide. Yet childbearing increased in all other subgroups of the population. It increased slightly (from 3.06 to 3.12) among non-ultra-Orthodox Jewish

TABLE 1. AVERAGE COMPLETED FERTILITY AMONG ISRAELI WOMEN ALIVE IN 1995

Ethnicity	Years of first marriage			
	Prior to 1955		1970–1980	
	Ultra-Orthodox	Others	Ultra-Orthodox	Others
Jews—Mizrahi origin	6.86 (217)	5.23 (10,491)	5.11 (274)	3.46 (20,427)
Jews—Ashkenazi origin	2.76 (571)	2.29 (14,954)	5.88 (275)	2.88 (17,164)
Jews—Israeli-born parents	3.03 (40)	3.06 (793)	5.91 (57)	3.12 (2,636)
Arabs—Non-Bedouins	8.41 (2,290)		5.55 (6,055)	
Arabs—Bedouins	7.32 (34)		9.02 (139)	

In parentheses: the number of women in each cell in the sample.

women whose parents were born in Israel, and moderately (from 2.29 to 2.88) among non-ultra-Orthodox women of Ashkenazi origin (whose father was born in Europe or America). In a phenomenon that can only be described as a *reverse fertility transition*, fertility increased rapidly and substantially (from 3.03 to 5.91, and from 2.76 to 5.88) among ultra-Orthodox women of these origins.

What explains the dissimilar levels and trends in completed fertility among different groups of Israeli women? It is easy to suggest a set of private and social forces that may have combined to yield the complex pattern of fertility depicted in Table 1. The potentially relevant sources of time-series and cross-sectional variation in fertility include the following¹:

The Fertility Transition: Standard economic and social-cultural explanations of the fertility transition can readily rationalize the substantial declines in fertility among Israeli Arabs and among Jews of Mizrahi origin. In both groups, women whose first marriage occurred prior to 1955 lived in largely traditional societies characterized by low income and high child mortality, but women who first married in 1970–1980 lived in the western-oriented society of modern Israel, characterized by relatively high income and low child mortality.

Religiosity: Ultra-Orthodox Jewish women follow the directives of their rabbis, who encourage high fertility and discourage the use of contraceptives. Other Jewish women may hold varying preferences for family size and are not subject to strong religious strictures against contraception.

Social Interactions: The various ethnic and religious subpopulations of Israeli society interacted to different degrees over the past half-century, and so may have been subject to different social norms for childbearing. Adoption of the fertility norms of Ashkenazi Jews by the Mizrahi Jews, most of whom migrated to Israel in the 1950s, may partly explain the fertility decline in the latter group. Increased integration between these two groups may partly explain the increased fertility among the former. Yet, throughout the past half-century, the Jewish ultra-Orthodox, the Jewish non-ultra-Orthodox, and the Arab subpopulations of Israel have largely resided, been schooled, and worked in separate, almost isolated communities.

Recent History: The traumatic decimation of European Jewry in the Holocaust may have affected fertility in the period following World War II. Thus, it is sometimes argued that the Holocaust reduced the fecundity of the women who were directly affected, but increased the desire to have children among survivors and the kin of nonsurvivors. It has also been suggested that the relatively high

1. For a discussion of the underlying sources of the Israeli fertility pattern see Friedlander and Feldmann (1993).

death rates experienced or expected by some subgroups of Israelis in repeated wars may have increased the desired number of children, as a form of insurance.

Child Allowances and Related Public Welfare Programs: Changes in public policy have generated time-series variation in private incentives for fertility. In 1970–1975, the Israeli government introduced a universal (non-income-tested) child allowance program that ranks among the most generous in the world. Modest monthly payments are made to families with one or two children under age 18, but the payments for each child from the fourth up are substantial. The 1960s and early 1970s also saw the introduction of other major welfare programs that reduced the private cost of childbearing and childcare. On the whole, women whose first marriage occurred prior to 1955 bore their children before the introduction of the child allowance and other welfare programs, whereas women who first married after 1970 were subject to them throughout their childbearing years.

Juxtaposition of the child allowance program and the Israeli fertility patterns in Table 1 suggests that enactment of the allowance program in 1970–1975 may explain some of the trends shown, particularly the reverse fertility transition among the ultra-Orthodox of Ashkenazi origin. However, the same allowance program applied to other segments of the Jewish population, who experienced a variety of different fertility trends.

These and perhaps other disparate forces have somehow combined to determine fertility in Israel. The puzzle, or set of puzzles, is how to disentangle these forces and identify the socioeconomic process at work. From a public policy perspective, there is particular interest in learning how the child allowance and other welfare programs have affected, and continue to affect, fertility. From a social science perspective, there is a general interest in learning how private and social incentives interact to determine childbearing. This paper addresses the puzzles posed by Israeli fertility patterns and makes some, albeit incomplete, progress towards its resolution.

We begin with the public policy question, focusing on the child allowance program. Section 2, which draws on Mayshar and Manski (2000), describes the history of the program, compares it with similar ones in Europe, and recalls Malthus' hypothesis on the fertility effects of child allowances. We also present in this section a more complete picture of how Israeli fertility evolved over time for different social groups.

The empirical associations that we find between the child allowance program and the time-series variation in fertility are intriguing but necessarily inconclusive. The challenge is to dig beneath these associations to learn how private and social incentives may have combined to determine fertility. Section 3 poses a model of family-size decisions that flexibly represents some aspects of decision making while abstracting from others. The model assumes that, at some point after marriage, a woman chooses a family size that is optimal, given the

information available at the time.² The woman's valuation of bearing a given number of children depends on her private preference for family-size, her preference for conformity to social norms in childbearing, and the child allowance that the mother would receive. The model permits private and social preferences to vary across women. Perhaps the main respect in which the model abstracts from actual childbearing behavior is its neglect of the dynamics of birth timing, a subject that has been a focus of some recent economic research [see Hotz and Miller (1988) and Walker (1995)]. Abstracting from timing issues enables us to determine how optimal family size varies with fertility preferences, social norms, and child allowances.

Our model of family-size decisions, although simple in many respects, is still too complex in its general formulation to permit much in the way of theoretical analysis. After making further simplifying assumptions, we are able to characterize the social dynamics of fertility with considerable clarity. Given these assumptions, we find that a proportional allowance formula, in which the magnitude of the benefit is proportional to the number of children, generates a simple dynamic in which the distribution of family size converges to a unique steady state. In this case, the value of steady-state mean family size varies linearly with the magnitude of the allowance; it is however independent of the magnitude of preferences for conformity to social fertility norms and of the dispersion of the distribution of private fertility preferences.

A piecewise linear convex allowance formula of the type enacted in Israel is then shown to generate rather different dynamics. It continues to be the case that the distribution of family size converges to a steady state. However, steady-state mean family size now does depend on preferences for conformity to social norms and on the shape of the distribution of private fertility preferences. Moreover, some configurations generate multiple steady states, some being locally stable and others not. Performing numerical calculations in particular cases, we are able to demonstrate that small changes in the magnitude of allowances or in private fertility preferences can, in principle, yield rather large changes in mean fertility; on the order of the changes observed in Israel between 1950 and 1980.³

Section 4 describes our efforts to use the fertility model developed in Section 3 to interpret the actual decisions made by Israeli women. The goal was to estimate a version of the model that could be applied to examine how Israeli child allowance policy has affected family-size decisions. Unfortunately, we

2. We follow the long tradition in demography of attributing fertility decisions to women rather than to couples. This attribution is merely semantic in the large majority of cases in Israel, which has been characterized by low rates of nonmarital childbearing.

3. Previous research on endogenous social interactions, wherein each decision maker is influenced by the decisions of others, has shown that such interactions can produce social multipliers that magnify the individualistic effects of policy interventions (see Manski 1993). Previous research has also shown that endogenous interactions can generate complex dynamics with potentially multiple steady states (see Brock and Durlauf 2001). Our findings are based on an interaction model that differs in some respects from those studied earlier.

were not able to credibly separate the effects on decision making of child allowances, private fertility preferences, and preferences for conformity. Our attempts to fit reasonably flexible versions of the model yielded unstable, unreliable estimates. A qualitatively ‘sensible’ estimate of the effect of child allowances on fertility emerged only when we maintained assumptions that we feel are too strong to be believable.

The main contribution of Section 4 is to explain the identification problem that we faced when attempting to estimate the fertility model. The basic problem, common to all econometric analysis of choice behavior, is that observation of a woman’s chosen family size only partly reveals her preferences; hence the fertility preferences of Israeli women cannot be learned in an entirely empirical manner. A specific problem in our analysis is that the information about preferences revealed by data on completed fertility depends qualitatively on the shape of the child allowance formula that women face; revealed preference analysis is more difficult when the benefit is proportional to family size than when it is nonlinear. The Israeli allowance has, to varying degrees over the years, been nonlinear in family size. However, we found in practice that nonlinearity of the formula provides an insubstantial foundation for empirical analysis.

Viewing our work in toto, we see a descriptive empirical contribution in our depiction of the complex time-series and cross-sectional pattern of Israeli fertility (Section 2). Our theoretical analysis of the joint determination of fertility dynamics by private preferences, preferences for conformity, and child allowances helps to understand how such a complex pattern of fertility may have come to be (Section 3). However, our structural empirical analysis of fertility decisions does not successfully disentangle the forces at work (Section 4). This being so, we can at best conjecture on the processes that have generated Israeli fertility. Section 5 presents our concluding thoughts.

2. Child Allowances and Fertility in Israel

Section 2.1 describes the history of the Israeli child allowance program. Section 2.2 compares this program with ones in Europe and summarizes empirical research on the Malthusian hypothesis that child allowance programs increase fertility among the poor. Section 2.3 then examines fertility patterns in Israel.

2.1 The Israel Child Allowance Program

The Israeli welfare state has subsidized childbearing and childcare through a number of mechanisms, including provision of free education and health insurance, generous maternity grants, tax benefits for working mothers, and housing benefits that depend on the number of children.⁴ In the major tax reform of 1975, Israel instituted a generous child allowance program that quadrupled the allow-

4. See Doron and Kramer (1991) for background on the development of Israel’s welfare system.

ances granted up to 1969.⁵ Since 1975, the National Insurance Institute (NII) has annually allocated more than 1.6 percent of GDP to this program. A notable feature of the program is that the size of the allowance varies substantially with the birth order of the child, with the first two children under age 18 receiving minimal benefits, and each child from the fourth on receiving a large benefit.

Enactment of the child allowance program was in part a response to growing ethnic discontent among Mizrahi Jews who had immigrated to Israel from Arab countries after 1948 or from North Africa in the early 1960s. These Jews were often poor and had large families. A universal child allowance was thought at the time to be an effective way to reduce their poverty, without creating the work disincentives thought to be characteristic of income-tested benefits.⁶ In fact, enactment of the child allowance program appears to have had a dramatic short run impact in reducing the incidence of poverty among Israeli children in the mid-1970s. Since then, however, the percentage of Israeli children living below the poverty line has steadily increased.⁷

The child allowance program uses a credit point system under which mothers receive a monthly, tax-exempt allowance from the NII, the allowance equaling the number of credit points to which her children are entitled, multiplied by the cash value of a credit point. The number of credit points depends on the number of children in the family who are below the age of 18; prior to 1997, it also depended on whether the family had the status of army veterans.⁸ Table 2 presents the credit point formula in effect from 1975 through 2000.

5. The child allowance program enacted under the 1975 tax reform replaced a complex system of benefits for children that included tax credits, small mandatory child payments by private employers, minor allowances by the NII to large families, and a more substantial allowance to children of army veterans that had been enacted in 1971.

6. The report of the national commission that recommended the 1975 tax reform (State of Israel 1975) indicates that the possible impact of child allowances on fertility was not an important consideration at the time. Moreover, an earlier commission of demographers (State of Israel 1966), set up to study pronatalist policies, had concluded in 1966 that: "it would not be appropriate to recommend monthly grants for families as an incentive to have children for the following reasons: small grants do not have any impact on increased fertility . . . if the grant . . . would be . . . large, then considerable sums of money will flow to large families, for whom it would not provide a birth incentive, because they do not practice birth control anyway, but it could turn out to be a factor that discourages breadwinners from working."

7. According to NII annual reports, the incidence of child poverty rose from 9.6 percent to 13.9 percent between 1979 and 1984. Using somewhat different definitions, between 1987 and 1997, the incidence of child poverty rose from 9.5 percent to 11.8 percent among families of wage earners, and from 20.5 percent to 21.8 percent in the population at large. The increase in the overall incidence of child poverty, despite the significant decline of family size among the relatively poorer two main subgroups of the population, Arabs and Mizrahi Jews, can be explained by the increased weight of these two groups in the population, and, in addition, by the increased family size and population weight of ultra-Orthodox Jews.

8. The conditions for obtaining veteran status prior to 1993 were such that almost all Jews received it, including the ultra-Orthodox who by and large did not serve in the army. Almost all Arabs, for whom army service is voluntary, did not receive it. Beginning in 1993, the dependence of the allowance on veteran status was gradually phased out over four years. From 1997 on, all children receive the full allowance.

TABLE 2. SUMMARY OF CHILD ALLOWANCES CRITERIA

Child order (k)	Allowance credit points					Percent of official standard adult poverty line income					Marginal cost of care for k th child
	Non-Veterans	Veterans				Marginal allowance granted for k th child, for those with veteran status					
		1975–1982	1983–1993	1994–1999	2000	1965	1970	1975	1992	1997	
1	1	1	1	1	1	6.6	11.0	24.2	0 ¹	14.2	70
2	1	1	1	1	1	6.6	13.2	24.2	0 ¹	14.2	55
3	1.25	2	2	2	2	6.6	13.8	48.5	33.6	28.4	55
4	1.25	2.25	3.75	4.05	4.4	11.0	20.4	54.0	96.6 ¹	57.4	50
5	1.25	2.5	3.25	3.4	5	13.2	20.4	54.5	54.5	48.2	50
6	1.25	2.5	3.5	3.75	5	14.3	21.5	60.6	58.8	53.2	45
7+	1.25	2.5	3.5	3.5	5	16.5	21.5	60.6	58.9	49.6	40

Source: National Insurance Institute Statistical Quarterly and Yearbooks (various years).

¹In 1992 the allowances for the first two children were withdrawn for families with less than four children, unless family income was below a low threshold.

Figure 1 shows how the New Israel Shekel (NIS) value of Israeli child allowance benefits has evolved over the longer period 1965–1999.

As Figure 1 and Table 2 indicate, various adjustments have been made in the child allowance program since 1975. The real value of credit points eroded initially due to inflation, but has been fully linked to the CPI since 1987. In addition, repeated changes have been made in the credit point formula, with the number of credit points for children of high birth-order gradually increasing

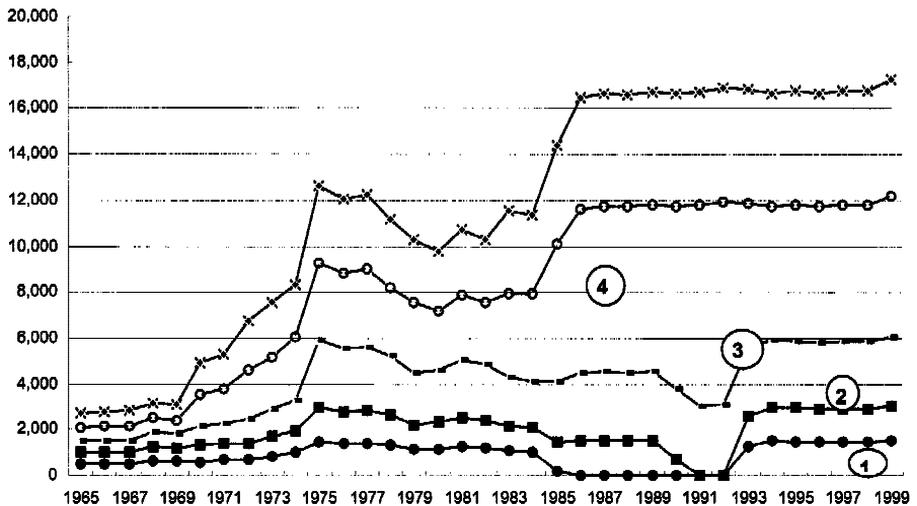


FIGURE 1. Annual Value of Child Allowances per Family, by Number of Children under Age 18 (NIS in constant 1995 prices. \$1.00 = NIS 3)

over time. In 1983, the number of credit points allocated to families with veteran status and four or more children increased by 50 percent.⁹ In 1994 and 2000, the credit points for children from the fourth up were raised further.

As Table 2 shows, the marginal allowance granted for children from the fourth up fully covers the marginal cost of caring for these children in a manner that averts poverty, as calculated using Israel's official poverty line for families of different sizes. To give a specific sense of the magnitude of the allowances, consider families with six children under age 18. In 1999, such families received a monthly tax-exempt allowance of NIS 2,566, equivalent to about \$640 per month, and to 39 percent of the income level officially defined to mark the poverty line for families of that size.

2.2 European Child Allowances and the Malthusian Hypothesis

Allowances for children are not new. Already in the second century AD, an extensive alimenta system of public subsistence payments to children was in force in Rome. According to Duncan-Jones (1974, p. 295), this system was "evidently intended to encourage a rise in the birthrate." The motivation apparently was the shortage of Romans to run the empire, but other historians emphasize other motives.

The English Poor Laws, initiated during the reign of Elizabeth I, were clearly motivated by welfare considerations rather than by a pronatalist sentiment. As is well known, it was a proposal by William Pitt to increase family benefits for the poor that prompted Thomas Malthus' 1798 *Essay on the Principle of Population*. In the fifth chapter, Malthus contended that providing financial support for poor families as a function of family size would not contribute toward alleviating poverty; to the contrary, in the long run it would exacerbate poverty by encouraging poor families to have more children or to work less. According to Malthus, well-intended allowances for families in need may thus become a trap that could contribute toward perpetuating poverty in general, and among children in particular.¹⁰ The empirical validity of Malthus' hypothesis on the fertility effects of the English Poor laws was subject to much criticism, but a recent study by Boyer (1989) seems to confirm Malthus.

Modern-day child allowances seem to be motivated by a mixture of concern for the welfare of children and concern with forecasts of declining population.¹¹ The latter was clearly paramount in the case of France and Belgium, which

9. Between 1985 and 1992, fiscal concerns led to suspension of the allowance for the first child to families with fewer than four children and with income above a certain low threshold; in 1991–1992 this suspension was extended to the second child. This created some work disincentives and administrative confusion due to the income threshold, and also implied that the marginal allowance for the fourth child, whose total allowances were kept intact (for families with income above the income threshold) increased from 3.75 to 4.75 and then to 5.75 credit points.

10. Under Malthus' influence, Pitt withdrew his proposal. Later on, Britain radically amended its poor laws, still much influenced by Malthus' ideas.

11. See Teitelbaum and Winter (1985) for an insightful account of the fear of population decline and of policies that were enacted in different countries to reverse this decline.

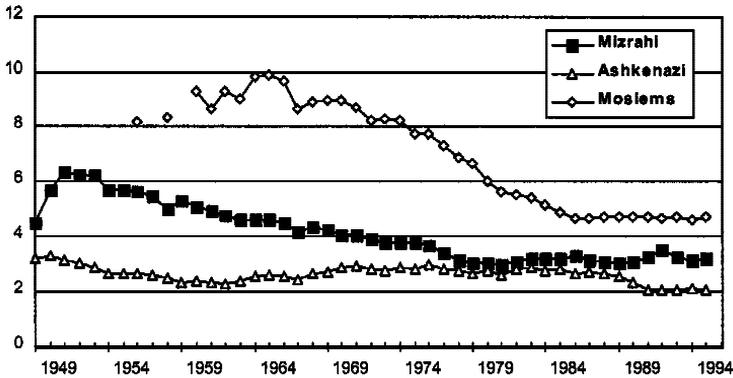


FIGURE 2. Total Fertility Rates of the Main Subpopulations in Israel, 1949–1996

introduced universal child allowances in the 1930s. Following their lead, all European countries and almost all other developed countries in the world followed suit—with the notable exception of the United States.¹² Notwithstanding the large magnitude of these programs, there has been no definitive study of the demographic effects of modern European programs of child and family support. Yet the consensus seems to be that these programs have had at most a negligible impact on fertility. It is often noted that France, with possibly the most generous (and complex) child support program in the world, has an overall birth rate no higher than that of neighboring countries, where the allowance program is less extensive. Based on an econometric study of twenty-two industrialized countries from 1970 to 1990, Gauthier and Hatzius (1997) have reported that child support programs may have had a positive effect on birth, but one of a very small magnitude. Specifically, they estimate that a 25 percent increase in child support benefits would increase total fertility from an average of 1.71 children per woman to 1.78 children per woman. Several studies on European countries have concluded that pronatalist policies have encouraged women to give birth earlier, yet have had practically no effect on total lifetime births (see Walker 1995; Gauthier and Hatzius 1997).

2.3 Fertility Patterns in Israel

The standard source of information on fertility in Israel has been annual data on total fertility, a construct that measures the average lifetime number of births a woman would have if she were to give birth over time according to the current birth rates of women of different age cohorts. Figure 2 presents the time series of total fertility in Israel among the three major ethnic subpopulations: Ash-

12. See Watson (1954) for an extensive survey of the French program, in its earlier years. For a comparison of the terms of child, or family, allowances across countries see Social Security Administration (1999).

kenazi Jews, Mizrahi Jews, and Arab Moslems.¹³ The figure shows a significant decline up to the middle 1970s and middle 1980s within the latter two groups, and a fluctuating total fertility rate among Ashkenazi Jews.¹⁴

The trends depicted in Figure 2 do not suggest that the child allowance program, enacted in 1970–1975, has had any material effect on fertility. However, total fertility is a synthetic-cohort construct that may not adequately represent the childbearing decisions of actual women. Moreover, the Malthusian hypothesis concerning the fertility effects of child allowances should apply primarily to poor families, who form just part of the large subpopulations depicted in Figure 2. For these reasons, we decided to obtain data that would enable a less aggregated and more accurate portrayal of Israeli fertility patterns than has previously been possible.

At our request, the Central Bureau of Statistics prepared a special data file merging information from the 1983 Census of Population and Housing and the 1995 Population Registry. We begin with the random sample of all households in which the husband and wife completed the ‘long form’ of the 1983 Census. Within these households, we identified all women born between 1920 and 1960 who were married at the time of the census, whose husbands were also enumerated, and whose national identity card numbers appear in the 1995 Population Registry. The Population Registry links the identity numbers of children to those of their mothers; hence we were able to determine the number and age of children born to each woman until 1995.¹⁵

To understand Israeli fertility patterns, we disaggregated these women into marriage cohorts defined by year of first marriage.¹⁶ It has been more common in the demography literature to disaggregate women into birth cohorts defined by their year of birth. However, we felt that the timing of fertility decisions may be more closely related to year of marriage than to year of birth.

We also felt it important to disaggregate the three major ethnic subpopulations into more homogeneous groups. Information collected from respondents to the 1983 Census made it straightforward to distinguish Bedouins from other

13. The trend in total fertility of Arab Christians, not shown in the figure, differs from that of Arab Moslems. The total fertility rate of Christians initially was similar to that of Moslems, but declined by the 1990s to levels below that of Ashkenazi Jews.

14. Berman (2000) has used data from surveys of the Israeli labor force to estimate total fertility among Israeli ultra-Orthodox in the periods 1980–1984 and for 1994–1996. He finds (Table 6) that total fertility increased from 6.9 to 7.8 among Ashkenazi ultra-Orthodox, and from 4.6 to 7.2 among Mizrahi ultra-Orthodox. Berman interprets this recent rise in total fertility, as well as increased religious stringency and reduced labor participation, as a response of a “club to a change in wages and [non-child-specific] transfers” (p. 940, parenthetical explanation added).

15. The linkage of children’s identification numbers to those of their mothers is quite accurate for children born after 1960, but incomplete for children born in the 1950s and almost nonexistent for children born earlier. Fortunately, the 1983 census itself provides information on the total number of live births up to that year. A detailed explanation of the manner in which we combined Population Registry and Census data to determine each sampled woman’s births is available in a technical memorandum that can be obtained from the authors upon request.

16. In cases where a woman had multiple marriages, here and throughout the paper we use “year of marriage” to mean “year of first marriage.”

Arabs, and to distinguish Jews whose parents were born in Israel from those whose parents immigrated to the country.¹⁷ We also thought it important to disaggregate Jewish Israelis by degree of religiosity, particularly to distinguish ultra-Orthodox Jews from others. This was more problematic, because respondents to the 1983 Census were not asked to report their religious practices. We ultimately used a schooling-based criterion, previously applied by Dahan (1998), whereby families are identified as ultra-Orthodox if the last place of study that was reported by the husband in the 1983 census was a postsecondary *yeshiva* (an Orthodox institute of religious study).¹⁸

Table 1, introduced in Section 1, uses the merged Census and Population Registry data to present completed fertility rates within each of eight groups defined by ethnic origin and religiosity. Given that some women who married before 1980 were still of childbearing age in 1995, we used ethnic-group, age-specific fertility rates for 1994–1995 to predict childbearing by these women after 1995. The table provides a simple ‘before-and-after’ perspective on the association between fertility and the 1970–1975 child allowance program. To accomplish this, we distinguish women whose first marriage occurred prior to 1955, who generally completed childbearing before initiation of the program, from those first married in the period 1970–1980, who were eligible for the allowances throughout most or all of their childbearing years. The table excludes members of the 1955–1970 marriage cohorts because their childbearing period straddles the date of the policy change. It excludes those who first married after 1980, who married only shortly before the 1983 census and had less than 15 years of marriage up to 1995.

The disaggregated data on completed fertility in Table 1 present a more nuanced portrayal of Israeli fertility patterns than do the more aggregated data on total fertility presented in Figure 2. As discussed in Section 1, the table shows that Israel manifests the worldwide trend towards reduced childbearing only within the Mizrahi–Jewish and Arab–non-Bedouin subpopulations. Within all other groups, completed fertility among women who first married in the period near the 1975 reform is higher than among women who married earlier on. The increase in completed fertility among ultra-Orthodox Jews of non-Mizrahi origin is little short of astonishing. On average, women in these groups approximately doubled their childbearing, from about three children per women to about six.

Other analyses of our sample data reinforce the impression that the completed fertility rates of ultra-Orthodox Ashkenazi Jews substantially increased in

17. In principle, we could further disaggregate those with Israeli-born parents into Ashkenazi and Mizrahi subgroups. We do not do so because the samples of ultra-Orthodox segments of these subgroups are quite small.

18. See Mayshar and Manski (2000) for analyses of fertility patterns that use alternatives to this criterion for identification of religiosity.

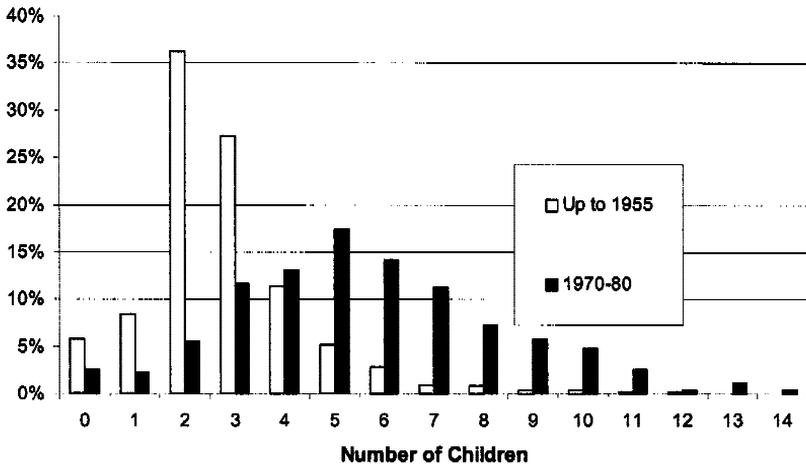


FIGURE 3. Distribution of Completed Fertility among Ultra-Orthodox Ashkenazi Women Pre-1955 and 1970–1980 Marriage Cohorts

the period around the 1975 tax reform.¹⁹ The discussion here focuses on Jewish Israelis because Arab Israelis, not having veteran status, generally did not receive significant child allowances at any point during the period of our analysis.

Figures 3 and 4 supplement the data on averages in Table 1, while concentrating on the fertility among ultra-Orthodox Ashkenazi women. In Figure 3 we present the entire distribution of completed fertility, comparing members of the pre-1955 marriage cohorts with those who first married in 1970–1980. The figure demonstrates the substantial shift rightward in the distribution. The mode increased from two to five children. The percentage of women with three children or less fell from 78 percent to 23 percent.

Figure 4 moves away from completed fertility and compares the kernel-smoothed, age-specific fertility rates of ultra-Orthodox Ashkenazi women in the five-year periods that just straddle the 1970–1975 allowances reform: 1965–1969 and 1975–1979. For women of age 21 and below, the age-specific rates are essentially identical in the two periods. However, the rates for women of age 22 and above are persistently higher in the latter period, indicating a marked change in fertility behavior.

The final descriptive evidence that we present in Figure 5 shows kernel-smoothed completed fertility rates by year of marriage, for each of the four subgroups of Israeli Jews whose parents were not born in Israel. The overall message of Figure 5 is surprisingly simple to describe, albeit not to explain. During 1945–1950, completed fertility varied by ethnicity, with Mizrahi women

19. The sample of ultra-Orthodox Jews with Israeli born parents is too small to support the analyses to be discussed, so we now focus on the ultra-Orthodox Ashkenazi Jews.

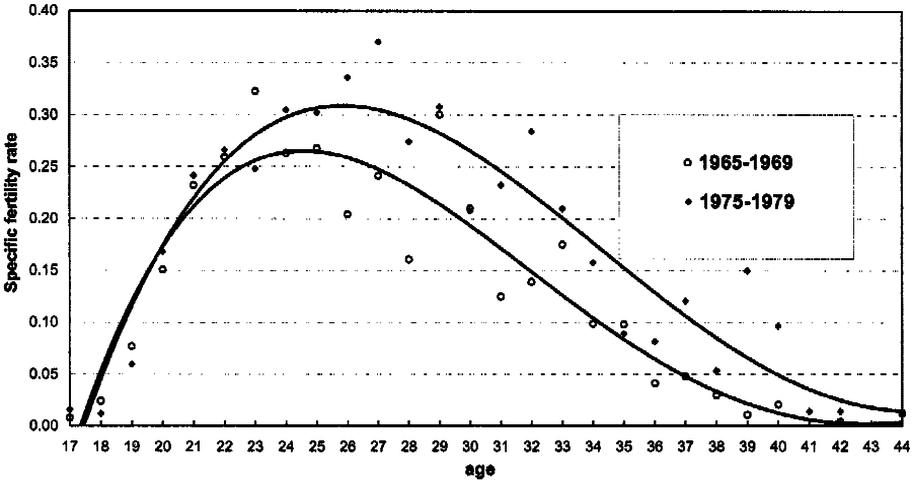


FIGURE 4. Comparison of Age-Specific Fertility Rates for Ultra-Orthodox Ashkenazi Women: 1965-1969 versus 1975-1979

bearing about three times as many children as did Ashkenazi women. Thirty years later, during 1975-1980, completed fertility again separates the population into two groups, but now the relevant dimension is religiosity rather than ethnicity, with ultra-Orthodox women bearing about twice as many children as

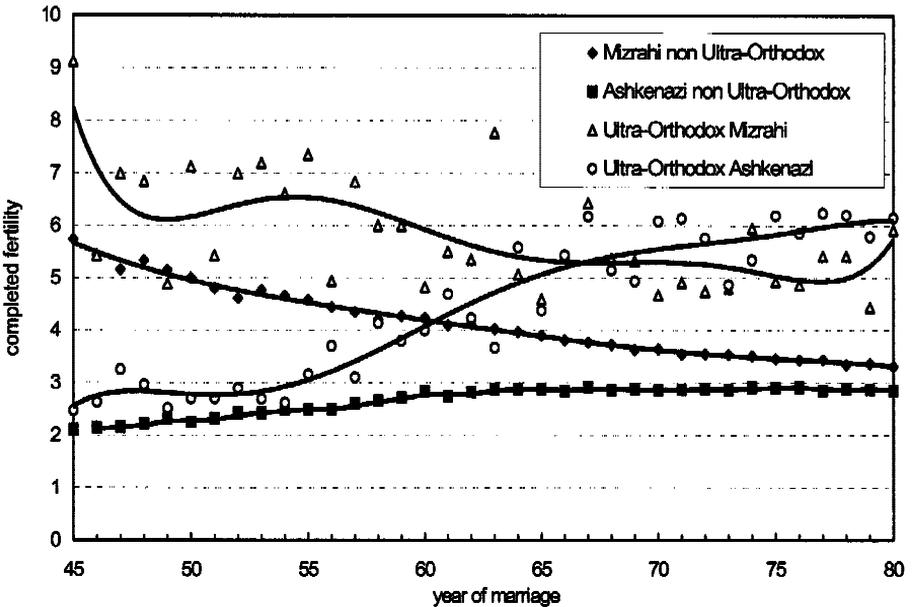


FIGURE 5. Completed Fertility among Jewish Subgroups, by Year of Marriage

non-ultra-Orthodox Jewish women. Such large, group-specific changes in fertility present an intriguing puzzle, or set of puzzles.

3. Family-Size Decisions: Private Preferences, Child Allowances, and Social Interactions

To do more than speculate about the fertility patterns described in Section 2, it seems essential to model the family decision processes that generate the population statistics. Combining basic elements of economic and sociological thinking, we think it plausible to suppose that women choose family size to maximize utility functions that recognize both private and social incentives for fertility. The private incentives depend in part on government policy toward children, one aspect of which is the child allowance program. The social incentives depend in part on the childbearing behavior of peers.

We pose a decision model with these properties here. The model aims to balance realism and tractability to provide a suitable vehicle for analysis of fertility decisions.

3.1 *Maintained Assumptions*

We assume that a woman chooses family size irrevocably sometime after marriage, based on the information available at the time. Thus we abstract from revisions to family planning that may occur if new information later becomes available. Let t denote the date of the fertility decision, j denote a woman choosing family size (i.e., number of children) at that date, and k denote family size. We assume that woman j assigns utility $U_{ij}(k)$ to having k children and acts to maximize utility. Thus the woman chooses

$$k_{ij} = \operatorname{argmax}_{k=0,1,\dots} U_{ij}(k) \quad (1)$$

children, assuming that the argmax is unique. We assume no infertility and no child mortality, or similar sources of uncertainty that could prevent some women from realizing their family-size plans. We also take the spacing of the k births as predetermined.

We now further assume that the preceding utility function has three additively separable components, the first two expressing private incentives and the third expressing social incentives:

- (a) a private utility of family size, assumed quadratic in the number of children. This is $a_j k - b_j k^2$, where (a_j, b_j) are woman-specific parameters that may vary with covariates such as age and schooling. We assume that $a_j > 0$ and $b_j > 0$.
- (b) a government incentive for childbearing embodied in the child allowances formula prevailing at the time of the fertility decision. This is $A_{ij}(k)$, the discounted life-cycle allowance that woman j would receive under the formula prevailing at the decision date, if j were to have k children; we assume that $A_{ij}(0) = 0$. The magnitude of the allowance

may vary with characteristics of the woman, such as veteran status. Let I_{ij} denote the family's expected life-cycle income from date t on; here we assume, for simplicity, that income does not depend on family size. Then expected life-cycle consumption if woman j were to have k children is $A_{ij}(k) + I_{ij}$. The contribution to utility of such consumption is $c_j[A_{ij}(k) + I_{ij}]$, where $c_j > 0$ is a woman-specific parameter.

- (c) a social interaction wherein the woman compares k with the actual family sizes of other women in the population. We assume that the population is composed of M mutually exclusive and exhaustive strata, or reference groups. Consider group m . Woman j incurs a utility loss that grows with the squared deviation between k and the size of each family in group m . The woman averages these losses across the members of m , yielding the average loss $\sum_h(k - h)^2 P_{tm}(h)$, where $P_{tm}(h)$ is the fraction of women in group m having h children. Finally, the woman aggregates the average losses across the M groups, yielding the aggregate loss $-\sum_m w_{mj}[\sum_h(k - h)^2 P_{tm}(h)]$, where $w_j \equiv (w_{mj}, m = 1, \dots, M)$ are parameters that measure the utility weight the woman attaches to each group. We assume that $w_{mj} \geq 0, m = 1, \dots, M$.

Combining the private, government, and social determinants, the form of the utility function that we assume is

$$U_{ij}(k) = a_j k - b_j k^2 + c_j[A_{ij}(k) + I_{ij}] - \sum_m w_{mj}[\sum_h(k - h)^2 P_{tm}(h)]. \quad (2)$$

3.2 Optimal Family Size

Let $S_{tm} = \sum_h h P_{tm}(h)$ be the mean number of children in group m at date t . For each m ,

$$\begin{aligned} \sum_h(k - h)^2 P_{tm}(h) &= \sum_h[(k - S_{tm}) + (S_{tm} - h)]^2 P_{tm}(h) \\ &= (k - S_{tm})^2 + \sum_h(h - S_{tm})^2 P_{tm}(h) \\ &= k^2 - 2kS_{tm} + S_{tm}^2 + \sum_h(h - S_{tm})^2 P_{tm}(h). \end{aligned} \quad (3)$$

It follows that

$$\begin{aligned} U_{ij}(k) &= U_{ij}(0) + a_j k - b_j k^2 + c_j A_{ij}(k) - (\sum_m w_{mj})k^2 + 2(\sum_m w_{mj} S_{tm})k \\ &= U_{ij}(0) + (a_j + 2\sum_m w_{mj} S_{tm})k - (b_j + \sum_m w_{mj})k^2 + c_j A_{ij}(k). \end{aligned} \quad (4)$$

Hence

$$k_{ij} = \operatorname{argmax}_{k=0,1,\dots} (a_j + 2\sum_m w_{mj} S_{tm})k - (b_j + \sum_m w_{mj})k^2 + c_j A_{ij}(k). \quad (5)$$

The model yielding equation (5) expresses many core features of economic and sociological thinking about fertility decisions. The model acknowledges that family size is a discrete choice problem and recognizes that women may have heterogeneous preferences, expressed through the woman-specific parameters (a, b, c, w).

The model flexibly expresses the idea that social norms may affect fertility. We suppose that each woman may pay attention to childbearing in multiple reference groups and place different weight on the practices of each group. A woman may be concerned with the entire distribution of family size within each group. Yet, as shown in Equations (3) and (4), the model has the simplifying property that decision making ultimately depends only on mean family size in each group, not the entire distribution.

The model yields an especially simple solution if the child allowance is proportional to the number of children in a family. Thus, suppose that $A_{ij}(k) = \pi_{ij}k$ for some $\pi_{ij} \geq 0$. Then (5) becomes

$$k_{ij} = \operatorname{argmax}_{k=0,1,\dots} (a_j + 2\sum_m w_{mj} S_{tm} + c_j \pi_{ij})k - (b_j + \sum_m w_{mj})k^2. \quad (6)$$

Observe that both composite coefficients are positive. Hence utility is quadratic in family size with its maximum at

$$k_{ij} = \operatorname{INT} \left[\frac{a_j + 2\sum_m w_{mj} S_{tm} + c_j \pi_{ij}}{2(b_j + \sum_m w_{mj})} + \frac{1}{2} \right], \quad (7)$$

where $\operatorname{INT}[\cdot]$ denotes the integer component of the real number in brackets.

Of course, the model does not express all potentially important features of fertility decisions. In particular, it neglects associations between fertility and family income that form part of the Malthusian hypothesis. It ignores possible connections between family-size decisions and female labor-supply decisions, which would make life-cycle income a function of family size. It also ignores the possible contributions of child labor to life-cycle income. Nevertheless, we feel that Equation (5) provides an effective starting point for analysis of the interaction of private and social incentives for fertility.

Section 3.3 approaches the problem from a theoretical perspective, making further idealizations in order to achieve some understanding of the dynamics and steady-state distributions of family size that may be generated by the model. Section 4 examines the problem from an empirical perspective, with particular attention to the difficulty of identification of private and social incentives.

3.3 The Social Dynamics of Fertility in Two Special Cases

The fertility model introduced here, although simple in many respects, is still too complex to permit much in the way of theoretical analysis. An interesting theoretical analysis becomes possible if we make four further assumptions, as follows:

- (a) The number of children is continuous, rather than integer-valued.
- (b) Each woman is influenced only by women in her own group who made their choices in the preceding period.
- (c) All women in a given ethnic-religious group have the same values of utility parameters (b , c , w), so that only the parameter a_j may vary within the

group.²⁰ The within-group distribution of this parameter is continuous and time-invariant, with density $f(a)$.

- (d) The child allowance formula is piecewise linear in the number of children, with one kink. The allowance formula is time-invariant and is the same for all members of a given group.

Given these assumptions, we can consider each group in isolation and characterize the evolution of family size and the steady state distribution of family sizes. We begin with the special case in which there is no kink, so the allowance is proportional to family size.

3.3.1 Allowances Proportional to Family Size. Suppose that the allowance formula is $A(k) = \pi k$, where $\pi \geq 0$ is the time-invariant allowance per child. Supposing that children are continuous and invoking the preceding additional assumptions, Equation (7) for the optimal family size for woman j who chooses family size at date t becomes:

$$k_{ij} = \frac{a_j + 2wS_{t-1} + c\pi}{2(b + w)}, \quad (8)$$

where S_{t-1} is the mean family size of women who chose their family size at date $t - 1$. It follows that the mean family size of women who choose at date t is

$$S_t = S(S_{t-1}, \lambda) = E\left[\frac{a + 2wS_{t-1} + c\pi}{2(b + w)}\right] = \frac{\mu + 2wS_{t-1} + c\pi}{2(b + w)}, \quad (9)$$

where $\mu \equiv E(a) > 0$, and where $\lambda \equiv [b, c, w, A(\cdot), f(\cdot)]$ is the full set of parameters that characterize preferences and child allowances.²¹

Equation (9) determines the social dynamics of mean family size. With $b > 0$ and $w \geq 0$, S_t converges monotonically to the unique steady state S^0 that solves the equation

$$S^0 = S(S^0, \lambda). \quad (10)$$

The steady state in this case of proportional allowances is:

$$S^0 = \frac{\mu + c\pi}{2b}. \quad (11)$$

This simple result is qualitatively unsurprising in many respects. The steady-state mean family size S^0 increases with the magnitude of the allowance per child π and with the parameter c measuring the contribution to utility of the

20. Letting c vary within the group does not qualitatively change the analysis, but complicates the notation somewhat. Given that behavior is unaffected by a positive multiplicative change in the four utility parameters (a, b, c, w), we could also normalize b or w . The essential part of our assumption is thus that both (b, w) are constant within the group.

21. Equation (9) shows that if allowances are proportional to family size, sufficient statistics for $A(\cdot)$ and $f(\cdot)$ are π and μ .

allowance. S^0 also varies as expected with the private determinants of utility, increasing with μ and decreasing with b .

Observe, however, that S^0 does not vary with the social interaction parameter w ; hence there is no social multiplier effect. Furthermore, the steady state is unique and the fertility gap between any two groups who differ only in the parameter μ is independent of the magnitude of the allowances. These results turn out to be peculiar to the case of proportional child allowances. We show below that when the allowance is piecewise linear, steady-state mean family size may not be unique and may depend on w in a somewhat subtle manner.

With S^0 determined, it is straightforward to characterize the entire steady-state distribution of family sizes. Inserting $S_{t-1} = S^0$ in Equation (8), we obtain that the steady-state optimal family size of a woman with private utility parameter a is

$$k^0(a, \lambda) = \frac{a + 2wS^0 + c\pi}{2(b + w)} = \frac{1}{2b} \left(\frac{ba + w\mu}{b + w} + c\pi \right). \tag{12}$$

Observe that $k^0(a, \lambda)$ is linear in a . Hence the density function for optimal family size has the same shape as $f(a)$, the density of a . Observe that the child-allowance π always shifts $k^0(a, \lambda)$ by the same amount, regardless of the value of a . Hence child allowances affect the central tendency of the distribution of family sizes but not its dispersion. Finally, observe that $k^0(a, \lambda)$ is determined by a weighted average of a woman's own utility parameter and the mean μ of this parameter in the group, with weights $b/(b + w)$ and $w/(b + w)$, respectively. Thus the social interaction parameter w , which was earlier shown not to affect the central tendency of the distribution of family size, does affect the dispersion of the distribution. All else equal, the larger the value of w , the more concentrated is the distribution around its mean S^0 .

3.3.2 Piecewise Linear Child Allowances. Now suppose that the child allowance formula has the piecewise linear form

$$A(k) = 0 \text{ if } k \leq K, \quad A(k) = \pi(k - K) \text{ if } k \geq K,$$

where $K \geq 0$ is a specified threshold family size and $\pi \geq 0$ is the time-invariant allowance per child from the K^{th} onward.²² Although the Israeli child allowance formula has changed over the years, the formula has always been reasonably well approximated by a function of this form.

22. The present analysis generalizes easily to formulas that begin with a positive allowance per child and rise in marginal value after K children; that is, to formulas of the type

$$A(k) = \pi_1 k \text{ if } k \leq K, \quad A(k) = \pi_2(k - K) \text{ if } k \geq K, \text{ where } 0 \leq \pi_1 \leq \pi_2.$$

However, the analysis does not apply to piecewise linear formulas that initially give larger allowance per child and then smaller ones; that is to formulas of the above type but with $0 \leq \pi_2 < \pi_1$. Such formulas have qualitatively different implications for family-size distributions.

Given assumptions (a)–(d) and a piecewise linear allowance formula, the optimal family size is

$$k_{ij} = \operatorname{argmax}_{k \geq 0} \{ (a_j + 2wS_{t-1})k - (b + w)k^2 + c\pi(k - K) \cdot 1[k \geq K] \}, \quad (13)$$

where $1[k \geq K]$ is the indicator function taking the value one if $k \geq K$ and zero otherwise. The objective function in (13) is quadratic everywhere except at the point $k = K$, where its first derivative jumps by $c\pi$ units. Holding $(b, w, c, \pi, K, S_{t-1})$ fixed, optimal family size can be shown to vary discontinuously with the utility parameter a_j , the single discontinuity occurring at a certain pivot value $a(S_{t-1}, \lambda)$, which is

$$a(S_{t-1}, \lambda) = 2(b + w)K - 2wS_{t-1} - c\pi/2. \quad (14)$$

In particular,

$$k_{ij} = \frac{a_j + 2wS_{t-1}}{2(b + w)} \leq K \quad \text{if } a_j \leq a(S_{t-1}, \lambda) \quad (15a)$$

$$k_{ij} = \frac{a_j + 2wS_{t-1} + c\pi}{2(b + w)} \geq K \quad \text{if } a_j \geq a(S_{t-1}, \lambda). \quad (15b)$$

A woman whose parameter a equals the pivot value is indifferent between these two solutions.

Recall that the distribution of the parameter a has been assumed continuous; hence the borderline case $a = a(S_{t-1}, \lambda)$ occurs with probability zero. It follows that the mean family size among women who choose fertility at date t is

$$\begin{aligned} S_t = S(S_{t-1}, \lambda) &= E \left[\frac{a + 2wS_{t-1}}{2(b + w)} \mid a \leq a(S_{t-1}, \lambda) \right] \cdot \operatorname{Prob}[a \leq a(S_{t-1}, \lambda)] \\ &\quad + E \left[\frac{a + 2wS_{t-1} + c\pi}{2(b + w)} \mid a \geq a(S_{t-1}, \lambda) \right] \cdot \operatorname{Prob}[a \geq a(S_{t-1}, \lambda)] \quad (16) \\ &= \frac{1}{2(b + w)} [\mu + 2wS_{t-1} + c\pi \cdot \operatorname{Prob}[a \geq a(S_{t-1}, \lambda)]]. \end{aligned}$$

Inspection of Equation (16) shows that, given any value of the parameters λ , the function $S(S_{t-1}, \lambda)$ is increasing and continuous in S_{t-1} ; hence the recursive equation $S_t = S(S_{t-1}, \lambda)$ generates a monotone time path for mean family size.²³

23. Equation (16) simplifies substantially if the support of the distribution of utility parameter, a , lies either entirely to the left or entirely to the right of $a(S_{t-1}, \lambda)$. In the former case, where $\operatorname{Prob}[a \geq a(S_{t-1}, \lambda)] = 0$, all women choose to have the family size that they would prefer in the absence of child allowances. In the latter case, where $\operatorname{Prob}[a \geq a(S_{t-1}, \lambda)] = 1$, all choose a family size above the threshold K , making the allowance effectively proportional to the number of children.

A steady state, where $S^0 = S(S^0, \lambda)$, is defined implicitly by the condition

$$S^0 = \frac{1}{2b} [\mu + c\pi \cdot \text{Prob}[a \geq a(S^0, \lambda)]]. \quad (17)$$

Inspection of (17) shows that this equation necessarily has solutions, and these satisfy the inequalities $\mu/2b < S^0 < (\mu + c\pi)/2b$. That is, the steady-state solutions with piecewise linear allowances lie between the steady state $\mu/2b$ that would prevail in the absence of allowances and the steady state $(\mu + c\pi)/2b$ that would prevail if allowances were proportional to family size.

Implicit differentiation of Equation (17) reveals that $\partial S^0/\partial \pi > 0$ and $\partial S^0/\partial K < 0$. Thus, as should be expected, a more generous child allowance formula generates increases in steady-state mean family size. When $\pi > 0$ and $K > 0$, similar differentiation also reveals that $\text{sgn}(\partial S^0/\partial w) = \text{sgn}(S^0 - K)$. Thus, stronger social interactions reinforce the deviation of steady state fertility away from the threshold K . We think this last result particularly important, and so state it as a proposition:

PROPOSITION: *Let our assumptions hold. Under a child allowance program with allowances proportional to family size, social interactions have no effect on mean steady state fertility S^0 . Under a program with allowances piecewise linear and threshold K , stronger social interactions reinforce the deviation of S^0 from K .*

This result is the key to our findings below on the differential effect of allowances on different social groups and on the possibility of multiple equilibria.

3.3.3 Findings for Specific Distributions of Utility Parameters. The above analysis partially characterizes steady-state fertility when the allowance formula is piecewise linear, but important questions remain open. In particular, we would like to know whether the steady-state equation (17) can have multiple solutions and we would like to understand the role of the social interaction parameter w in determining the steady state. Answers to these questions could shed light on the role that the piecewise linear Israeli child allowance system may have played in generating the sharply different fertility trends experienced by different ethnic-religious groups. As we have shown, a proportional allowance program, even when coupled with strong social interactions, cannot explain the reverse-fertility transition among Israeli ultra-Orthodox woman, nor why their fertility behavior diverged so much from that of other groups in the population.

To partially address the open questions, we now examine the dynamics of fertility under two hypotheses about the shape $f(a)$ of the distribution of utility parameters. First we suppose that $f(a)$ is uniform, a simple specification that yields a closed-form solution of the dynamic Equation (16). Then we suppose that $f(a)$ is log-normal, a specification that may be more realistic but that does not yield a closed-form solution of this equation.

Uniform Distribution: When the utility parameter a is uniformly distributed on the interval $[0, 2\mu]$, the dynamic process (16) becomes:

$$\begin{aligned}
 S_t = S(S_{t-1}, \lambda) &= \frac{1}{2(b+w)} (\mu + 2wS_{t-1}) && \text{if } a(S_{t-1}, \lambda) > 2\mu \\
 &= \frac{1}{2(b+w)} \{ \mu + 2wS_{t-1} + c\pi \cdot [2\mu - a(S_{t-1}, \lambda)]/2\mu \} && \text{if } 0 \leq a(S_{t-1}, \lambda) \leq 2\mu \\
 &= \frac{1}{2(b+w)} (\mu + 2wS_{t-1} + c\pi) && \text{if } a(S_{t-1}, \lambda) < 0.
 \end{aligned} \tag{18}$$

The function $S(S_{t-1}, \lambda)$ is continuous in S_{t-1} and piecewise linear in its three ranges.

We are interested particularly in the possibility of a steady state S^0 in the intermediate range, where some women choose to have fewer children than the allowance threshold K , and some have more than this threshold. The intermediate-range linear equation has the unique solution

$$S^0 = K + \frac{1}{2b\mu - wc\pi} [(\mu + c\pi/2)(\mu + c\pi/2 - 2bK)]. \tag{19}$$

It can be shown that this steady state is locally stable if $2b\mu > wc\pi$. Moreover, this intermediate solution satisfies the required boundary conditions $0 \leq a(S^0, \lambda) \leq 2\mu$ if $w\mu + wc\pi \leq b[2(b+w)K - c\pi/2] \leq w\mu + 2b\mu$. If the boundary conditions are not satisfied or the intermediate-range steady state is not stable, the dynamic process converges to one of two polar steady states in which all women choose to have families of size less than or greater than K . Inspection of (19) shows that if the intermediate steady state exists and is locally stable, then $S^0 > K$ if and only if $\mu + c\pi/2 > 2bK$.

Figure 6 displays how the steady state varies with the magnitude π of the allowance. To approximate the Israeli formula, we set the threshold at $K = 4$. We consider two groups, both with $b = 0.5$, $c = 1$, $w = 0.4$. The groups differ only in the value of the parameter μ . The first group, meant to represent the ultra-Orthodox Ashkenazi, has $\mu_1 = 2.5$; the other group has $\mu_2 = 2$. These figures were selected so that $S_1^0 = 2.5$, and $S_2^0 = 2$ in the absence of allowances; this conforms with the pre-1955 mean fertility rates reported earlier. It is evident that increasing the magnitude of the allowance π rapidly widens the fertility gap between the two groups. With $\pi = 3.5$, for example, the steady-state fertility levels of the two groups become 5.6 and 3.0, respectively, increasing the fertility gap between the two groups from 0.5 to 2.6. However, further increases in π narrow the gap.

The figure illustrates that a piecewise linear child allowance program combined with social interactions can generate fertility behavior that is qualitatively different from the linear case. It can be shown that when $w = 0$ for both groups, the fertility gap is hardly affected by the magnitude of the allowance. It

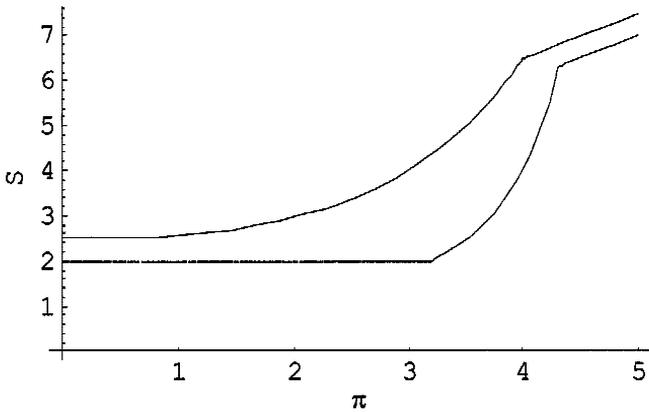


FIGURE 6. Steady-State Fertility as a Function of the Allowance for Two Groups, Where: $c = 1$, $w = 0.4$, $K = 4$, $b = 0.5$, and a Is Distributed Uniformly over $[0, 4]$ and $[0, 5]$

is the force of the social interaction that initially (up to $\pi = 3.2$) holds back women of the second group from adapting to the higher allowances by increasing their number of children to above the threshold $K = 4$. And it is the same force that operates in the opposite direction for the first group.

Log-Normal Distribution: Whereas the uniform distribution implies a unique steady state for all values of π and K , a log-normal distribution may generate multiple steady states. Since we cannot explicitly solve the steady-state Equation (17) when $f(a)$ is log-normal, we resort to a simple illustrative example. Figure 7A plots the dynamic function $S_t = S(S_{t-1}, \lambda)$ for particular values of λ . A steady state, being a solution to the equation $S^0 = S(S^0, \lambda)$, occurs when $S(\cdot, \lambda)$ crosses the 45°-line shown on the figure. A steady state is stable if $[\partial S(\cdot, \lambda)/\partial S_{t-1}] < 1$ when evaluated at $S_{t-1} = S^0$.

Figure 7A plots $S(S_{t-1}, \lambda)$ for four values of λ . In all four cases, $K = 4$, $c = 1$, $w = 0.5$, and $f(a)$ is log-normal with mean 2.5 and standard deviation 1, for both groups. The two social groups differ only in the parameter b , where $b_1 = 0.5$ and $b_2 = 0.625$, selected so that in the absence of allowances, the steady-state mean number of children for the two groups are 2.5 and 2.0, respectively. The figure displays three plots for the first group, for three alternative allowance values $\pi = (3.25, 3.75, 4.25)$, and one plot for the second group, for the middle allowance value: $\pi = 3.75$.

We find that when $\pi = 3.75$, the second group has a unique steady state at $S^0 = 2.4$ children. In sharp contrast, the first group has in this case three steady states, $S^0 = (3.69, 5.2, 6.25)$, of which only the first and third are stable. However, when $\pi = 3.25$ or $\pi = 4.25$, the first group has a unique steady state at $S^0 = 3.25$ and $S^0 = 6.7$ children, respectively. Thus, relatively small changes in allowances can produce very significant qualitative and quantitative changes in mean fertility.

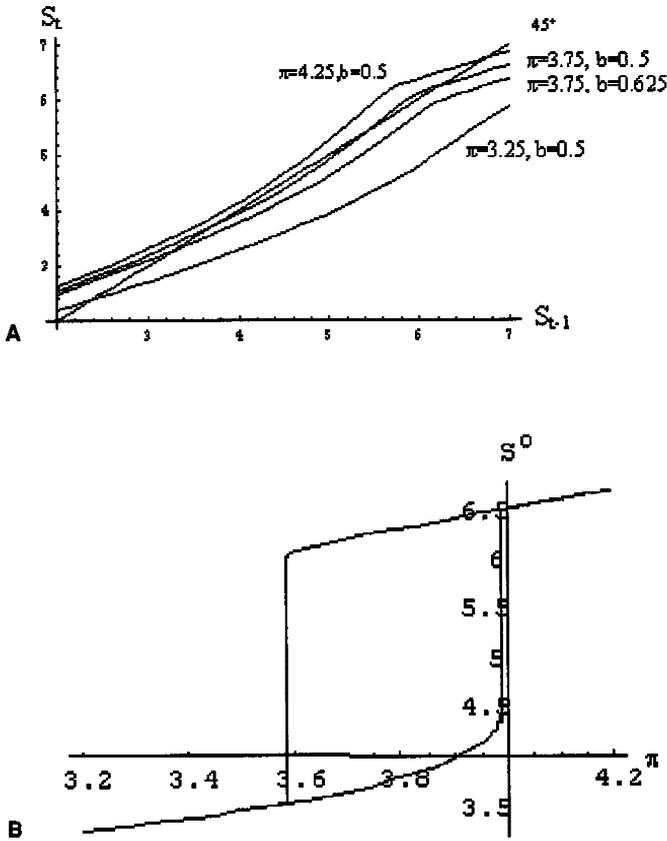


FIGURE 7A AND 7B. Dynamic Fertility Behavior and Steady State with a Log-normal Distribution

Figure 7B shows more fully how π affects the steady-state fertility level for the first group. It shows that the multiplicity of equilibria appears only when $3.59 < \pi < 3.99$.

4. Attempted Structural Empirical Inference

This section describes our efforts to use the fertility model developed in Section 3 to interpret the actual decisions made by Israeli women. Our objective was to estimate a version of the model that could, with some degree of credibility, be applied to examine how Israeli child allowance policy has affected family-size decisions. We were unable to achieve this objective, but the attempt was instructive nonetheless. We explain here.

4.1 Identification in Theory and Practice

Even though it is idealized in many respects, the model developed in Section 3 still has more than enough degrees of freedom to generate the main observed

features of Israeli fertility. The different rates of completed fertility observed in different ethnic-religious groups at different times could be produced by

- (1) cross-sectional and time-series variation in the woman-specific utility parameters
- (2) cross-sectional and time-series variation in the child allowance formula
- (3) social interactions within and across groups.

Indeed, the central problem that arises in empirical analysis of family-size decisions is that many alternative combinations of these forces could have generated the observed patterns.

The basic identification problem is evident in Equation (5), which shows that observation of a woman’s chosen family size only partly reveals her preferences. Observation that woman j chooses family size k_{ij} implies only that the composite utility parameters $(a_j + 2\sum_m w_{mj} S_{tm}, b_j + \sum_m w_{mj}, c_j)$ satisfy the revealed preference inequalities

$$(a_j + 2\sum_m w_{mj} S_{tm})(k_{ij} - k) - (b_j + \sum_m w_{mj})(k_{ij}^2 - k^2) + c_j[A_{ij}(k_{ij}) - A_{ij}(k)] \geq 0, k \neq k_{ij}. \tag{20}$$

The identification problem is especially stark when the child allowance formula provides benefits proportional to the number of children. In that case, Equation (7) shows that choice of k_{ij} children is consistent with any value of the parameters $(a_j, b_j, c_j; w_{mj}, m \in M)$ such that

$$k_{ij} - \frac{1}{2} < \frac{a_j + 2\sum_m w_{mj} S_{tm} + c_j \pi_{ij}}{2(b_j + \sum_m w_{mj})} < k_{ij} + \frac{1}{2}. \tag{21}$$

Thus, the fertility preferences of Israeli women cannot be learned from observation of choice behavior alone. The available data on fertility choices must be combined with a priori restrictions on the distribution of preferences.^{24,25}

24. There are further difficulties that exacerbate the identification problem inherent in revealed preference analysis. Our discussion has presumed that a researcher is able to observe not only the family sizes that women choose but also the child allowance formulas that they face and the mean family sizes of the women in their reference groups. These are nontrivial problems in practice. Consider the child allowance formula. Our fertility model assumes that women choose family size irrevocably sometime after marriage, based on the information available at the time. As discussed in Section 2, Israeli child-allowance benefits have varied over time. Hence the particular timing that we assume for family-size decisions affects the allowance functions that we assume women use when computing the utility of alternative numbers of children. Moreover, throughout its existence, the Israeli child allowance program has employed a formula that predicates benefits on the number of children currently under age 18. This feature of the program, as well as the variation of the formula over time, implies that the life-cycle magnitude of the child allowance associated with any number of children actually depends on the spacing of their births.

25. In principle, the identification problem may be ameliorated by enriching the available data. Specifically, Manski (2000) recommends collection of subjective data eliciting the preferences and/or expectations of decision makers. Not having such subjective data, we followed the traditional econometric practice of imposing a priori restrictions on the distribution of preferences.

In an attempt to cope with the identification problem, we refined the model in a manner that combined realism with tractability. To specify the timing of fertility decisions, we supposed that women choose family size six to ten years following marriage, and base their utility calculations on the average annual benefits in effect during that five-year period. This conjecture flows from the idea that, although women may begin to bear children soon after marriage, they need not decide so soon on completed family size.

We defined reference groups by the same criteria of ethnic origin and religiosity as were used in our description of fertility patterns in Israel in Section 2.3; moreover, we distinguished members of these groups by their year of marriage. We focused on five ethnic-religious groups: ultra-Orthodox Mizrahi, non-ultra-Orthodox Mizrahi, ultra-Orthodox Ashkenazi, non-ultra-Orthodox Ashkenazi, and non-ultra-Orthodox with Israeli-born parents. We neglected the ultra-Orthodox with Israeli-born parents because the available sample of such persons is too small to support empirical analysis. We neglected the two Arab groups because their members generally did not receive substantial child allowances during the period of our analysis.

To simplify the analysis, we constrained the vector $(w_m, m \in M)$ of social interaction parameters in two important ways. First, we assumed that women who marry in a given year are influenced only by women who married in the five preceding years. Second, we assumed that they give equal weight, say w_0 , to the fertility of all Jewish Israeli women married in the past five years and an additional weight, say w_1 , to members of their own ethnic-religious group.

After considering a variety of specifications for the distribution of fertility preferences, we chose to work with a specification that is flexible in some respects but restrictive in others. We permitted the distribution of preferences to vary freely across ethnic-religious groups and year of marriage. Thus “ultra-Orthodox Ashkenazi women married in 1955” could have one distribution of fertility preferences and “non-ultra-Orthodox Jews whose parents were born in Israel and who married in 1970” could have another. The flexibility afforded by this aspect of the model is appealing because there is reason to think that fertility preferences in Israel have varied both across groups and over time.²⁶

The restrictive part of the specification is that we assumed all women who belong to the same ethnic-religious group and who marry in the same year to have the same values of the utility parameters (b, c, w_0, w_1) . We permitted heterogeneity only in the utility parameter a , which we assumed to be distributed normal across women who belong to the same group and marry in the same year. This aspect of the specification was chosen for reasons of tractability; it implies that a simple ordered-probit model describes the probability that a woman choose any given family size.

26. Social scientists often hypothesize that fertility preferences vary with women's schooling, family earnings capacity, and perspectives on the morality of various forms of contraception. All of these factors, and more, have varied both across ethnic-religious groups and over time in Israel.

These stated assumptions, as strong as they are, only identify the model of fertility decisions if the child allowance is nonlinear in the number of children. They do not identify the model if the benefit is proportional to family size. Consider women who belong to a given reference group, say g , who marry in a given year, say t . For such women, let the parameter a be distributed $N(\mu_{tg}, \sigma_{tg}^2)$. Let S_{tg} be the mean family size of members of group g who married in the past five years and let S_{t0} be the mean family size of all other Jewish Israeli women who married in this period. If the benefit is proportional to family size and is the same for all women who marry in year t , (21) implies that the probability of choosing k children is

$$\begin{aligned} P_{tg}(k; \gamma_{tg}, \delta_{tg}) &= \text{Prob} \left[k - \frac{1}{2} < \frac{a + 2w_{0tg}S_{t0} + 2w_{1tg}S_{tg} + c_{tg}\pi_t}{2(b_{tg} + w_{0tg} + w_{1tg})} < k + \frac{1}{2} \right] \\ &= \Phi \left[\left(k + \frac{1}{2} \right) \gamma_{tg} - \delta_{tg} \right] - \Phi \left[\left(k - \frac{1}{2} \right) \gamma_{tg} - \delta_{tg} \right], \end{aligned} \quad (22)$$

where

$$\gamma_{tg} \equiv 2(b_{tg} + w_{0tg} + w_{1tg})/\sigma_{tg} \quad (23a)$$

$$\delta_{tg} \equiv (\mu_{tg} + 2w_{0tg}S_{t0} + 2w_{1tg}S_{tg} + c_{tg}\pi_t)/\sigma_{tg}. \quad (23b)$$

Equation (22) shows that the only identified quantities are the composite parameters γ_{tg} and δ_{tg} . Thus, a necessary condition for identification of the parameters $(\mu_{tg}, \sigma_{tg}, b_{tg}, c_{tg}, w_{0tg}, w_{1tg})$ describing the distribution of utility is that the benefit be nonlinear in family size.²⁷

The Israeli child allowance benefit has, to varying degrees over the years, been piecewise-linear in family size. Nevertheless, when we attempted to estimate the model, we learned that piecewise-linearity of the formula provides an insubstantial foundation for empirical analysis. In particular, we found that the objective functions used in either maximum likelihood or least squares estimation of the ordered-probit model were close to flat when evaluated at alternative parameter values that would formally be observationally equivalent if the allowance benefit were proportional to number of children. Hence, under the maintained assumptions, we were unable to obtain reliable, stable estimates of the parameters describing the distribution of utility.

Although unpleasant, this empirical finding is instructive. It cautions against reliance on nonlinearities in child allowance formulas to identify fertility preferences. The difficulties that we encountered are reminiscent of those that have long afflicted attempts by empirical researchers to learn the

27. One of these parameters may be fixed in value to normalize the scale of women's utility functions. We do not choose a particular normalization here.

effects on labor supply of benefit/tax schedules that are nonlinear in hours worked.²⁸

4.2 Structural Interpretation of “Approximate” Reduced-form Estimates

Although the assumptions imposed here did not yield usable estimates of the structural parameters (μ_{tg} , σ_{tg} , b_{tg} , c_{tg} , w_{0tg} , w_{1tg}), we recognized that strengthening the assumptions could, in principle, make estimation possible. A simple, flexible way to explore the implications of adding further assumptions is to estimate the composite parameters (γ_{tg} , δ_{tg}) for the various reference groups g and years t . This done, we could attempt to structurally interpret the cross-sectional and time-series variation in these estimates, via equation (23).

For each reference group g and year of marriage t , our merged Census and Birth Registry data make it straightforward to estimate (γ_{tg} , δ_{tg}). In years when the child allowance benefit was proportional to family size, Equation (22) follows directly from our structural model of fertility decision. Hence (22) is a reduced-form of the structural model and (γ_{tg} , δ_{tg}) are reduced-form parameters related to the structural parameters (μ_{tg} , σ_{tg} , b_{tg} , c_{tg} , w_{0tg} , w_{1tg}) by Equation (23). When the benefit was piecewise-linear in family size, Equation (22) does not follow directly from the fertility model; hence (γ_{tg} , δ_{tg}) are not formal reduced-form parameters. Nevertheless, it may be reasonable to view (γ_{tg} , δ_{tg}) informally as “approximate” reduced-form parameters, with π_t defined to be the allowance per child averaged over a suitable range of family sizes.

We used a least-squares criterion to estimate (γ_{tg} , δ_{tg}) for group g and year of marriage t . Specifically, we chose (γ , δ) to minimize the criterion function $\sum_k [f_{tg}(k) - P_{tg}(k; \gamma, \delta)]^2$; here $f_{tg}(k)$ is the observed frequency with which women of type (t , g) choose k children and $P_{tg}(k; \gamma, \delta)$ is the choice probability given in (22).²⁹ We found that our estimated choice probabilities fit the observed family-size frequencies very well. We achieved almost perfect fits for the non-Orthodox groups, where our observed choice frequencies are based on large samples of women per year; in these cases, the value of the criterion function (i.e., the sum of square residuals) was generally in the range [0.001, 0.01]. The fits were less exact but still good for the Orthodox groups, where the observed choice frequencies are based on samples of only ten to fifty women per year; here the sum of squared residuals was generally in the range [0.01, 0.05].

With estimates of the reduced-form parameters in hand, we used Equation (23) to interpret their cross-sectional and time-series variation. We first explored specifications that impose no cross-group restrictions on the distribution of

28. See Moffitt (1992) and the articles collected in the *Journal of Human Resources* Special Issue on Taxation and Labor Supply in Industrial Countries, Volume 25, No. 3, 1990.

29. If Equation (22) was a formal reduced-form throughout the period of analysis, conventional considerations of asymptotic efficiency would suggest maximum likelihood as an alternative to least-squares estimation. However, with (22) sometimes being only an approximate reduced-form, least-squares estimation seemed advantageous because it does not emphasize a model’s ability to predict small probability events, as maximum likelihood does.

utility, but do constrain the time-series variation in each group's preferences. The findings were disappointing. Estimates of the structural parameters were statistically imprecise. Contrary to the basic theoretical prediction that utility should rise with increasing child benefits, we often obtained negative estimates of the parameter $c_{t\mathbf{g}}$.

A qualitatively more "sensible" estimate of the effect of child allowances on fertility emerged only when we imposed two assumptions that we find much too strong to be believable: (a) the parameters ($\sigma_{t\mathbf{g}}$, $c_{t\mathbf{g}}$, $w_{0t\mathbf{g}}$, $w_{1t\mathbf{g}}$) are constant over time and across groups and (b) the parameter $\mu_{t\mathbf{g}}$, which describes the central tendency in private preferences for fertility, varies over time only with the rate of high school completion by women in group \mathbf{g} .³⁰ Specifically, we set $\sigma_{t\mathbf{g}} = 1$ to normalize the scale of utility and let equation (23a) have the form

$$\delta_{t\mathbf{g}} \equiv (\mu_{\mathbf{g}} + \beta H_{t\mathbf{g}} + 2w_0 S_{t0} + 2w_1 S_{t\mathbf{g}} + c\pi_t), \quad (24)$$

where $H_{t\mathbf{g}}$ is the fraction of women of type (t, \mathbf{g}) who complete high school and β is a parameter. This done, we used our estimates of $\delta_{t\mathbf{g}}$ to obtain least-squares estimates of the common structural parameters (β_δ , w_0 , w_1 , c) and the group-specific parameters $\mu_{\mathbf{g}}$. Now, finally, the estimate of c was positive and statistically precise by conventional standards.

5. Conclusions

To reiterate the main aspects of our work, we first called attention to the complex time-series and cross-sectional pattern of Israeli fertility, which is much at odds with the fertility transition that has occurred elsewhere in the world (Section 2). We next presented a theoretical analysis of the social dynamics of fertility that shows how private preferences, preferences for conformity, and piecewise linear child allowances could have combined to yield such a complex fertility pattern (Section 3). We then explained the identification problem that makes it so difficult to infer the actual Israeli fertility process from data on completed fertility (Section 4).

Although much about Israeli fertility remains puzzling at the end of our investigation, we believe that our work may cast some new light on what has happened in the ultra-Orthodox community. In a recent study of fertility and labor supply in the Israeli ultra-Orthodox community, Berman (2000) has interpreted the trends toward increased fertility, decreased labor supply, and increased supply of time to religious studies in this community as the behavior of a "club" that has strengthened its norms of religious stringency in an attempt

30. Demographers have long viewed women's schooling as a primary determinant of fertility preferences. Israeli women experienced very substantial increases in their rates of high school completion during the period covered by our data. Hence we thought it reasonable to focus on this as a source of time-series variation in fertility preferences. There are two levels of high school completion in Israel, a lower level in which a student obtains a diploma and a higher level in which a student also passes a national examination known as the Bagrut. We measure high school completion by successful completion of the Bagrut.

to maintain exclusion. Berman traces this increased stringency back to the nineteenth century reaction of ultra-religious Jews to their exposure to modernity, and he implicitly views recent changes as continuation of a process that has been underway since that early time. Our descriptive empirical analysis casts doubt on Berman's interpretation. We have found that in the marriage cohorts prior to 1955, the fertility of Ashkenazi ultra-Orthodox women was almost the same as that of non-ultra Orthodox Ashkenazi women. Thus, the reverse fertility transition among the Israeli ultra-Orthodox is a relatively recent phenomenon.³¹ We conjecture, but cannot prove, that the reverse fertility transition that has occurred among the ultra-Orthodox would not have happened were it not for the emergence of the Israeli welfare state.

Our work may also help to understand the conventional fertility transition. To recall, demographers have questioned why the fertility transition has begun in different societies under different economic conditions, and why it has been so rapid. It is intriguing that the reverse fertility transition experienced by the Israeli ultra-Orthodox has similarly been rapid. We have conjectured that this reverse transition arose out of the combination of two key factors: noncontinuity in private behavior, due to the piecewise linearity of child allowances, and social interactions as each woman is influenced by the fertility decisions of other women. The same two factors are central to the recent analysis by Kohler (2001) of the conventional fertility transition. In his work, noncontinuity in behavior is generated by a binary decision process, whether or not to use contraceptives, and the social interaction occurs when each woman is influenced by the contraception decisions of other women. The specifics of our research and that of Kohler differ in many respects, but each demonstrates how noncontinuity in behavior and social interactions may combine to explain the social dynamics of fertility.

References

- Andorka, Rudolf (1978). *Determinants of Fertility in Advanced Societies*. London: Methuen.
- Becker, Gary S. and Murphy, Kevin M. (2000). *Social Economics: Market Behavior in Social Environments*. Cambridge, Massachusetts: Harvard University Press.
- Berman, Eli (2000). "Sect, Subsidy and Sacrifice: An Economist's View of Ultra-Orthodox Jews." *Quarterly Journal of Economics*, 115, pp. 905–953.
- Bongaarts, John and Watkins, Susan Cotts (1996). "Social Interactions and Contemporary Fertility Transitions." *Population and Development Review*, 22, pp. 639–682.
- Boyer, George R. (1989). "Malthus Was Right After All: Poor Relief and Birth Rates in Southern England." *Journal of Political Economy*, 97, pp. 93–114.
- Brock, William and Durlauf, Steven (2001). "Discrete Choice with Social Interactions." *Review of Economic Studies*, 68, pp. 235–260.

31. The recent noncontinuity in behavior is illustrated by a quote from an address by a noted Rabbi to young women in an ultra-Orthodox seminary: "Our mothers . . . have absorbed too much from the odor of European culture . . . Our slogan must be: 'back to grandmother'" [cited in Friedman (1988)].

- Cleland, John and Wilson, Christopher (1987). "Demand Theories of Fertility Transition: An Iconoclastic View." *Population Studies*, 41, pp. 5–30.
- Dahan, Momi (1998). "The Ultra-Orthodox Jews and Municipal Authority." Part I, The Jerusalem Institute for Israel Studies, Jerusalem, Research Series No. 79 (in Hebrew).
- Doron, Abraham and Kramer, Ralph M. (1991). *The Welfare State in Israel: The Evolution of Social Security Policy and Practice*. Boulder, Colorado: Westview Press.
- Duncan-Jones, Richard (1974). *The Economy of the Roman Empire: Quantitative Studies*. Cambridge, United Kingdom: Cambridge University Press.
- Friedman, Menachem (1988). "Back to Grandmother: The New Ultra-Orthodox Woman." *Israel Studies*, Spring, pp. 21–26.
- Friedlander, Dov and Feldmann, Carole (1993). "The Modern Shift to Below-Replacement Fertility: Has Israel's Population Joined the Process?" *Population Studies*, 47, pp. 295–306.
- Gauthier, Anne Helene and Hatzius, Jan (1997). "Family Benefits and Fertility: An Econometric Analysis." *Population Studies*, November, 51, pp. 295–306.
- Hammel, E. A. (1990). "A Theory of Culture for Demography." *Population and Development Review*, 16, pp. 455–485.
- Hotz, V. Joseph and Miller, Robert A. (1988). "An Empirical Analysis of Life Cycle Fertility and Female Labor Supply." *Econometrica*, 56, pp. 91–118.
- Kohler, Hans-Peter (2001). *Fertility and Social Interaction: An Economic Perspective*. Oxford, United Kingdom: Oxford University Press.
- Malthus, Thomas Robert (1996). *An Essay on the Principle of Population* (1798), Reprint of 1st ed. with a new introduction by Samuel Hollander. London: Routledge.
- Manski, Charles F. (1993). "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60, pp. 531–542.
- Manski, Charles F. (2000). "Economic Analysis of Social Interactions." *Journal of Economic Perspectives*, 14, pp. 115–136.
- Mayshar, Joram and Manski, Charles F. (2000). "Child Allowances and Fertility in Israel: Preliminary Findings." *The Economic Quarterly*, 47, pp. 535–565 (in Hebrew).
- Moffitt, Robert (1992). "Incentive Effects of the U.S. Welfare System: A Review." *Journal of Economic Literature*, 30, pp. 1–61.
- The National Insurance Institute, *Quarterly Statistics and Annual Reports*, Jerusalem, various issues.
- Social Security Administration, U.S. Department of Health and Human Services (1999). *Social Security Programs Throughout the World—1999*. Washington, D.C.: U.S. Government Printing Office.
- State of Israel, Committee for Natality Problems (1966). *Report Submitted to the Prime Minister*, Jerusalem (in Hebrew).
- State of Israel, Commission on Tax Reform (1975). *Proposals for Changing the Direct Tax System*, translation of Hebrew original, Jerusalem.
- Teitelbaum, Michael S. and Winter, Jay M. (1985). *The Fear of Population Decline*. Orlando, Florida: Academic Press.
- Walker, James R. (1995). "The Effect of Public Policies on Recent Swedish Fertility Behavior." *Journal of Population Economics*, 8, pp. 223–251.
- Watson, Cicley (1954). "Population Policy in France: Family Allowances and Other Benefits." *Population Studies*, 7, March, pp. 263–286, and July, pp. 46–73.

Copyright of Journal of the European Economic Association is the property of MIT Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.