WORLD FERTILITY SURVEY

TECHNICAL BULLETINS

SEPTEMBER 1982

NO. 8

Progressive Fertility Analysis

NORMAN B. RYDER

INTERNATIONAL STATISTICAL INSTITUTE
Permanent Office. Director: E. Lunenberg
428 Prinses Beatrixlaan, PO Box 950
2270 AZ Voorburg
Netherlands

WORLD FERTILITY SURVEY
Project Director:
Halvor Gille
35–37 Grosvenor Gardens
London SW1W 0BS, UK

The World Fertility Survey is an international research programme whose purpose is to assess the current state of human fertility throughout the world. This is being done principally through promoting and supporting nationally representative, internationally comparable, and scientifically designed and conducted sample surveys of fertility behaviour in as many countries as possible.

The WFS is being undertaken, with the collaboration of the United Nations, by the International Statistical Institute in cooperation with the International Union for the Scientific Study of Population. Financial support is provided principally by the United Nations Fund for Population Activities and the United States Agency for International Development.

This paper is one of a series of Technical Bulletins recommended by the WFS Technical Advisory Committee to supplement the document *Strategies for the Analysis of WFS Data* and which deal with specific methodological problems of analysis beyond the Country Report No. 1. Their circulation is restricted to people involved in the analysis of WFS data, to the WFS depositary libraries and to certain other libraries. Further information and a list of these libraries may be obtained by writing to the Information Office, International Statistical Institute, 428 Prinses Beatrixlaan, Voorburg, The Hague, Netherlands.

WORLD FERTILITY SURVEY

TECHNICAL BULLETINS

Progressive Fertility Analysis

NORMAN B. RYDER

Office of Population Research
University of Princeton

SEPTEMBER 1982 NO. 8

Contents

FIGURE

Ackowledgements

The author wishes to acknowledge the helpful criticisms of John Hobcraft of the WFS **staff during the** preparation of this work.

1 Styles of Measurement

1.1 TEMPORAL VARIABLES

This work is concerned with the measurement of the fertility of aggregates of individuals, especially temporal aggregates, based on survey data. The account is intended not as a comprehensive and even-handed review of measures customarily used for this purpose, but as an advocacy of a new style of measurement.

Most of the work is limited to consideration of what can be done with the elementary facts about a reproductive history: the number of births, their dates of occurrence, the dates of respondent's birth and marriage, and the date of interview. Dated information is crucial to the subject at hand, and in two distinct ways.

On the one hand, fertility is here construed not – as is often the case – as simply the number of births occurring over the reproductive span, but as a process occurring through time. Dates provide a framework for description of the process. The familiar variables based on dates are ordinarily differences between some reference time and the time of occurrence of an event of interest, eg age, marital duration, age at first marriage, and length of birth interval. In this use, the dates represent information about the passage of personal time. Although it is not uncommon to speak of the influence of, say, age on fertility, the sense of the presentation is that such primary pieces of information are to be considered as part of the definition of the subject of interest.

On the other hand, the dates associated with reproductive events are evidently also required for the study of temporal variations in fertility. In this use, they signify historical rather than personal time. A central concern of this work is the distinction between two systems of temporal comparison, two ways of writing the history. Respondents can be ordered in time by reference to the calendar periods within which their histories, or particular phases of their history, begin, such as their date of birth or date of first marriage. These exposure-initiating events are said to define cohort membership; the comparative fertility of successive cohorts of one or another kind represents one kind of history, based on what may be called the cohort mode of temporal aggregation.

Alternatively, the reproductive events occur at particular times, and measures of fertility may be devised to characterize the experience in each historical period. This is the conventional way in which data are organized for the study of change over time, by what may be called the period mode of temporal aggregation. Such measures are appropriate to the study of the consequences of reproductive change; one such measure is the most common of all, the crude birth rate. Yet they are usually constructed so that in form they resemble aspects of the reproductive history of a cohort. In such a guise they are here characterized as measures for synthetic cohorts.

The important fact about reproductive indices assembled from the same data set for a succession of real cohorts and for a succession of synthetic cohorts is that the two modes tell different stories. The nature of this difference is an important facet of the following account.

1.2 STYLES OF MEASUREMENT AND SOURCES OF DATA

Styles of measurement are responsive to types of data collection as well as to conceptualization of the phenomenon being analysed. The measurement style characteristically associated with birth registration data is the period mode of temporal aggregation. One theme of the present account is that the appropriate way to approach survey data is in the cohort mode.

In principle, conceptualization should inform the measurement style, but practical considerations may be dominant. To date, the analysis of fertility has been guided at best only implicitly by theory. The primary task has been seen as the description of the phenomenon and its correlates, and the provision of a body of information which can be responsive to the needs of different kinds of theory, rather than the collection of data designed to test particular hypotheses. One consequence of this orientation is that we are inclined to allow the form of the data to influence our choice of measure. This is not to denigrate the accomplishments of the discipline: every science has begun with a primitive stage of description.

For many decades, the fertility information available to demographers was a by-product of official enumeration and registration systems. Consider first the parity question on a census (ordinarily given the clumsy title of 'children ever born to women ever married'). This may be thought of as the beginning of the kind of information collected in a fertility survey. A woman's parity, a summary index of her lifetime reproductive experience, can be studied for its covariation with other census items.

Problems with this kind of analysis are well known.

1 The only women with complete records are those past menopause.
2 The events they are reporting are remote in time and thus subject to problems of recall.
3 The women available for enumeration are the survivors of the processes of mortality and migration, and may have been selected on criteria of relevance for what is being studied.
4 Since most of the recent reproduction is contributed by women whose histories remain incomplete at time of enumeration, and the dates of occurrence of births are not usually obtained, there is no way to make a statement about what is currently happening to fertility.
5 For the same reason, time series follow the cohort mode, not because of an analytic choice but because there is no way to reassemble the data period by period.
6 With enumerations ordinarily ten years apart, there is little scope for temporal precision concerning fertility variations.

The other major secondary source is the birth registration system. As such systems grew in coverage and reliability, they became the preferred source of information about fertility as a function of time. Since comprehensive data are provided year by year, the natural mode of temporal aggregation is the period. In order to produce indices which can be considered as summarizing personal histories, the synthetic cohort device is employed, combining the information for each stage of the life cycle (the behaviour of each successive real cohort) to make a synthetic history characterizing the fertility of the period. The synthetic cohort tactic is feasible only when the birth rates for women of all the different ages in the period in question are available.

It has proved difficult to achieve completeness of birth registration. A decline in the degree of incompleteness may be misinterpreted as a rise in fertility. Registration data provide only the numerators for birth rates; the denominators come from enumeration sources. Any measure based on data from more than one source is only reliable when the definitions in the various sources are comparable. The information about other characteristics which may be associated with fertility is restricted to the contents of the birth registration form, and that may be ungenerous with respect to socio-cultural and socio-economic information.

The registration system produces a complete set of rates over all ages (for all cohorts) in each period in which it is in operation. In some countries, as the time series of registration data on a comparable basis has lengthened, cohort fertility tables have been developed, as a re-organization of the form in which the conventional period-by-age rates are displayed. In any time series of registration data, there are many cohorts with records incomplete, either for the beginning or for the end of their reproductive span. Period indices are the obvious choice for summarizing registration data temporally, again on grounds of convenience. The registration system also yields the kind of information needed to prepare a population projection, and thus display the consequences of population change, although the form of that information is not quite right, since the desideratum is data organized on a period-by-cohort rather than period-by-age basis (on which more later, p 16).

Registration systems can yield births classified by more than age. Various aspects of the reproductive history of the mother may be identified: the number of her previous births, the length of time since she first married, and the length of time since her preceding birth. But by themselves such pieces of information are almost unusable. Responsible analysis requires a way of estimating the distribution of these aspects of all women, and not just those giving birth in the period in question. Although such information may be collected by periodic enumeration, this is rarely the case, and even then the lengthy interval between enumerations dates the information quickly. Procedures have been developed to produce the required denominators in part from registration data themselves, but successful examples are rare because a lengthy series of reliable data of the requisite detail is called for.

What are the distinctive characteristics of registration and enumeration data on fertility?

1 With registration data, one begins with the occurrence of a birth, and then determines the associated characteristics of those to whom the births occur. Thus one focuses on the set of circumstances surrounding the event. The beginning of the measurement process is the end of the experience, and the challenge is to work backwards in time, trying to reconstruct the circumstances which led to the experience. The focus of enumeration data, on the other hand, is the set of characteristics of people, including the circumstances of events, like births, which may have occurred to them.
2 With enumeration data, the population is defined as those present at the time of the census or survey. With registration data, the cohort becomes a population, subject to change through mortality and migration. In short, the definitions of the universe differ.
3 The coding conventions also differ. With registration data it is customary to code the data by age; with enumeration data, age is used to signify cohort membership.
4 The classic registration measure is the central age-specific rate, while the enumeration calculation can be a direct observation of the proportion changing status from one enumeration to the next. That is feasible in part because in this type of data the cohort is not subject to change in size or composition.
5 The configurations of the data sets differ. With registration data, one has a series of complete periods, and thus a series of incomplete cohorts, in the form of a rhombus. With enumeration data, a triangle of information results, with no difference between the availability of data by period or by cohort (see p 15).

The reason for devoting this attention to a comparison of registration and enumeration is that the survey is an elaboration of the enumeration procedure, whereas most measurement techniques have been developed on the basis of registration data. In the survey, one obtains a comprehensive record of the fertility of a woman, with the events recorded successively by date. It follows that exposure to risk of occurrence of any event is automatically displayed in the history. There is no formal problem of developing highly specific conditional probabilities of birth, because the same record provides both the numerators and the denominators. Although synthetic cohort measures of comparable kind are feasible, the procedure is generally awkward and complex, relative to the same calculations for real cohorts. The major problem with the exploitation of survey data for fertility measures has been the tendency to use the configuration of data provided by the survey as a surrogate for the kinds of data produced by registration systems. The message of this work is that straightforward and simple measures are available in survey data for real cohorts and they are preferable for many analytic purposes to those for synthetic cohorts.

1.3 DEFINITION OF THE UNIVERSE

At the outset of a fertility survey, criteria are established for determining eligibility for interview, criteria which define the universe. There is ordinarily a compromise between comprehensive coverage, an analytic desideratum in the abstract, and considerations of research economy and a realistic approach towards obtaining reliable information. The choice of criteria has important implications for measurement.

The first consideration is the obvious circumstance that those who are interviewed are members of the population at the time of interview. Comprehensive coverage of the reproductive behaviour of a population over time would encompass all those who had been in the reproductive ages at any point during the time span. In principle, the registration procedure provides such coverage. In a survey, however, those interviewed are a non-random selection of the identified population to the extent that their relevant characteristics may have been affected by the volume and selectivity of the processes of mortality and migration. Although in practice it is usually assumed, implicitly, that mortality and migration are either non-selective or that their selectivity can be ignored because the proportion of respondents affected is sufficiently small, this bias is present in all surveys. Considering the survey as evidence about the comparative behaviour of a series of birth cohorts, the relevance of the point is that the various cohorts are likely to be differently biased in this respect, since the greater the lapse of time between birth and interview (a difference intrinsic to the cohort distinction), the greater the opportunity for selection to operate.

In the second place, the interview takes place at a particular time, and therefore, from the standpoint of the cohorts being interviewed, at a different stage in their life cycles. The histories of those interviewed are said to be censored by interview. The survey can therefore provide data which is progressively less complete, the more recent the birth cohort, ie the younger it is at interview. It follows that a comparison of the records of successive cohorts must accommodate itself in some way to this circumstance in order to avoid an obvious bias. On the one hand, one can compare equivalently incomplete information for two cohorts by deleting from the more extensive record the information which is necessarily unrepresented in the less extensive record. This is the procedure generally followed in the present work, although, regrettably, some information is ignored. On the other hand, one may use the available information for each cohort as a basis for projecting a complete record. Since that requires judgement about the form of extrapolation to employ, and the assumption that the behaviour of the cohort subsequent to interview will correspond to the behaviour observed for other cohorts in other times and places, it is to be expected that different demographers will carry out the task in different ways, obviating any codification of the procedure. Important though such activities are, there would seem to be advantages in separating them from the task of describing the results of a survey.

In the third place, the survey is normally restricted to interviews with those under 50 years of age. There are three reasons for this restriction.

1 The concern about selectivity increases with advancing age; the limit represents a sensible discretion.
2 There is more interest in recent behaviour than in behaviour more remote in time. Cohorts older than age 50 at time of survey have contributed little to the fertility of the preceding decade.
3 As age increases, the length of time between events of interest and time of interview increases. It is generally agreed that the quality of information varies inversely with the length of recall required. This is another respect in which the records of successive birth cohorts may be differentially biased.

The consequence of imposing an age limit on the survey is that progressively less information in the sense of coverage of the age span, is available, the further back in time one goes. What happens is that the histories of synthetic cohorts are censored. In this case, the censoring is more severe the further one goes back in time, whereas with real cohorts, the censoring is more severe the further one goes forward in time. Synthetic cohorts face the same kind of problem of comparison as real cohorts, with respect to differential incompleteness, with the same alternative resolutions.

In the fourth place, a common practice in surveys is to collect information on all women in the household schedule, but restrict eligibility for individual interview to those who have ever been married. Although one would ideally prefer to include in a survey all women who may have been exposed to the risk of fertility, direct questions to determine which of the never-married have in fact been exposed to risk may threaten rapport. In most populations it is likely that there is sufficient correspondence between those who are married and those who are exposed to risk to make the former a good approximation of the latter. In all societies there are norms designed to achieve that outcome, although they differ in the extent to which the reality approximates the ideal. Because of those norms, it is offensive to ask a never-married woman whether she has been exposed to risk. Furthermore, it may be judged that the yield of information of interest, for those never-married who have ever been exposed to risk, is too small to justify the cost of an interview. In brief, the ever-married stipulation is a practical decision.

In surveys in which a marital status criterion is employed, there are important consequences for measurement. A survey contains answers to questions administered at a particular time to a set of individuals who satisfied the criteria of the universe definition. The answer to any question is a censored measurement for the individual – meaning a measurement conditional on the time the interview took place – if that individual, interviewed at another time, were to give a different answer (quite apart from inconsistency of response). Demographic processes are inherently time-dependent in this sense.

Considering the interview as a random event from the standpoint of an individual's history, those to whom a particular event occurs earlier in their life are more likely to be interviewed after than before that event, compared with those to whom the same event occurs later in their life. The stipulation that the members of the sample be ever-married is another way of saying that they must have experienced the event of first marriage prior to interview. It follows that the members of any birth cohort who are interviewed are those with marriages which occurred prior to their age at interview. Accordingly there is a different possible age at marriage distribution for each cohort. This bias requires careful treatment in devising fertility measures. The effect is called a bias because it originates not in the behaviour of the respondents but, in a sense, in the behaviour of the researcher, through the act of defining the universe.

Many other kinds of bias can be found in survey data, but they are not the subject of the present work. Random error arises from the circumstance that the sample of individuals observed is but one of a population of such samples which would result from repeatedly conducting the same survey in the same way. Bias may arise in the selection of individuals for the sample, perhaps because of selective non co-operation, compromises made in any randomization procedure, failure to execute instructions, and errors in determining eligibility. Non-random error may also arise in the process of producing data concerning individuals. Two sources in particular are noteworthy. The first is recall bias, a consequence of the lapse of time between the event in question and the time at which information about it is obtained. Little is known but much suspected and surmised about the extent and characteristics of this phenomenon. The subject has attracted a specialized literature. Although it is relevant to judgements about reproductive histories and the comparison of successive cohorts, it will not be discussed here. The second is a tendency for events and characteristics to be systematically misreported so that the ostensible behaviour comes closer to normative expectations than the actual behaviour. This problem is to be expected in the investigation of any behaviour which is heavily endowed with normative content, in effect any socially important behaviour. Nothing beyond this warning is contributed to the subject here.

1.4 INSTRUMENTAL AND EXPLANATORY VARIABLES

A survey ordinarily produces a wealth of information about reproductivity in much greater biometric detail than the elementary data listed at the beginning of this chapter. This is particularly so for the instrumental variables: exposure to risk, fecundability, foetal mortality, fertility regulation, lactation, and associated topics. (The term 'instrumental' is proposed as an alternative to the more common designation 'intermediate', because the latter has an unevocative quality, and is used elsewhere with other connotations.) These are properly the subject of separate specialized discussion elsewhere. The instrumental variables ideally provide a total explanation (at the biometric level) of the primary data, since they are the sole pathways along which any explanatory variable in some deeper sense can have an influence on fertility. Part of the philosophy underlying a survey is the study of fertility as a function of the instrumental variables, and the study of the instrumental variables as a function of the array of explanatory variables of socio-cultural and other kinds.

The instrumental variables are explicitly mentioned because they are investigated with a unit of analysis as the interval between one birth and the next, with respect to length of pregnancy, duration of lactation, episodes of contraceptive use, and so forth. In order to link research at this level with fertility measures based on primary data, it is important to consider a research design which yields interval-specific fertility measures.

Most of the account which follows is concentrated on the measurement of variations in fertility as a consequence of the passage of cohort or period time. While this is an important, perhaps the most important, subject for a fertility survey, many interesting questions can be asked which are not concerned with the development of a time series. One important class of characteristic to which attention is devoted in the last chapter is that which is invariant for an individual over the course of the life cycle. Any such characteristic may be thought of as defining membership in a subpopulation. The principles of measurement for such subpopulations are identical with those for measurement of the population as a whole, as discussed throughout the body of this work.

Other characteristics mark the location of an individual on one or another dimension at a particular time, but are subject to change over the course of the individual life cycle. As such, they are processes which may be adaptable to investigation by demographic procedures like those presented here, and, moreover, their study requires attention to the same problems of bias associated with the definition of the universe, and the censoring of the data configuration, outlined above. The analytic approach associated with the relationship between locational characteristics and reproductive behaviour is commonly individual-specific, whereas that associated with inquiries employing subpopulation identifications is macro-analytic. Although it is not uncommon for a survey to be regarded as an instrument for collecting information about individuals, in order to test hypotheses at the individual level (for which aggregate calculations are at best inefficient), a lot of faith is required to sustain that view in the face of the meagre results of such inquiries to date. The position taken here is that the measurement of the fertility of aggregates is not only in some sense necessary, but moreover analytically attractive in its own right.

1.5 DETERMINANTS AND CONSEQUENCES

The principal message in this introductory chapter is that the style of measurement should be adapted to the form of data collection. Although the period mode of temporal aggregation, and associated constructions, has the attractiveness of convenience with registration data, the cohort mode is the orientation of choice with enumeration (and thus survey) data. If one thinks of the survey as a vehicle for observation of the instrumental variables, it is sensible, again on grounds of convenience, to use an interval-specific design to process the primary demographic data which are the consequence of the operation of those instrumental variables.

A major source of support for fertility surveys is the premise that they can contribute to the formulation of a programme to achieve the objectives of a population policy. One element in such work is the development of alternative fertility projections. The style of measurement proposed here has the advantage of permitting the incorporation of assumptions into such projections, concerning the time pattern of fertility, and the instrumental variables, in ways that more closely correspond with the relevant behaviour patterns than is feasible with conventional projection procedures.

Underlying such assumptions is some sense of the determinants of reproductive behaviour, and orienting a fertility survey towards understanding these determinants is in the long run a practical proposal.

To those long accustomed to conventional fertility analysis, in the measurement style inherited from work with registration data, the approach taken here is unfamiliar. However, it can accomplish the same objectives as traditional measures in a less clumsy fashion, and various other valuable objectives quite out of their reach. Although the approach is certainly novel, in the sense that none of the measures proposed can be found in the literature, in precisely the form proposed here, they are the outcome of systematic application of principles of responsible measurement which are well established throughout the reaches of the discipline.

Nor are the recommendations advanced in the spirit of substituting more complex for simpler procedures. Indeed, complexity arises only at the point at which one form of data is used to calculate measures designed for application to another form of data. Nor is the position tenable that refinement of the measurement procedure is futile in the presence of data of suspect quality; whatever that quality, remediable bias should be remedied. And since the changes in fertility to be observed may often be small and subtle, there is a premium on the acuity of the measuring instrument. We propose to look at the output of a survey in its own terms, not as a substitute for something else, and recognize that what it does readily provide is more valuable than what it can only with difficulty be forced to yield.

2 Coding

The account begins with a discussion of coding procedures. Although the issues involved are elementary, they are in a sense basic. In defence of the treatment, it is asserted that most surveys have been less than efficient in exploiting their data because of inept decisions concerning one or another of the issues raised.

The first proposal is to code the time of occurrence of the events in a respondent's history not in terms of calendar years, but on a scale which uses the time of interview as the reference point. In the typical survey, interviews are conducted over a span of some months; that time span is located variously from one survey to another with respect to the months of the calendar. Since the sample criteria apply to characteristics ascertained at time of interview, they have a somewhat different temporal referent for each respondent. If one were to follow the practice of allocating events to calendar years, the most recent interval of time would constitute a fraction of a year, and a different fraction for each individual, as well as for each survey. To use the experience of the most recent calendar segment to stand for the entire year would be to risk seasonal bias, a common feature of fertility, as well as bias associated with non-random aspects of the order in which interviews happen to be conducted. Moreover, given that the sample criteria are specific to time of interview (with respect to whether the respondent has ever married and is less than age 50), the implied limits of age and marital duration for each preceding calendar period would themselves be fractional. In summarizing such partial information with respect to the life cycle, there are two unattractive alternatives: either to truncate the record at an exact age, or other life cycle boundary, and sacrifice the remaining information, or to engage in imaginative extrapolation.

The way out of these difficulties is to employ the convenient fiction that all interviews take place at one time point (analogous to census practice) and to date all events relative to time of interview, respondent by respondent. The consequence is that the data are presented for what may be called fiscal rather than calendar periods, a circumstance that can be noted, once and for all, at the beginning of the analysis of the results. In brief, the recommendation is to follow enumeration rather than registration practice in this respect. Although this means that the results will not be temporally aligned with data from registration sources, the same holds for enumeration data generally, and interpolation procedures to effect alignment are well known. The problems of comparison on other grounds – the form of the measure, the reliability of the data, the definition of the universe, and so forth – are much more difficult to resolve than the discrepancy of temporal location.

The next question concerns the coding of intervals of time. Conventional procedures have been developed to process data from registration systems. The births which occur in a period are ordinarily coded by age of mother, using exact limits for age boundaries. Now for an individual, to any degree of precision required, the age of the mother at occurrence of the birth in question is identical to (and is defined as) the difference between the time of occurrence of the birth and the time of the mother's birth. But when grouping is employed, the choice of any two variables from this triad to be coded explicitly leaves the third identified implicitly (and cut on the bias) as an unavoidable cost. The convention with registration data is to code the period of occurrence and the age explicitly, and leave the time of birth (the cohort identification) implicit and overlapping. With survey data for a cross-section of the reproductive histories of a sequence of cohorts, the practice is inefficient and destroys information.

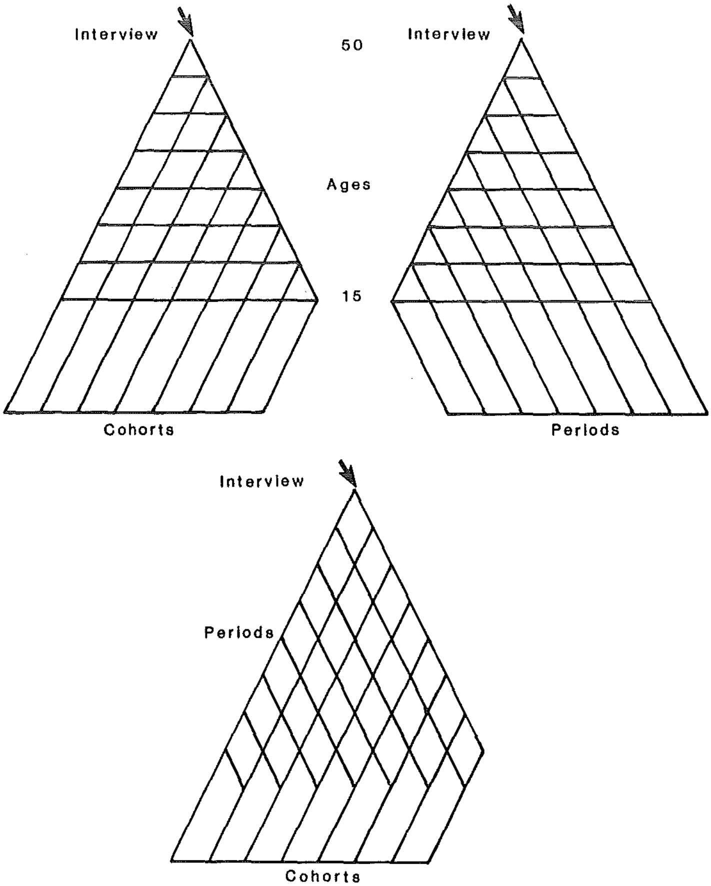

Figure 1 Survey grids based on three combinations of codes

Figure 1 elucidates the points to be made. The three panels depict the consequences of three choices of coding scheme by which to classify the data collected in a survey. The boundaries of the overall figure are established by the time of interview (the right-hand diagonal side), the upper age limit (the left-hand diagonal side), the lower age limit, and the time of the respondent's birth (the horizontal axis). This is a variant of the familiar Lexis diagram. However, as that diagram is usually employed, the relationships among three interlocked variables, such as period, cohort and age, are presented in the form of a right isosceles triangle. That necessitates the choice of which two of the three variables to show on the horizontal and vertical axes, and which to relegate to hypotenuse status. Not only is the choice arbitrary, but there is the further unfortunate consequence that one of the lengths of time is represented geometrically as different from the other two, whereas the three are the same. Both of these difficulties are obviated by the adoption of the equilateral triangle form.

In the equilateral version of the Lexis diagram (as in figure 1), the record for each respondent may be represented by a 'life' line, originating in the base at a point signifying the time of birth, and proceeding upward at a 60° angle to the point of intersection with the upper right diagonal (the time of interview). Events occurring to the respondent at particular times are signified by points on the life line. For the purposes of aggregate analysis, one may consider the number of life lines originating within any specified segment of the base line (giving the size of the birth cohort), and the frequency of points of any particular class of event which are located within any subdivision of the overall space.

Throughout this work, we use quinquennial coding of temporal location variables; the parallel lines in figure 1 are to be thought of as five years apart. This provides adequate cell frequencies to limit random error, in most situations, and is also a circumspect choice in the common situation in which the reporting of dates is approximate.

There are two different ways of using the time of occurrence of events. First, one may want to identify those individuals to whom an event occurred within an interval of time as members of some aggregate whose subsequent behaviour is to be examined; such aggregates are called cohorts (of one or another kind). For that purpose one needs a code of temporal location; it constitutes a particular kind of independent variable. Secondly, one may want to characterize the length of time elapsing between one event and another, in an individual's history. For that purpose, one may want to use a temporal measure coded to the feasible degree of precision, directly from the raw data. In this case the temporal measure constitutes a kind of dependent variable, and grouping would only be employed if the data were considered highly unreliable, ie as a form of rounding. However, when such intervals are employed as control variables, ie as part of the definition of the dependent variable, coded values are appropriate, and widely used with registration materials. In what follows, we suggest that there are substantial advantages with survey materials in coding such interval variables implicitly, as the intersection of the codes established for the times of occurrence of their beginning and ending events.

In the upper left diagram in figure 1, the boundaries of birth cohorts are represented by positive diagonals, running from birth (age zero) up to time of interview. The exact age limits of each episode of experience are represented by horizontal lines. The reproductive history of each cohort consists of the events in a series of cohort-by-age rhombuses, topped by a triangle of events. That triangle of experience cannot be used in conjunction with the data for adjacent cohorts without judgmental estimation, since it is not only incomplete, but off centre (unevenly distributed by age within the triangle).

In the upper right diagram in figure 1, the same experience is presented using another temporal format, the period of occurrence of births to respondents. The boundaries of periods are represented by negatively sloping diagonals, running from age zero up to the line corresponding to the upper age limit of the survey. Again because of the configuration of the data, the record for each period consists of a sequence of period-by-age rhombuses, topped by a triangle. As before, one either truncates the experience and wastes information, or indulges in questionable extrapolation beyond the available evidence.

Furthermore, if the cohort-by-age scheme is used for some purposes and the period-by-age scheme for other purposes, it is evident that there is redundancy in the alternative modes of organizing the available information. The same experience is summarized in two ways, with the component segments of the two orientations overlapping. This is inefficient.

One can avoid both the triangular embarrassment (which leads to ignoring or inventing data) and the redundant inefficiency by employing period-by-cohort coding, with age in the implicit role, as shown in the lower panel of figure 1. The same segments are available for either a period or a cohort orientation to the data set viewed as a time series. The coding practice exhausts the area of experience without residual triangles. The reason this outcome is fitting is that the universe defined for the survey has a period boundary (the time of interview) and a cohort boundary (the upper age limit at interview).

The same diagrammatic representation can be adapted to illustrate two comparable problems of summarizing the records collected in a survey. Consider a marriage record for birth cohorts. Using the cohort-by-age format (the upper left diagram), the evidence for the most recent birth cohort would be its marriages up to age 15/20; for the next most recent birth cohort would be its marriages up to age 20, and from 20 to 20/25; and so forth. Using the period-by-age format (the upper right diagram), the same outcomes would eventuate for the earliest period; for the next earliest period; and so forth. With a cohort-period diagram, the available segments of the experience are all shaped the same.

To generalize the point, consider the study of fertility with respect to time of marriage. In this case the cohorts would be marriage cohorts (coded by date of marriage) and the life cycle identifier indicated by horizontal lines in the upper two diagrams would be marital duration. Explicit coding of marital duration with exact boundaries (0, 5, 10, . . .) as is the conventional practice would lead to a triangular duration segment at the end of the history for each marriage cohort, or for each period of experience; again the problems are resolved by resort to the lower configuration.

Throughout the remainder of this work, the following coding practices are used:

Time of birth of respondent is coded $k = 1, 2, \ldots, 7$. This code is based on age at interview, so that $k = 1$ stands for respondents in ages 45/50 at interview, $k = 2$ for respondents in ages 40/45 at interview, and so forth. Respondents who are less than age 15 at interview are considered to be coded $k = 7$.

Time of marriage of respondent is coded $j = 4, 5, \ldots, 10$. This code is based on marital duration at interview, more precisely on the time elapsed between date of first marriage and date of interview. Thus $j = 4$ stands for those marriage cohorts in durations 30/35 at interview, $j = 5$ for those in durations 25/30 at interview, and so forth. Respondents who are married more than 35 years at interview are coded $j = 4$. The code for j is staggered by three units (15 years) relative to the code for k to provide the same temporal identification for period of respondent's marriage and period of respondent's birth. More generally, the code j is employed for time of entry into exposure to risk of one or another specified event, such as the time of occurrence of the x^{th} birth, which initiates exposure to risk of the occurrence of an $x + 1^{th}$ birth.

The same quinquennial periods are used to identify the time of occurrence of a birth to a respondent, with the label $i = 4, 5, \ldots, 10$, again based on the difference between that date and the date of interview. Births which occur to respondents more than 35 years prior to interview are coded $i = 4$. The focus of this work is on a survey of ever-married women, because that is the most common universe definition. While some of these women may report births in a (quinquennial) period prior to the period of first marriage, it is not improbable that the frequency of such cases will be small enough to justify the convenience of coding them as if they occurred in the same period as first marriage. Should the universe be defined as all women, no such approximation would be justified.

The labels (i, j, k) are used to identify the basic data which are employed in index construction. Thus in chapter 3 the basic data are N(k), the number of respondents in birth cohort k, and B(i, k), the number of births occurring in period i to members of birth cohort k. In chapter 4, the basic data are M(j, k), the number of marriages occurring in period j to members of birth cohort k, and B(i, j, k), the number of births occurring in period i to members of birth cohort k who married in period j. In chapter 5, the basic data are B(x +1, i, j, k), the number of births of order x + 1 occurring in period i to members of birth cohort k who had a birth of order x in period j.

The above descriptions serve to identify locations of the basic information on the data tape. Where measures are calculated from these data, for analytic use, it is more convenient to employ identifiers of life-cycle intervals as the labels. Thus age of respondent at time of occurrence of birth is defined as $a = i - k$. Age of respondent at time of entry into marriage or parity x is defined as e = j–k. Duration of marriage at time of occurrence of birth, or interval since entry into parity x, is defined as y = i–j.

Two points deserve emphasis at this juncture. In the first place, three codes suffice to identify both the three temporal locations, and the three temporal intervals, rather than the six codes ordinarily employed. In the second place, if one were to create explicit codes for age at occurrence of birth, age at occurrence of marriage, and marital duration at occurrence of birth, for example, the obligatory identity of the three for any individual would not always hold in coded form. With the proposed implicit coding, on the other hand, it is always the case that a = e + y, since (i–k) = (j–k) + (i–j).

Although there is an initial cost of unfamiliarity with a code which, to use the example of age, has groups with diagonal boundaries from, say, 20/25 to 25/30, rather than the much more common exact age limits, comparable grids play a prominent role in demography already. When a population projection of the customary component type is undertaken, the basis for the calculations, although often referred to loosely as age-specific birth and death rates, is actually a set of fertility and mortality values representing cohort behaviour period by period; each element has diagonal age boundaries. The mortality measures used to 'survive' an age group (a cohort) from one time point to the next are not ${}_5p_x$ values but ${}_5L_{x+5}/{}_5L_x$ values. Likewise fertility rates are used to estimate the births in each successive period as the sum of the reproductive outputs of each cohort in that period, and the births of the period which that sum constitutes form the new cohort.

The unfamiliarity of the coding recommendations is a small cost to pay for the substantial advantages achieved. This is one example of the way in which measurement practices should be adapted to the form of the data collection procedure at hand, rather than borrowed from practices developed to handle data produced in a different way.

3 Mode of Temporal Aggregation

3.1 PRELIMINARY DEFINITIONS

The first specific assignment is to make meaurements of what may be called the quantum of fertility, with the intention of identifying changes in such measurements over time. The quantum of fertility is measured by the number of births per woman occurring to those in a specified category over an extended interval of time, as a comprehensive measure of the reproductive experience over at least a substantial part of the relevant life cycle. It is distinguished from the tempo of fertility, measures of which concern the distribution of the quantum over time.

Time has two separate implications. The first sense, captured in the concept of tempo, is personal time, the location of a birth occurring to an individual by reference to the time elapsed since some reference event in the individual's life. That lapse of time is called age if the reference event is the date of the respondent's birth, (marital) duration if the reference event is the date of the respondent's (first) marriage, and (birth) interval if the reference event is the date of the preceding birth to the respondent. A comprehensive measure of the quantum of fertility is a cumulation of the reproductive experience across the several ages or durations. One of the challenges of a fertility survey is that respondents at interview are arrayed across personal time, with evident consequences for the length of experience for which comprehensive measures can be calculated.

The second sense of time is the one implied in the concept of a time series of measures. Once one has specified a particular index of the quantum of fertility, the task is to order such indices in what may be called historical time, and compare their respective values. To do so, there are two possible procedures, distinguished by mode of temporal aggregation. The cohort mode uses the calendar dates associated with the reference events in personal time to identify the membership of the respondents in successive cohorts – birth cohorts if the reference event is their birth; marriage cohorts if their reference event is marriage.

The other procedure uses the calendar dates associated with the occurrence of the fertility itself. Such measures are said to be achieved by the period mode of temporal aggregation. Although there are many ad hoc ways in which such indices may be constructed, those to be considered here are isomorphic with cohort indices. A principal task is to determine the extent to which indices of the same form, calculated in the different modes, may differ from one another, and therefore tell a different story about temporal variations in fertility.

Although there is much room for dispute concerning the appropriate mode of temporal aggregation to employ in a particular analytic situation, most theorizing about the causes of temporal variations in fertility employs concepts for which the cohort is the appropriate mode, whereas most measures of temporal variations are in fact in the period mode. In the literature, preponderant attention is devoted to different forms of period-specific output, especially for the time immediately prior to a survey, but the problems involved in their measurement, both the effects of censoring and the distinction between the two kinds of time series, are to a considerable extent ignored.

3.2 AGE-SPECIFIC MEASURES FOR BIRTH COHORTS

First we consider possible measures of comprehensive fertility which cover the available life span. The procedure followed is to define a measure for a cohort, and then develop the specifications for a period analogue to that measure. Arrangements of the data for a period in such a way that they resemble in form the history of a cohort are frequently called synthetic cohort constructions. Such measures are attractive because they convey the experience in a particular period in a form which resembles the values usually associated with reproductive behaviour.

Although the main interest is the development of approaches to fertility measurement for a survey of ever-married women, it is convenient to begin with the case of a sample of birth cohorts, unrestricted by marital status. With age limits of 15 and 50 as sample criteria, the basic data for such analysis consist of the number of women in birth cohort k(= 1, 7), N(k), and the number of births in period i(= k + 3, 10) to women of birth cohort k, B(i, k). The quinquennial coding practice is followed here as elsewhere, and age is defined implicitly as the difference between period and cohort, a = i–k.

In table 1 the basic data to be used in this and the following chapter are displayed as they might appear on a data tape. For present purposes, attention is restricted to the summary rows for each horizontal panel, aligned with the cohort identification k, and consisting of cohort size, N(k), and the numbers of births occurring in period i to cohort k, B(i, k). Thus cohort k = 1 consists of 1000 respondents, with 93 births in period 4, 467 births in period 5, and so forth. (These are illustrative data, and do not represent actual survey results.)

Fertility cumulated to interview for birth cohort k is readily obtained by summing its births over the span of periods, and dividing by the number of women in the cohort:

$$FC(k) = \sum_{i=k+3}^{10} B(i, k)/N(k)$$

If there were no further births subsequent to interview, this would be the total fertility rate. The value for cohort k = 1 is FC(1) = 3455/1000 = 3.455.

In order to develop the synthetic cohort analogue of this measure, it is necessary to distinguish the component behaviour of each cohort in each period, ie in each age, since age is defined implicitly. The fertility rate for birth cohort k in age a is

$$f(a, k) = B(k + a, k)/N(k) = B(i, k)/N(k)$$

These rates are shown in table 2. The display follows the triangular format of the lower diagram in figure 1 (as do all the tables of rates presented here).

$$FC(k) = \sum_{a=3}^{10-k} f(a, k)$$

It is evident that the entries in table 2 can also be construed as the fertility rates in age a in period i.

$$f(a, i) = B(i, i-a)/N(i-a) = B(i, k)/N(k)$$

Thus the cumulated fertility in period i is

$$FP(i) = \sum_{a=3}^{i-1} f(a, i)$$

The value for period i = 10 is FP(1) = 2.655, the sum of the entries in the right-hand negatively sloping diagonal of table 2.

Table 1 Basic data table

k	j	N(k)	M(j, k)	B(i, j, k) i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
1		1000		93	467	880	792	585	449	189
	4		200	93	161	171	155	113	59	28
	5		608		306	648	517	385	314	130
	6		96			61	104	60	50	14
	7		29				16	22	15	10
	8		13					5	9	4
	9		6						2	3
	10		2							0
2		1100			86	449	893	822	604	447
	5		198		86	150	149	160	96	33
	6		658			299	677	527	405	326
	7		122				67	117	75	62
	8		39					18	23	18
	9		18						5	7
	10		6							1
3		1200				79	424	898	844	627
	6		192			79	126	154	138	82
	7		706				298	667	565	403
	8		151					77	121	107
	9		51						20	29
	10		22							6
4		1300					69	391	893	868
	7		182				69	102	146	114
	8		749					289	653	588
	9		185						94	140
	10		66							26
5		1400						60	358	883
	8		168					60	92	127
	9		788						266	655
	10		222							101
6		1500							51	318
	9		150						51	76
	10		824							242
7		1600								41
	10		128							41
			M(j)			B(i, j)				
	4		200	93	161	171	155	113	59	28
	5		806		392	798	666	545	410	163
	6		946			439	907	741	593	422
	7		1039				450	908	801	589
	8		1120					449	898	844
	9		1198						438	910
	10		1270							417

Table 2 Age-specific birth rates for cohorts – 1000 · f(a, k)

i = 10							a(= i − k)
189							9
449	406						8
585	549	523					7
792	747	703	668				6
880	812	748	687	631			5
467	408	353	301	256	212		4
93	78	66	53	43	34	26	3
k = 1							

The two cumulated fertility indices, FC(k) and FP(i), where the letters C and P identify them as cohort and period respectively, are isomorphic in the sense that the former is achieved by cycling the rates over the available periods for each cohort and the latter by cycling the rates over the available cohorts for each period.

Some notes may be useful.

1 The calculation of age-specific birth rates is necessary for the period but not for the cohort calculation, essentially because the denominator, N(k), is fixed for a cohort but not for a period. To remove the influence of cohort size from the period calculation is referred to as removal of the influence of the age distribution.
2 Since cohort size is fixed, there is no obligation, as there would be with registration materials, to calculate person-years of exposure to risk for particular cohorts within particular time periods.
3 Each successive real cohort is censored by interview at one younger age than its immediate predecessor. Comparison of two successive real cohorts requires, for comparability of age span, that the final period of experience for the earlier real cohort be deleted from the calculation.
4 Each successive period is censored by the age limit at interview at one older age than its immediate predecessor. Comparison of two successive synthetic cohorts requires, for comparability of age span, that the experience of the first real cohort be deleted from that summation for the later period.

The principle underlying such adjustments in the interest of comparability is obvious once stated. What makes the point noteworthy at all is that such precautions are often ignored. Frequent resort to the triangular diagram is a worthwhile aid, and may help avoid such blunders.

3.3 COMPARISON OF REAL AND SYNTHETIC COHORT RESULTS

For anything approaching comprehensive coverage of the life cycle, it is evident that neither real nor synthetic cohorts, considered separately, offer much in the way of temporal scope. Accordingly, there is considerable interest in the validity of a comparison of the earliest experience represented in the survey, that of the first cohort, and the latest experience, that in the last period. Perhaps the commonest way in which the question is put, as implied in the 'synthetic cohort' appellation itself, is the extent to which the period parameters are valid proxies for cohort experience, as a kind of average of the behaviour of those real cohorts contributing to reproductive output in the period in question. That is the orientation adopted in what follows, and the justification for characterizing departures of period parameters from their cohort counterparts as distortions.

It deserves emphasis, however, that the cogency of the argument to be presented does not depend on acceptance of that orientation. A defensible case can be made for considering a time series of period measures as the analytic desideratum, and therefore inquiring into the extent to which cohort parameters represent a kind of average of the experience occurring in the periods through which the cohort passes during its reproductive years. From either viewpoint, the sources of distortion are formally the same: it is irrelevant to the statistical argument underlying them whether one views the cohort record as a period array or the period record as a cohort array. One can accordingly view the orientation adopted here either as a manifestation of analytic preference or simply as an arbitrary choice between twin presentations to avoid repetition and reduce confusion.

From the cohort perspective, one can think of the configuration of fertility rates for a cohort, period by period and thus age by age, as a series of proportions of the cohort's eventual fertility total, characterized *in toto* as the temporal distribution of that fertility. This permits a statistical distinction between the quantum of fertility (the total) and the tempo (its distribution). If both quantum and tempo remain fixed from cohort to cohort, the record for each period will display that same quantum and tempo (when calculated in analogous fashion from the fertility rates in the period). If the tempo remains fixed from cohort to cohort, but the quantum varies, the period quantum will be a weighted average of quantum for the cohorts represented in that particular period, where the weights are the (fixed) temporal distribution, and add up to unity. The period tempo, however, will tend to be shifted upwards (in age) by a downward trend in cohort quantum, and downwards (in age) by an upward trend in cohort quantum, relative to the (fixed) cohort tempo. Variation in cohort quantum creates distortion in period tempo. As a first approximation, the difference between the means of the cohort and period age distributions of fertility is equal to the product of the relative change in cohort quantum, and the variance of the cohort age distribution of fertility.

If the quantum remains fixed from cohort to cohort but the tempo varies, the period quantum will ordinarily differ from that value because, in such circumstances, it will be a weighted average in which the sum of the weights departs from unity. Thus a tendency for earlier cohorts in a period to allocate larger proportions of their total fertility to older ages, and for later cohorts in the same period to allocate smaller proportions of their total fertility to older ages (and thus larger proportions to younger ages), will yield a sum of weights greater than unity and thus an upward distortion of period quantum. Conversely, if successive cohorts shift toward an older age distribution of fertility (perhaps because of a rising age at marriage), there will be a downward distortion of period quantum, *ceteris paribus.* Variation in cohort tempo creates distortion of period quantum. As a first approximation the ratio of the quantum of period to that of cohort fertility is equal to the complement of the time derivative of the mean of the cohort age distribution of fertility. Although the quality of this approximation (and of that for distortion of tempo) depends on the extent to which concomitant changes are under way in both the quantum and the tempo of cohort fertility – as is ordinarily the case empirically – the formulae capture the sense of the major sources of distortion.

Most demographic analyses focus on change in the quantum of fertility. Given that emphasis, it is clear from the foregoing account that co-ordinate attention to change in the tempo of fertility is a sensible precaution when cohort and period parameters are being compared. Moreover, the tempo of fertility is an interesting and important dimension of reproductive behaviour in its own right. The growth rate of a population is the key parameter in considerations of population policy. In a closed population with fixed reproductive processes and a stable age distribution, the rate of natural increase is the ratio of the natural logarithm of the net reproduction rate (a quantum phenomenon) to the generation length (a tempo phenomenon). Admittedly the relative change in the numerator exceeds that in the denominator during the course of a typical demographic transition. On the other hand, among high fertility populations an important source of differentiation is their distinctiveness with respect to the time pattern of childbearing; among low fertility populations, the principal source of temporal variation in the quantum of fertility is modification of the tempo of fertility, as discussed in the following section.

3.4 PERIOD SOURCES OF FERTILITY VARIATION

We have been considering the consequences, for comparison of period and cohort indices, of long-term change in the quantum or tempo of cohort fertility. As noted above, analogous propositions could have been derived from an orientation to the consequences of long-term change in the quantum or tempo of period fertility. In either case, the basic statistical element is a record of the performance of a particular cohort in a particular period. The conceptual distinction between a cohort and a period orientation to that record can be considered as a judgement of the relative importance of the past experience of the cohort, on the one hand, and the distinctive environmental stimuli characterizing the period, on the other, in determining that performance.

While the choice between a cohort and a period format for the assessment of long-term change is an arguable judgement, it is empirically evident that short-term changes in the environment produce responses in the same direction and to much the same degree by the constituent cohorts in such periods. For any one cohort, such experiences tend to be counter-balancing over time: the cohort displaces its fertility away from the bad years and towards the good years. In terms of the parameters of the preceding account, there is response in cohort tempo but not necessarily in cohort quantum. From a period standpoint, however, it is the quantum of fertility that is sensitive to the change. One may expect, therefore, short-term distortion of the quantum of period fertility relative to the quantum of cohort fertility. One palliative for this well-known phenomenon is the practice of calculating an average period total fertility rate for a span of years; the recommendation of a five-year temporal unit in coding survey data is in this spirit.

Beyond such short-term interperiod variations, there is another empirical possibility for a kind of change originating in the environment and most clearly manifest in period by period comparison, viz, a discontinuous change, unreversed subsequently, associated, for example, with major modifications in technology or in legislation, and impinging on the subsequent experience of all cohorts from the time of its inception. From the standpoint of the period time series, there will be an abrupt shift in the quantum of fertility; from the standpoint of the cohort time series, there will be a gradual modification in the quantum of fertility, since earlier cohorts will manifest the response only in their oldest ages, whereas later cohorts will spend almost their entire reproductive span in the new environment.

3.5 ANALYSIS OF INCOMPLETE HISTORIES

To summarize the argument to this point, we have noted that the synthetic cohort construction is often an inept proxy for parameters of real cohort fertility, either because there may be a transitory rise or fall in period fertility, synchronized because the stimulus is common to all cohorts, or because there may be a drift in the tempo of cohort fertility, distorting period quantum in the interim, or a drift in the quantum of cohort fertility, distorting period tempo in the interim.

In the light of this, one alternative would be to compare successive real cohorts, and pay no attention to the synthetic cohort alternative. The problem with this approach is that each successive cohort has one fewer age segment than its predecessor. Although one can obtain life-cycle comparability in a mechanical sense, by deleting from the earlier cohort the age not found for the later cohort, the outcome is less than satisfying. The sacrifice of evidence is regrettable when the total available supply is expensive and limited. But beyond that, the comparability may be merely superficial. If one could be assured that successive cohorts had the same age distribution of fertility, then, by definition, the cumulated fertility to any age, relative to the total, would be the same for each. But in that circumstance, one could proceed to compare real and synthetic cohort results without concern. It is highly likely that temporal variation in the quantum of fertility is accompanied by temporal variation in its tempo. A comparison of incomplete real cohort histories is subject to the same failings as a comparison of a real history with a synthetic history.

An alternative to the loss of information involved in truncating the experience of the earlier cohort to make it comparable with that for the later cohort is to use the available information on the fertility-age function to make a best estimate of what that function will look like once it is completed. Although considerable ingenuity has been expended on this assignment, there are bases for viewing the outcome with a jaundiced eye. Such efforts rely on the examination of complete histories for other times and places. The assumption underlying the extrapolation is, in short, that we know the pattern of fertility, that it is part of our current stock of knowledge. Yet the purpose of investigating fertility in another particular time and place is to find out something new. There is ample empirical evidence that the future in some societies has turned out to be unlike the past in other societies. One can make a best guess, *faute de mieux*, but the assumption underlying the guess contradicts the purpose of the survey.

This comment is not intended to denigrate the employment of assumptions concerning missing data, and like activities, but rather to suggest that the summarization of the results of a fertility survey ought to be kept distinct from the activity of model construction, in which judgement and other kinds of evidence are used. The latter activity is intrinsically idiosyncratic to the person doing the work, and competitive models have a considerable virtue. To convert the process into a routine procedure for elaborating survey results would be unfortunate for the survey as a body of fresh evidence, and unfortunate for progress in model construction as well.

All that has been said about the difficulties of temporal comparison of real cohort results applies, *mutatis mutandis*, to the comparison of synthetic cohort results. The question of what might be gained from a comparison of data for successive periods deserves careful attention because of the greater interest in recent change. It is true that, although the quantum of period fertility is distorted by changes in the tempo of cohort fertility, those changes may be of comparable magnitude in the periods being compared, permitting an inference about relative change in the quantum of cohort fertility from relative change in the quantum of period fertility. On the other hand, the quantum of period fertility is peculiarly sensitive to short-term variations, and successive periods tend to show divergences in opposite directions. The problem is not intractable, but its solution requires more detail about the reproductive process than we have so far considered. That is the intention in the next two chapters.

4 Measures for a Sample of Ever-Married Women

4.1 FERTILITY SPECIFIC FOR AGE AND MARITAL STATUS

In the typical WFS sample, the universe is defined as ever-married women under the age of 50. Stipulation of a marital status criterion introduces some complexity into the design of synthetic cohort measures analogous to those for real cohorts. Moreover, some care is required to avoid censoring bias. The purpose of this chapter is to present procedures which resolve these problems.

For a real birth cohort, an appropriate measure of the fertility of ever-married women is straightforward. For experience up to interview, one sums the number of births occurring to cohort k,

$$BB(k) = \sum_{i=k+3}^{10} B(i, k)$$

and the number of marriages likewise

$$MM(k) = \sum_{j=k+3}^{10} M(j, k).$$

The required index is the ratio of the former to the latter.

$$GC(k) = BB(k)/MM(k)$$

This is the mean parity at interview for ever-married women of birth cohort k. It is comparable to the result of the question on a census asking the number of children ever born to women ever married. For cohort $k = 1$, the calculation gives $3455/954 = 3.622$. (The reason why the number of births to the cohort, within marriage, is the same as all births is that we have adopted the simplifying assumption that never-married women are infertile.)

There is no formal difficulty in developing a synthetic cohort analogue to this measure, although a superior proposal will be advanced shortly. The essence of the procedure is to identify the contribution to the real cohort index which is made in each separate period. Then one re-assembles the contributions by period rather than by cohort, across the available ages (cohorts).

The introduction of marital status considerations into the measurement system can be thought of as permitting a partitioning of the age-specific fertility rate into a nuptiality component and a marital fertility component:

$$f(a, k) = mc(a, k) \cdot g(a, k)$$

In this formula, $g(a, k)$ is the ratio of births to the number of ever-married women, in age a, for cohort k; $mc(a, k)$ is the proportion of women in the age who are ever married. Calculation of these measures requires a digression on nuptiality. If eligibility for individual interview requires that the woman be ever-married, information about the size of the birth cohort, $N(k)$, must be inferred from data collected in the household questionnaire (or some comparable source of information such as a recent census). If one assumes that the proportion of the cohort ever married, in the individual interviews, say $EC(k)/N(k)$, is the same as the proportion of the cohort ever married, in the household schedule, say $EC'(k)/N'(k)$, it follows that $N(k) = EC(k) \cdot N'(k)/EC'(k)$. The estimate will be flawed if age or marital status are differently reported in the household schedule and in the individual questionnaire, and if non-response to the individual questionnaire is related to marital status or age. Accordingly it is fortunate that many of the important calculations to be proposed in subsequent chapters depend not at all on the value of $N(k)$.

With registration data, the customary way to produce a nuptiality function is to calculate central first marriage rates by age, for the never-married population (those exposed to risk), and transform the rates into parameters of 'survival' in the single state, by analogy with the conventional life table. With survey data, a more direct procedure is available. For women never married at the beginning of age e, one can calculate the proportion who remain never married at the end of age e, or what may be termed a celibate survival ratio.

$$p(e, k) = \frac{N(k) - \sum_{j=4}^{k+e} M(j, k)}{N(k) - \sum_{j=4}^{k+e-1} M(j, k)}$$

The values $M(j, k)$, representing marriages in period j to members of birth cohort k, are shown in the left-hand column of table 1. $p(e, k)$ is shown in table 3. The celibate survival ratios, $p(e, k)$, serve to define the nuptiality function.

Then the proportion of the cohort ever married by interview is

$$EC(k) = 1 - \prod_{e=3}^{10-k} p(e, k)$$

The proportion single at the beginning of age a is

$$s(a, k) = \prod_{e=3}^{a-1} p(e, k)$$

where $s(3, k) = 1$.

The problem of determining the person-years single (and thus the person-years ever married) during age a, relative to the proportion single at the beginning of the age, is analogous to that involved in calculating ${}_1L_x/1_x$ values in a life table. On the assumption that the probability of marriage during age a is constant at the level implicit in $p(e, k)$, the required value is

$$L(a, k) = (1 - p(a, k))/(\text{cln}\, p(a, k))$$

Table 3 Celibate survival ratios for cohorts – 10 000 p(e, k)

j = 10	e(= j − k)
9583	9
8889 9077	8
8060 7831 7800	7
6979 6803 6623 6413	6
5000 5000 5000 4986 5000	5
2400 2646 2996 3301 3604 3896	4
8000 8200 8400 8600 8800 9000 9200	3

Table 4 Marital age-specific birth rates for cohorts – $1000 \cdot g(a, k)$

i = 10							a(= i − k)
198							9
473	430						8
622	588	564					7
862	821	784	753				6
1021	962	914	864	818			5
814	747	691	627	570	550		4
894	843	798	742	698	671	631	3
k = 1							

where cln stands for the natural cologarithm. The person-years ever married, in age a, is therefore

$$mmc(a, k) = 1 - (s(a, k) \cdot L(a, k))$$

Accordingly, the total fertility rate (or rather the cumulated fertility up to interview) can be expressed as follows:

$$FC(k) = \sum_{a=3}^{10-k} (mmc(a, k) \cdot g(a, k))$$

and the mean marital parity at interview is

$$GC(k) = FC(k)/EC(k)$$

From a real cohort standpoint, the purpose of the partitioning is to display the separate patterning of the nuptiality and marital fertility components of the age-specific fertility rates:

$$f(a, k) = mmc(a, k) \cdot g(a, k)$$

The values g(a, k) are shown in table 4.

The partitioning also provides a way to construct analogous measures for the synthetic cohort, as follows. The celibate survival ratios are used to calculate nuptiality weights for periods, as well as for cohorts.

$$s(a, i) = \prod_{e=3}^{a-1} p(e, i-e)$$

where $s(3, i) = 1$

$$L(a, i) = (1 - p(a, i-a))/(\text{cln}\, p(a, i-a))$$

$$mmp(a, i) = 1 - (s(a, i) \cdot L(a, i))$$

$$EP(i) = 1 - \prod_{e=3}^{i-1} p(e, i-e)$$

$$FP'(i) = \sum_{a=3}^{i-1} (mmp(a, i) \cdot G(a, i-a))$$

$$GP'(i) = FP'(i)/EP(i)$$

The prime is attached to the value of $FP'(i)$ to distinguish it from the outcome at the age-specific level, $FP(i)$. The prime is attached to the value of $GP'(i)$ to indicate its affiliation with $FP'(i)$. Such identifications are unnecessary with cohorts since the index is the same whatever the level of specificity employed.

The respective values for cohort k = 1 and period i = 10 are as follows:

mmc(a, 1)	g(a, 1)	a	mmp(a, 10)	g(a, 10 − a)
0.104	0.894	3	0.041	0.631
0.574	0.814	4	0.404	0.550
0.862	1.021	5	0.742	0.818
0.919	0.862	6	0.855	0.753
0.940	0.622	7	0.898	0.564
0.949	0.473	8	0.915	0.430
0.953	0.198	9	0.920	0.198

Finally we have

$$FC(k) = 3.455 = (GC(k) \cdot EC(k)) = 3.622 \cdot 0.954$$

and

$$FP'(i) = 2.581 = (GP'(i) \cdot EP(i)) = 2.799 \cdot 0.922$$

The synthetic cohort total fertility rate based on age-specific fertility rates, $FP(i) = 2.655$. The general point is that there is no unique synthetic cohort total fertility rate, but rather an array of such measures, dependent on the specificity of rates used in their construction.

The idea of the synthetic cohort is that the experience of successive real cohorts, in their respective life-cycle stages within a particular period, is treated as if it were the consecutive experience of a cohort. In the first illustration, in the preceding chapter, the only information specified for the population exposed to risk was birth cohort membership (the minimum specification requisite to construction of a synthetic cohort). In the present case, this has been supplemented by a further specificity, the distinction between those ever married and never married in the age in question. For the real cohort, this step increases the amount of information displayed, without changing the value of the index. For the synthetic cohort, the value of the index is particular to the level of specificity.

Nuptiality for the synthetic cohort, in the above procedure, has been based on the cohort-specific celibate survival proportions for the period in question, p(e, k). The case for distinguishing between the ever-married and the never-married, when one devises an appropriate exposure denominator for the former's births, is that the never-married are not exposed to risk. By the same token, an age-specific marriage proportion which had a denominator consisting of all women would be inferior to a calculation recognizing that never-married women are the only ones exposed to risk of first marriage.

Why is the more highly specific index preferred? The assumption underlying the synthetic cohort is that the experience of different real cohorts in a particular period may be treated as if it were consecutive. But the proportion married at the end of one age segment (for one cohort) ordinarily differs from the proportion married at the beginning of the next age segment (for the next earlier cohort). The sequence is incoherent unless the observed proportions married for the various cohorts are replaced by those derivative from one nuptiality function. The function based on marriage rates observed in the period is the evident choice.

The preference can be argued in another way. The synthetic cohort construction is an attempt to characterize experience within a period, uncontaminated by the experience of prior periods. The proportions married, for each cohort, represent the outcomes of their previous marital histories. Accordingly, on this criterion, they should be replaced.

The preceding account has taken it for granted that the most apt characterization of the universe is a sample of (ever-married) women from a sequence of birth cohorts. But it could equally be characterized as a sample of women (under age 50) from a sequence of marriage cohorts, ie respondents married within particular time periods. Since there is nothing in principle to prevent the use of the marriage cohort as the unit of temporal analysis, we propose to investigate the possibilities.

Suppose the respondents are sorted into marriage cohorts on the basis of their period of marriage, $j (= 4, 10)$, so that the size of the marriage cohort is

$$M(j) = \sum_{k=1}^{j-3} M(j, k)$$

and births are tabulated by period of occurrence and period of mother's marriage, $B(i, j)$. These values are shown in the lower panel of table 1.

In the same way that one defines the age-specific rate for a birth cohort

$$f(a, k) = B(k + a, k)/N(k) = B(i, k)/N(k)$$

one can define the duration-specific fertility rate for a marriage cohort

$$h(y, j) = B(j + y, j)/M(j) = B(i, j)/M(j)$$

where y, the years of marital duration, is defined implicitly as the difference between the period of occurrence of the birth (i) and the period of marriage (j). This is analogous to the implicit definition of age as $a = i - k$. The code $y = 0$ stands for fertility up to duration 0/5, $y = 1$ stands for fertility between duration 0/5 and 5/10, and so forth. The values of $h(y, j)$ are shown in table 5.

Again by analogy with the treatment of age-specific fertility, one can define a cumulated fertility index, up to interview, for the real marriage cohort

$$HC(j) = \sum_{y=0}^{10-j} h(y, j)$$

and the parallel construction for the synthetic marriage cohort

$$HP(i) = \sum_{y=0}^{i-4} h(y, i - y)$$

Table 5 Duration-specific birth rates for marriage cohorts – $1000 \cdot h(y, j)$

i = 10							y(= i − j)
140							6
295	202						5
565	509	446					4
775	676	627	566				3
855	826	783	771	754			2
805	990	959	874	802	760		1
465	486	464	433	401	366	328	0
j = 4							

Thus from table 5, the value for HC(4) = 3.900, the sum of the entries in the left positive diagonal, and the value for HP(10) = 3.196, the sum of the entries in the right negative diagonal.

This seems to be an attractive alternative to age-specific fertility rates, particularly since one is working with an ever-married sample, as a simple and direct measure of marital fertility. It is not surprising that the index has attracted some attention, and it is important to identify several serious flaws in the procedure.

One problem is that the definition of the universe includes the stipulation that all women interviewed must be not only ever married but also under age 50. The effect of this stipulation is that the successive (real and synthetic) marriage cohorts are differentially censored by age at marriage. All those married in period j = 4 are obliged to have been married in the youngest age at marriage, e = 3, since otherwise they would have been too old by time of interview to have been included in the sample. Likewise those married in j = 5 could only have been married in ages at marriage e = 3 and e = 4. The marriages in the final period, j = 10, on the other hand, may encompass the entire range of ages at entry into marriage, e = 3, . . . 9.

From the standpoint of the synthetic marriage cohorts (the negative diagonals in the h(y, j) table), consider the sources of the duration-specific experience. For duration y = 6, the uppermost duration, the marriage cohort responsible, j = 4, is restricted to age at marriage e = 3; for duration y = 5, the marriage cohort responsible, j = 5, is restricted to ages at marriage e = 3 and e = 4; for duration y = 0, on the other hand, the full range of ages at marriage is feasible for the marriage cohort.

In brief, it is improper to compare the experience of successive real marriage cohorts without some control on the non-comparability of their age at marriage limits, and it is improper to aggregate the experience for a synthetic marriage cohort without the same kind of consideration with respect to the successive marital durations.

Parenthetically, one can draw a broad general inference from this observation, extensively applicable in survey research. If the universe is defined as ever-married women, with some upper age limit, any cross-classification of any variable by age or by marital duration alone will implicitly incorporate an age at marriage bias; if age at marriage has any relationship to the variable in question, corrective measures should be taken.

In order to resolve this problem, it is necessary to have information about births specific for age at marriage as well as for marital duration. In terms of the period coding employed throughout this work, it is necessary to specify k as well as i and j. That information is provided in the main body of table 1.

Before discussing the procedure, a further reason for specifying the birth cohort (k) should be mentioned. The numbers of marriages occurring in a period (and constituting a real marriage cohort) have an age distribution which depends in part on the relative sizes of the birth cohorts contributing marriages in that period, or what is ordinarily called the age distribution. It is considered standard practice in demographic measurement to remove from calculations the influence of the age distribution, because it has an arithmetical effect on the outcome which is unrelated to the phenomenon under investigation. If, for example, one population is growing more rapidly than another, it will tend to have marriage cohorts with a younger distribution of marriages by age than the other, because rapid growth is associated with increasing size of birth cohort. One may also find in particular populations considerable irregularity in the sequence of birth cohort sizes.

Unbiased measures of marriage cohort fertility require marriages specific for age at marriage as well as period of marriage, ie M(j, k), and births specific for age at marriage

Table 6 Birth ratios specific for age at marriage and marital duration, for birth cohorts[a]

e(= j − k)	1000·g(y, e, k)	y(= i − j)
9	000	0
	k = 1 i = 10	
8	500	1
	333 167	0
	k = 1 i = 10	
7	308	2
	692 389	1
	385 278 273	0
	k = 1 i = 10	
6	345	3
	517 486	2
	759 622 569	1
	552 486 392 394	0
	k = 1 i = 10	
	j = 7	
5	146	4
	521 508	3
	625 615 709	2
	1083 959 801 757	1
	635 549 510 508 455	0
	k = 1 i = 10	
	j = 6	
4	214	5
	516 495	4
	633 616 571	3
	850 801 800 785	2
	1066 1029 945 872 831	1
	503 454 422 386 338 294	0
	k = 1 i = 10	
	j = 5	
3	140	6
	295 167	5
	565 485 427	4
	775 808 719 626	3
	855 753 802 802 756	2
	805 758 656 560 548 507	1
	465 434 411 379 357 340 320	0
	k = 1 i = 10	
	j = 4	

[a] The symbols used in this table are identified in the text.

and marital duration as well as period of occurrence, ie B(i, j, k). The basic fertility ratio is

$$g(y, e, k) = \frac{B(k+e+y, k+e, k)}{M(k+e, k)} = \frac{B(i, j, k)}{M(j, k)}$$

These measures are shown in table 6. From a marriage cohort viewpoint (keying the calculations to period of marriage, j), the values g(y, e, k) would be construed as g(y, e, j − e).

The essential problem of creating comparable fertility indices for marriage cohorts is evident in the structure of table 6. With each higher age at marriage (e), there is one less period of marriage (j) available for inclusion, and one less marital duration (y). A compromise must be made between comprehensive coverage by age at marriage and comprehensive coverage of marital durations. One choice, for example, would be to restrict the span of marriage ages to e = 3, . . . 5 (or, in conventional terms, marriages at ages less than 25/30) and thus restrict the durations to be included to y = 0, . . . 4 (or, in conventional terms, durations less than 20–25). Note that the sum of the age at marriage limit and the duration limit is 50, because that is the upper age limit in the survey. These restrictions also limit the temporal span covered, since the marriage cohorts j = 4 and j = 5 are not represented at all three marriage ages. In terms of table 6, calculations are restricted to the lower three panels of the table, excluding the left most diagonal in the second to last panel and the two leftmost diagonals in the last panel.

For that tailored set of data, one can calculate the cumulative fertility ratios for real marriage cohorts (the left positive diagonal) and for synthetic marriage cohorts (the right negative diagonal) as in previous such exercises, except that in this case one has separate values for each age at marriage, as follows:

	e = 3	e = 4	e = 5
Real marriage cohort j = 6	3.015	3.395	3.010
Synthetic marriage cohort i = 10	2.636	2.976	2.575

These are unbiased results, although the achievement of that outcome required the sacrifice of a lot of evidence. Moreover, the answers are specific for age at marriage.

Consideration of what would be an appropriate system of weights to apply to these values, in order to have an index of marital fertility, leads to a last intractable problem with the marriage cohort orientation. For the synthetic marriage cohort, one can use the celibate survival ratios, p(e, j − e), as before, to create a distribution of marriages by age appropriate to period j. But for a real marriage cohort index, there is no sensible distributional choice. Of course a real marriage cohort has an age at marriage distribution (and it can be made independent of the influence of changing birth cohort size without difficulty), but it is constituted of a cross-section of the proportions marrying in the period for a series of birth cohorts. The age at marriage distribution for a real marriage cohort is a synthetic cross-section of the nuptiality of real birth cohorts. Accordingly, it reflects whatever change in nuptiality is occurring from one birth cohort to the next. The interesting twist to this outcome is that the conventional problem of devising comparable measures for cohorts and periods is how to make a period measure which is isomorphic with the cohort measure, whereas, in the present situation, the problem is how to make the cohort measure isomorphic with the period measure.

The source of the problem is the treatment of the temporal sequence of the reproductive process. In life-cycle terms, the respondent's birth (k) comes first, followed by marriage (j) and finally fertility (i), with the generally minor exception of premarital births. Following the registration style of measurement (the synthetic cohort approach), one can move back from the occurrence of birth (i) to the marriage cohort (j) and thence to the birth cohort (k), or, in more conventional terms, specifying marital duration (y = i − j) and then age at marriage (e = j − k). But if one orients one's calculations to the real marriage cohort (j), one is making calculations forwards with respect to births (i) but backwards with respect to time of respondent's birth (k). The subject recurs in the next chapter.

To conclude this section with a summary, the definition of the universe imposes special obligations if one wants unbiased measures of marital fertility for real and synthetic marriage cohorts. A three-dimensional classification of births and a two-dimensional classification of marriages is required. Some evidence must be sacrificed, by deletions, in the interest of time series comparability, and the data must be standardized for birth cohort size. Even so, the effort ends in failure, since there is no index which can be justified for both the real and synthetic marriage cohort. The conclusion is that the marriage cohort orientation, as a basis for marital fertility calculations of the kind discussed here, should be abandoned.

4.3 FERTILITY BY AGE AT MARRIAGE FOR BIRTH COHORTS

The difficulties arising from a marriage cohort orientation are all obviated simply by considering the data in table 6, g(y, e, k), from a birth cohort standpoint. In essence, the history of the birth cohort is conceptualized in two stages: (1) the occurrence of marriages in successive periods (ages); (2) for each marriage subcohort, ie for each period-specific subset of marriages, the occurrence of births in successive periods (marital durations).

The cumulated fertility ratio to interview is

$$FC(k) = \frac{\sum_{i=k+3}^{10} B(i, k)}{N(k)} = \frac{\sum_{j=k+3}^{10} \sum_{i=j}^{10} B(i, j, k)}{N(k)}$$

Define nc(e, k) as the proportion of the cohort marrying in age e.

$$nc(e, k) = \frac{M(k+e, k)}{N(k)} = \frac{M(j, k)}{N(k)} = \prod_{j=3}^{e-1} p(j, k) - \prod_{j=3}^{e} p(j, k)$$

$$\text{Then } FC(k) = \sum_{j=k+3}^{10} nc(j-k, k) \sum_{i=j}^{10} g(i-j, j-k, k)$$

$$= \sum_{e=3}^{10-k} nc(e, k) \sum_{y=0}^{10-k-e} g(y, e, k)$$

$$= \sum_{e=3}^{10-k} nc(e, k) \cdot GC(e, k)$$

where GC(e, k) is the cumulated marital fertility for age at marriage e. Thus the cumulated fertility ratio, FC(k), is a weighted sum of the age at marriage specific cumulated marital fertility ratios, where the weights are the proportions of the cohort marrying in each age. Finally GC(k) = FC(k)/EC(k) as before (where EC(k) is the proportion of the cohort married by interview).

The parallel construction for synthetic birth cohort i is as follows:

$$FP''(i) = \sum_{e=3}^{i-1} np(e, i) \cdot GP''(e, i)$$

$$\text{where } np(e, i) = \prod_{j=3}^{e-1} p(j, i-j) - \prod_{j=3}^{e} p(j, i-j)$$

$$\text{and } GP''(e, i) = \sum_{y=0}^{i-e} g(y, e, i-y-e)$$

Finally

$$GP''(i) = FP''(i)/EP(i)$$

where EP(i), as defined previously, is the proportion of the synthetic cohort married by the uppermost available age, as defined by p(e, k).

The values for real cohort k = 1 and for synthetic cohort i = 10 are shown in the following table.

nc(e, 1)	CG(e, 1)	e	np(e, 10)	GP″(e, 10)
0.2000	3.900	3	0.0800	2.943
0.6080	3.782	4	0.5616	3.190
0.0960	3.010	5	0.1792	2.575
0.0290	2.173	6	0.0643	1.794
0.0130	1.385	7	0.0253	0.970
0.0060	0.833	8	0.0082	0.667
0.0020	0.000	9	0.0034	0.000
EC(1) = 0.9540			EP(10) = 0.9220	
	FC(1) = 3.454			FP″(10) = 2.634
GC(1) = 3.621			GP″(10) = 2.856	

For the real cohort, as before, the indices are the same (within rounding error), whatever the level of specificity. The synthetic cohort values are, for cumulated overall fertility, 2.634 for FP″(10) above, compared with 2.581 for FP′(10), where age and marital status were specified, and with 2.655 for FP(10), where only age was specified.

The differences among the three synthetic cohort values are so small as to raise the question of whether the extra work was worth the trouble. One answer is that, in dealing with age-specific fertility measures, each age a(= i − k) is a compound of values for an unspecified mixture of age at marriage e(= j − k) and years of marital duration y(= i − j). Without proceeding with the calculations at the higher level of specificity, there is no way of knowing the magnitude of the difference; the pronounced variations of fertility by both age at marriage and marital duration make the question interesting.

But the case for increased specificity is primarily based on the analytic value of the comprehensive array of information contained in the above table, and the way it permits one to make a definitive distinction between the respective roles of nuptiality and marital fertility in determining overall fertility. No other approach to this important question reveals the structure of interactions, knowledge of which is essential to a definitive statement.

The approach can readily be extended to a consideration of the separate roles of nuptiality and marital fertility in the determination of the tempo of fertility. Associated with each age at marriage specific cumulated fertility ratio GC(e, k), there is a mean duration of fertility

$$DC(e,k) = \frac{\sum_{y=0}^{10-k-e} (y \cdot g(y,e,k))}{\sum_{y=0}^{10-k-e} g(y,e,k)}$$

Since age is the sum of age at marriage and marital duration, the mean age of fertility is

$$AC(k) = \frac{\sum_{e=3}^{10-k} ((e + DC(e,k)) \cdot nc(e,k) \cdot GC(e,k))}{\sum_{e=3}^{10-k} (nc(e,k) \cdot GC(e,k))}$$

$$= \frac{\sum_{e=3}^{10-k} (e \cdot nc(e,k) \cdot GC(e,k))}{FC(k)} + \frac{\sum_{e=3}^{10-k} (DC(e,k) \cdot nc(e,k) \cdot GC(e,k))}{FC(k)}$$

This partitions the mean age of fertility into that part contributed by nuptiality (say, MC(k)) and that part contributed by marital fertility (say, DC(k)):

$$AC(k) = MC(k) + DC(k)$$

Note that MC(k) is not the straightforward mean age at marriage, but rather a weighted mean, where the weights are the respective cumulated fertilities associated with each marriage age.

For real cohort $k = 1$, the values of these measures are: $AC(1) = 5.988$, $MC(1) = 3.916$ and $DC(1) = 2.072$. Converting these into the decoded values (by multiplying by five), one has a weighted mean age at marriage of 19.58, and a weighted mean duration of fertility of 10.36, for a mean age of fertility of 29.94.

Comparable formulae exist for synthetic cohorts. For synthetic cohort $i = 10$, the values are: $AP(10) = 6.323$, $MP(10) = 4.209$, and $DP(10) = 2.114$. Decoding, one has a weighted mean age at marriage of 21.04, a weighted mean duration of fertility of 10.57, for a mean age of fertility of 31.62.

Thus in this illustration, the responsiblity for the rise in the mean age of fertility is largely associated with the age at marriage. Note however that the mean duration of fertility for the real cohort (10.36 years) is associated with a mean marital parity of 3.621, whereas the mean duration of fertility for the synthetic cohort (10.57 years), although almost the same, is associated with a much smaller mean marital parity, 2.856. Since one would expect a higher mean marital duration with a higher parity, the implication is that the interval between births has been appreciably longer for the synthetic cohort than for the real cohort. The question can only be answered with more data, as explained in the next chapter.

4.4 DISTORTION AND SPECIFICITY

In the preceding chapter, there was a discussion of the non-comparability of real and synthetic cohort measures because of distributional distortion, when working with measures specific only for age. The question deserves reconsideration for the measures just presented, since they are specific for age at marriage and marital duration.

The proposed index has a nuptiality component and a marital fertility component. Considering the latter first, the basic data set that is summarized has, for each marriage age, a triangle of fertility by cohort and by period, in short, the same configuration found with age-specific fertility, in the preceding chapter. The problems of distributional distortion arise essentially because quantum variations with respect to one mode of temporal aggregation appear as tempo variations with respect to the other mode, and vice versa.

The problem persists whatever the level of specificity. A changing distribution of fertility by marital duration, for a marriage subcohort, produces a distortion in the cumulated fertility measure calculated for the period, for that marriage age. A quantum variation from period to period in marital fertility will produce little change in the quantum of cohort fertility unless it is of the non-compensatory type. There is nothing about the process of augmenting specificity that affects the basic principle of the asymmetric relationship between cohort quantum and tempo on the one hand and period tempo and quantum on the other.

The same issues can be repeated for nuptiality. In this case, the relevant surface is the p(e, k) table, the celibate survival ratios. This too is a cohort by period matrix. In this case, however, there is a further wrinkle, since the way the basic element is used in the calculations is multiplicative rather than additive. This is called a risk function, since those at risk of the event are erased from the denominator when they experience the event. Analysis of distortion with a risk function is more complex than with an additive process, but some generalizations can be drawn.

In the first place, the interest in nuptiality generally, and for the present exercise, is the tempo rather than the quantum of nuptiality. Were one working with an additive function, one would inquire into the extent to which change in the quantum of cohort nuptiality was occurring, since that is what would affect the tempo of period nuptiality. Since the quantum of cohort nuptiality ordinarily changes little over time, the inference is that this source of distortion is typically unimportant.

In the second place, consider the consequences of a period-specific perturbation in nuptiality. (Short-term variations in nuptiality are much greater than short-term variations in marital fertility.) Suppose there was a tendency for the marriage rates of the single at every age (in every cohort) to rise, because the period was propitious for marriage. If one were working with an additive function, the result, as outlined in the preceding chapter, would simply be a fluctuation in the quantum of period nuptiality. With a multiplicative function, on the other hand, the consequence of rising marriage probabilities is to shift the entire age distribution of marriages towards the younger ages. A phenomenon which is almost pure quantum becomes transmuted in its period nuptiality form into a tempo phenomenon. The converse is the case for a period unpropitious for marriage. Not only is it unsatisfactory to have an index which makes one phenomenon appear to be another, but the consequences extend to the fertility index as well. Since the age distribution of marriages is the weighting system for age at marriage specific cumulated fertilities, and since fertility tends to vary inversely with age at marriage, a younger age distribution of marriages, produced by a period-specific perturbation in the quantum of nuptiality, would lead to a higher fertility index.

Accordingly, we conclude that, in some circumstances, the increase in the level of specificity employed in calculating synthetic cohort fertility aggravates rather than ameliorates the problem of distributional distortion.

4.5 SOME DETAILS ON FERTILITY AND NUPTIALITY

Restriction of the universe to ever-married women omits some of the fertility of a population. Consider again a sample in which all women (under some age limit) are interviewed.

The element of the total fertility rate, on an age-specific basis, $f(i, k) = B(i, k)/N(k)$ can be thought of as having two components in the numerator, births to the ever-married and births to the never-married, and two components in the denominator, the ever-married women and the never-married women (at the time of the interview). Evidently the age-specific birth rate is a weighted average of two kinds of fertility. The weights, the proportions ever married and never married, are both defined by the celibate survival ratios, $p(e, k)$. There may be interaction: for example, a rise in the celibate survival ratios may be accompanied by a rise in the fertility of the never-married. Another implication is that, when the histories of two successive birth cohorts are compared, in a survey restricted to the ever-married, one should exclude those married for the first time in the final period for the earlier birth cohort, not only because that age at marriage is unrepresented in the record for the later birth cohort but also because those newly married women bring with them into the record whatever premarital births have occurred to them, and this cannot be the case for the latter birth cohort.

To this point, the discussion of nuptiality has been restricted to the occurrence of first marriage. A comprehensive treatment of the influence of nuptiality on fertility would require consideration of the pattern of marital dissolution, as a consequence of widowhood or divorce, and a distinction between the fertility of the still-married and the post-married at any age. The fertility of the post-married, in turn, could be dichotomized into the experience of those who do and those who do not remarry, with an accompanying remarriage function to provide an appropriate weighting scheme. The record would be further complicated by consideration of second dissolutions, second remarriages, and so forth.

While there is no formal obstacle to the adaptation of the procedures outlined above to production of appropriate subdivisions of the reproductive history to accommodate such details, and to the development of synthetic cohort analogues for each component, the result would be cumbersome, and the denominators for the new tabulations often so small as to drown out the result in sampling error. The preceding account must be faulted for incomplete coverage of the role of nuptiality as a component of fertility and for providing no guideline for satisfactory treatment of what is omitted. There are additional reasons for regarding the proposed calculations as only a first step in the analysis of the determinants of marital fertility. As one proceeds with inquiry into the instrumental variables (such as contraceptive use and efficacy, lactation and the like), questions of exposure to risk must be confronted in a much more direct fashion than above. In our judgement, the further consideration of topics like dissolution and remarriage deserve the same kind of scrutiny as the instrumental variables, and with a comparable approach to measurement, distinct from that outlined above.

4.6 CONCLUSION

In a survey of women who are ever married (and thus members of one or another marriage cohort) and under the age of 50 (and thus members of one or another birth cohort), two issues for fertility measurement need to be resolved: (1) Should the temporal format of the inquiry be based on birth cohorts or on marriage cohorts? (2) What is the appropriate level of specificity? Marriage cohorts appear to be an attractive alternative at first glance. The record for a woman begins with marriage, an event closely identified with the inception of exposure to risk of fertility, and proceeds with a summary of births, period by period, up to interview, essentially a duration-specific record. But each marriage cohort has a distribution by age at marriage (by birth cohort) which is crude in the sense that it reflects the age distribution of the population. To eliminate that undesirable feature, one is obliged to calculate fertility rates specific for age at marriage. Furthermore, the configuration of the data set imposes a distinctive limit on marriage age for each successive marriage cohort. Accordingly some evidence must be sacrificed to achieve temporal comparability. Finally, no comprehensive index can be devised which is justifiable both for real and for synthetic marriage cohorts, because of the cross-sectional character of the distribution of real marriage cohorts by age at marriage. In brief, the marriage cohort orientation would only be acceptable if one could convince oneself that age at marriage has little to do with fertility.

Several alternative forms of birth cohort measure have been presented. All forms yield the same index for real, but not for synthetic, cohorts. The issue really turns on the analytic utility of the components of the indices. The only construction which answers the question of the role played by nuptiality (more strictly first marriage) in marital fertility is the weighted average of fertility cumulated over marital durations, for marriage subcohorts of a birth cohort, where the weights are the distribution of the cohort by age at marriage. Although the number of calculations required may on first acquaintance seem excessive, there are actually, in the quinquennial format, only 7 values of birth cohort size, N(k), 28 values of marriages by period of occurrence and period of respondent's birth, M(j, k), and 84 values of births by period of occurrence, and period of respondent's marriage and birth, B(i, j, k), to exhaust the entire output of the survey. The remaining calculations are straightforward combinations of these 119 values.

The synthetic cohort index depends on the level of specificity of the fertility element used in its construction. As a general principle, the result for the higher level of specificity is preferred, not only for the detail provided, but also because the index for any lower level of specificity may be regarded as less coherent as a distributional sequence, and less effective in excluding the influence of the past history of the constituent cohort elements, as reflected in their distributions on the unspecified variables, in the period in question. On the other hand, the specification of age at marriage makes the synthetic cohort index peculiarly sensitive to short-term disturbances.

In this chapter we have tried to provide a reasonable answer to the question of the index of marital fertility most suitable for distinguishing the respective roles of nuptiality and marital fertility in the behaviour of real cohorts, and for characterizing the experience in a period, taking cognizance of the censoring implications of the definition of the universe. The outcome is incomplete in its treatment of the fertility of the never-married, and of the consequences for fertility of periods of non-marriage subsequent to first marriage. The former topic is inaccessible with the universe as defined; the latter topic is too complicated to be handled in a simple summary fashion. The outcome is also unsatisfactory in the sense that distributional distortion makes synthetic cohort indices non-comparable with real cohort indices, but that is an intractable feature of any analysis in which both modes of temporal aggregation are employed.

5 Measures of Fertility Specific for Parity and Interval

5.1 PARITY-SPECIFIC FERTILITY

One particular objective of the WFS is the study of changes in fertility which may be occurring as a consequence of intentional interference with the reproductive process. One can consider fertility regulation as divided into the regulation of the quantum of fertility, and the regulation of its tempo. Here we are concerned with the former.

The essential idea of regulation of the quantum of fertility is action taken to terminate childbearing at a particular parity, ie to introduce a difference between fertility at one parity and fertility at the next, conditional on the attained parity. Quantum regulation is parity-dependent reproductive behaviour. It has a proportional effect on all fertility beyond the parity in question (but obviously no effect on fertility prior to that parity). The study of quantum regulation evidently requires consideration of reproductive behaviour parity by parity rather than, as in the foregoing chapter, comprehensively across the life cycle.

The completed parity distribution of a real birth cohort is a straightforward tabulation of women by total number of births. The total fertility rate, considered in the preceding chapter, is the mean of that distribution. Before considering alternative ways to measure parity-specific fertility, a digression is required on the definition of parity, and of birth order. There are several possibilities; the choice among them depends on the kind of research one is undertaking. Thus the most direct link with the concept of fertility needed in models of population growth is the number of live births. If, on the other hand, one were interested in parity as a determinant of a reproductive decision, a better choice might be the number of co-resident living children. If one were interested in measuring the extent of failure in contraceptive efforts, or the physiological sequelae of pregnancy, the definition of choice would be the total number of pregnancies. However, in practice, the count of infertile pregnancies (those not ending in a live birth) is likely to be so incomplete that the conventional practice is to anchor the record to pregnancies which end in at least one live birth, so-called fertile pregnancies. Were it not for the occasional cases of multiple live births, this would be identical with the number of live births.

The definition selected here, in anticipation of interest in the analysis of birth intervals, is the fertile pregnancy, a concept which avoids the occurrence of intervals of zero length (as between twins) as well as the incompletely recorded cases of infertile pregnancies. In the subsequent account, parity will signify number of fertile pregnancies, and birth order the fertile pregnancy order. By implication, multiple births in the total fertility count will be scored as one fertile pregnancy. Measures keyed to this concept can readily be modified (by a correction multiplier) at any subsequent stage, to restore the concept of fertility appropriate for population growth models.

Although the parity distribution is a useful representation of the outcome of reproductive behaviour, suitable for studying the consequences of fertility, it is unsatisfactory for studying the determinants of that behaviour. The process of quantum regulation can be viewed as a series of steps by which the eventual parity is attained or, thinking collectively, as a progressive development of the parity distribution, based on choices made at each parity, and manifest in modification of subsequent reproductive behaviour. With this orientation, the measure of choice is the parity progression ratio, the proportion of women with a birth of any particular order who at least have a birth of the next higher order.

If the distribution of the members of a birth cohort by completed parity x is d(x), then the number of births of order x per woman is

$$F(x) = \sum_{i=x}^{w} d(i)$$

where w is the highest parity.

The parity progression ratio for parity x is

$$R(x) = F(x + 1)/F(x).$$

The symbol F(x) is chosen because it is a total fertility rate for births of order x.

$$F = \sum_{i=1}^{w} F(x) = \sum_{i=0}^{w-1} \left(\prod_{x=0}^{i} R(x) \right)$$

Although at first glance this would seem to be an untoward expression, it is in fact isomorphic with one of the most familiar expressions in demography, the expectation of life at birth

$$e_0 = \sum_{i=0}^{w-1} \left(\prod_{x=0}^{i} (L_{x+1}/L_x) \right)$$

In the life table, progression (survival) is measured from age to age. The mean age at death (e_0) is obtained by summing the successive products, ie the persons lèft alive by the process. In the fertility schedule, progression is measured from parity to parity. The mean parity (F) is achieved in exactly the same way, ie by asking what proportion progresses, and thus determining what proportion remains in the parity in question. The element common to the two calculations is the probability that a person will persist in an initial state.

To continue with the illustrative survey used in the preceding two chapters, we have the following new data for birth cohort k = 1.

Total fertility rates by order, F(x, k), parity distribution, d(x, k), and parity progression ratios, R(x, k), all per thousand (cohort k = 1)

x	F(x, k)	d(x, k)	R(x, k)
0	954	46	954
		105	895
		59	938
1	895	148	835
2	747	168	775
3	579	157	729
4	422	106	749
5	316		
		316	611
6+	496		

The table requires several clarifications

1 There is an entry for $F(0, k) = 0.954$. This is the proportion of the cohort married (by interview), previously symbolized by EC(k). It is convenient to regard the marriages as part of the birth sequence, as if they were births of zero order. Thus the women of zero parity in the d(x) column are divided in two groups, those who never marry and those who marry but remain infertile. In a more realistic example, some of those who never marry may be in parities higher than zero.
2 In the parity progression ratio column, $R(0, k) = 0.895$ is divided into the proportion marrying (0.954) and, of those, the proportion fertile (0.938).
3 The values in the d(x, k) column arise from successive subtractions on the F(x, k) column. The proportion in parities 5 + is the same as the births (per woman) of fifth order.
4 The values in the R(x, k) column arise from successive divisions on the F(x, k) column. The value of R(5 +) is obtained by calculating $F(6+, k)/(F(5, k) + F(6+, k))$. The sense of this calculation is that, with a fixed progression ratio for each parity beyond five, at some value R', the value of F(6 +) would be $F(5) \cdot R'/(1 - R')$; the expression for R(5 +, k) follows. In some populations, it may be worth while to extend parity specificity beyond parity five.

One particular advantage of parity progression ratios in studying quantum regulation is that the changes associated with a fertility transition tend to be concentrated in a narrow range of progression ratios for the intermediate parities (in the neighbourhood of mean completed parity). If, say, one-half of the births to a cohort are in the first three birth orders, then a decline of 20 per cent in R(3, k) means a decline of 10 per cent in F(k). This indicates that the parity progression ratio is a more sensitive index of quantum regulation than the total fertility rate. Such sensitivity is an advantage not only in detecting small changes in fertility (from an overall standpoint) but also in detecting small differences among subpopulations. The two applications converge, since fertility decline is unlikely to proceed at the same pace, or at the same time, in all subdivisions of the population. It is not uncommon that one class of the population acts as pioneer, producing variance in the reproductive pattern, and other classes follow, with various lags, so that the variance is eventually reduced again. The comparison of parity progression ratios is suitable for documenting such patterns of change.

The calculation of the measures in the above table, although routine for a real birth cohort, is somewhat tedious for the synthetic cohort analogue. The principle behind the synthetic cohort construction is to identify the basic behaviour by a cohort within each period of its reproductive history that leads to the observed outcome, and then reassemble the information in sequence for the synthetic cohort, based on the behaviour of successive real cohorts within that period.

Let B(x, i, k) be the births of order x in period i to cohort k, and BB(x, i, k) be the births of order x in periods up to and including i.

Then the number of women of cohort k who are in parity x at the beginning of period i are $BB(x, i-1, k) - BB(x+1, i-1, k)$ and, at the end of period i, are $BB(x, i, k) - BB(x+1, i, k)$. Exposure to the risk of an x + 1th birth in the period can be estimated by averaging these two values. Then the fertility rate specific for parity, period and cohort is

$$f(x, i, k) = \frac{B(x+1, i, k)}{((BB(x, i-1, k) - BB(x+1, i-1, k) + BB(x, i, k) - BB(x+1, i, k))/2}$$

Given these rates, one can develop an algorithm for the successive values of BB(x + 1, i, k) as follows:

$$BB(x+1,i,k) = \frac{((f(x,i,k)\cdot(BB(x,i-1,k)+BB(x,i,k)))+((2-f(x,i,k))\cdot BB(x+1,i-1,k)))}{(2+f(x,i,k))}$$

Thus, with marriages symbolized B(0, i, k), one can use the values p(e, k), the celibate survival ratios, as in the preceding chapter, to create the series BB(0, i, 10 − i). These, in combination with the observed f(0, i, 10 − i), are used in the algorithm, step by step from the latest to the earliest cohort (from the youngest to the oldest age), to create the series BB(1, i, 10 − i). These, in combination with the observed f(1, i, 10 − i), are similarly used in the algorithm to create the series BB(2, i, 10 − i), and so forth.

With this procedure, one can develop for a synthetic cohort a parity distribution and a set of parity progression ratios, as in the preceding table. Fertility rates which are specific for period and cohort, but not for parity (the age-specific rates used in chapter 3, table 2), may be thought of as weighted averages of these parity-specific rates, the weights being the parity distribution of the cohort in the period in question. Rates which are not parity-specific reflect not only the behaviour of the cohort in the period, but its behaviour in previous times (the events that produced its parity distribution); they are less purely contemporaneous than the parity-specific rates. From another standpoint, the parity distribution of one cohort, at the end of the period in question, will ordinarily differ from the parity distribution of the next earlier cohort, at the beginning of the period in question. Since the synthetic cohort model treats these as a single history, it follows that fertility rates which are not parity specific contain sequential incoherence with respect to the parity distribution. The general principle with respect to the increase in the specificity of fertility rates, in constructing synthetic cohort histories, is that one achieves an increase in the sequential coherence of the distribution, with respect to the newly specified variable, and an increase in contemporaneity (by purging the experience of that particular effect of the real cohort's past history).

This type of tabulation is not shown here because there is a preferable alternative to be provided below. Synthetic cohort parity distributions have rarely been constructed in this way, and there are cogent reasons against doing so. The same kind of argument supporting the employment of fertility rates specific for parity, in preference to those which are not, can be extended to the discredit of those rates as well, since they too leave uncontrolled other relevant aspects of the cohort history.

In particular, for each cohort in each period, the exposure within a particular parity has a distribution by interval since entry into that parity. The probability of a woman progressing from one parity to the next is highly dependent on that interval length. Exposure to risk, in the gross sense of the total length of time elapsing between one birth and the next, contains a phase following delivery in which the probability of conception may be very low because of post-partum amenorrhoea, which may be lengthened by lactation and is often accompanied by abstinence, and another phase in which the probability of conception in the aggregate may be expected to decline with increasing interval length, because there is likely to be heterogeneity with respect to fertility, within the aggregate. Those with higher fertility will be removed from exposure to risk at shorter intervals than those with lower fertility, systematically decreasing the aggregate fertility level as the interval lengthens. As in a population with a high proportion of couples using contraception successfully to delay the birth of the next child, the probability of progression is evidently linked to the length of the intended interval.

Since the parity-specific fertility rate depends in all these ways on the interval distribution within the parity, and since that interval distribution is a function of the past experience of the cohort, it follows that the calculation of parity-specific fertility rates, as a basis for constructing a synthetic cohort parity distribution, is unsuccessful in achieving the contemporaneity and sequential coherence which are the desiderata of the synthetic cohort concept.

5.2 INTERVAL SPECIFICITY

We have indicated various respects in which fertility is highly dependent on the length of the interval since the preceding birth. Beyond regulation of the quantum of fertility, it is evident that countries differ markedly in the extent to which they regulate the tempo of fertility, either intentionally or unintentionally, according to their own cultural practices. The behaviour in question may be the use of contraception to delay the next birth, or traditional lactation practices, or extended periods of abstinence, perhaps associated with labour migration or military service, or customs associated with marital dissolution and remarriage. Such behaviour may or may not be parity-dependent. On the other hand, there is a sense in which the regulation of the tempo of fertility intrinsically implies the regulation of its quantum, because the reproductive span is finite.

Since the subject is of interest and importance, it is reasonable in principle to consider the desirability of making fertility measurements which are specific not only for parity but also for the length of time since entry into that parity. A registration system may yield records of births which are specific for interval as well as age and parity, perhaps as a record of the time of occurrence of the immediately preceding birth to the mother, as well as the present birth. The practical problem with this approach to the measurement of interval-specific fertility is determination of the amount of exposure to risk of a birth so identified. If there were regular and frequent enumerations, in which each woman provided a record of the births she had experienced, and their dates of occurrence, the required exposure denominators would be generated. Alternatively, if a registration system with the requisite data had been in existence for a long time, it would be feasible to keep a running account, period by period, of the births occurring to each cohort, by order, with the added dimension of the length of time since the preceding birth.

Lack of the requisite data has prevented such an account being produced for any population. From this viewpoint, there is evident attractiveness to the output of a fertility survey, since the record of births by time of occurrence, for each respondent, can readily be employed to determine not only the numerators but the denominators of rates specific for interval as well as parity. If the survey is regarded as a substitute for a registration system, designed to yield birth rates at one or another level of specificity, the way to proceed is straightforward.

We know of no example of fertility rates specific for interval as well as age and parity, based on registration data. Although there are a few sets of such rates, based on enumeration or survey data, they are presented in the registration style, with a focus on the outcome, and a regressive movement back to the antecedent conditions, using an arbitrary grid of regularly spaced points in the overall life cycle.

The natural way to organize such information for a cohort is to begin at the beginning, and consider how each succeeding event is distributed over time, relative to the time of occurrence of the preceding event. In considering the outcome of exposure to risk, the natural order is to demarcate successive phases of the experience by events which change the exposure status, rather than by an arbitrary array of regularly spaced time points (like age). The reproductive process consists of a sequence of events, each representing an irreversible change of state. Before the fact, the circumstances determining whether or not a birth of order x will occur, and, if so, when, are various, and chance plays a large part. After the fact, the occurrence signals a new beginning, a new phase in the respondent's life. It is not merely coincidental that the way in which information is collected from respondents involves reference points in their lives, with a structure of questions oriented to the intervals between successive vital events, rather than particular periods of calendar time; the latter provoke problems of recall precisely because they lack the salience of the former.

Consider the role played by age in demographic analysis, perhaps the commonest variable apart from time itself. One basis for its ubiquity is its role in identifying periods for cohorts and cohorts for periods. In that respect, it is the pivotal concept in the component projection of a population, and the requisite for a synthetic cohort construction. But the point concerns age as an indicator of life-cycle stage, and of associated socio-economic, socio-psychological and physiological changes. Although one would scarcely gainsay the value of age to mark the approximate beginning and end of the reproductive span, it is at least questionable whether, in the interim, there would be much left of the potency of age as a correlate of reproductive behaviour, if those more immediately relevant items of information, such as time of departure from the educational system, time of marriage, or time of occurrence of the last intended child, were present in the analysis.

The procedure for analysis of the experience of a cohort, parity by parity, is a routine statistical exercise. However, as noted before, the only comprehensive record for a cohort is the record which is furthest in the past and the most subject to problems of reliability of recall. The data for the most recent synthetic cohort, while responsive to those difficulties, present the different kinds of problem associated with distributional distortion and lack of conceptual justification. Accordingly, it is proposed that we seek our way out of this dilemma by tailoring our ambitions realistically to the stock of information provided by the survey, and give the best available description of each successive phase of the reproductive process, eschewing the ambition of ending up with a parametrization of an entire history. No doubt such efforts will and should be made, but they differ in kind from the task of analysing the results of a survey, relying as they do on information from various sources, and assumptions to fill in the missing pieces of the puzzle. The one concerted attempt to make such model construction routine – for that is what the synthetic cohort algorithm represents – has not served us well.

5.3 PROGRESSIVE FERTILITY ANALYSIS

We have been presenting an argument for developing measures tailored to survey data rather than borrowed from the registration style of measurement, for the cohort rather than the period mode of temporal aggregation, for the use of parity rather than age (or marital duration) to identify successive reproductive phases, and for efforts to characterize behaviour in each phase, viewed as a subject of interest in its own right, at least as a complement to the generally frustrating efforts to encompass entire reproductive histories with a few indices. We have termed this new approach to measurement progressive fertility analysis because its distinguishing characteristic is the systematic orientation to the progression from one parity to the next.

The reproductive history for a cohort aggregate constitutes a contingent sequence. As the members of a birth cohort advance from period to period, they establish a succession of marriage subcohorts. Each of these, in turn, advances from period to period, establishing a further dimension of parity one subcohorts, and so forth. The sequence in general is responsive to the basic demographic principle of studying the relationship between two events, one of which is a necessary but not sufficient condition for the occurrence of the other. (The consequences of the circumstance that this is not strictly true for the relationship between marriage and first birth are noted subsequently.)

At the outset of this work, we identified the respondent's record in terms of a basic set of behavioural data: date of birth and of first marriage of respondent, and the dates of occurrence of each of her x births, together with date of interview. From a standpoint of specificity, one could, in principle, devise a scheme of analysis based on births of order x, specific for their date of occurrence, and the dates of occurrence of the $x - 1$ preceding births, as well as the mother's date of birth and marriage. Since such an exercise would be practicable only with a very large sample, it is incumbent upon us to make some selections.

As a minimum, the period of occurrence of birth is needed because the data set is censored by interview date, and the period of the respondent's birth is needed because the data set is censored by a maximum age criterion. Another justification for these two specifications is that they serve as surrogates for environment and experience respectively. Moreover, in conjunction, they identify age implicitly, and provide the basis for component projections.

Beyond the specification of cohort and period, we have argued above for inclusion of parity (x) and interval length. On the same logic that we coded period and cohort, and left age coded implicitly, we choose to code the period of occurrence of the preceding birth (j) as well as of the birth in question (i), and leave interval length ($y = i - j$) coded implicitly. With the combination of period of occurrence of preceding birth, and period of respondent's birth (k), we also have an implicit coding of age at occurrence of preceding birth ($e = j - k$), or what will be termed entry age. We have already made extensive use of one particular entry age, the age at first marriage. Entry age is a generalization of marriage age throughout the parity progression, and plays a similar role at each stage.

The form of record required is shown in table 7. The initial division of the table is by birth cohort. Each progression provides births of a given order, classified by their period of occurrence and the period of occurrence of the preceding birth. As a convenient device, marriages are considered as zero order births. The data are for the same artificial population used in the preceding chapters. Note that, for any birth cohort, the summation of the values in each column, for one parity, provides the row marginals for the next parity.

There is one small bit of unreality in the table. Because of the way in which the data were devised, no births occur in the third or higher period after entry. Such births would be quite rare since the implicit interval limit in the table is 10/15 years. Such lengthy intervals may occasionally arise, perhaps because of a lengthy hiatus between marriages, or as a reflection of misreporting (omitting one or more births because the child in question died or was adopted out). It should not be implied that the actual record be truncated at all. No procedure is affected by the presence of small non-zero entries in the upper right-hand part of each panel. Incidentally, one of the attractive features of this kind of specification is the limited interval length required for each progression, to encompass almost all of the subsequent experience within that parity.

In this form we can observe the reproductive history of a birth cohort, beginning with a distribution of (first) marriages over time, the left-hand colum of the uppermost panel. Each marriage subcohort, in turn, generates a distribution of first births over time. In cross-section, those first births constitute a series of parity one subcohorts, each of which generates a distribution of second births over time, and so forth.

Table 7 Births by order (x) and period of occurrence (j), B(x, j, k), and births of the next order (x + 1) by period of occurrence (i) and by period of occurrence of the preceding birth (j), B(x + 1, i, j, k) for birth cohort k

k = 1			B(x + 1, i, j, k)						
x	j	B(x, j, k)	i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
0	4	200	62	120	18				
	5	608		162	358	58			
	6	96			22	55	10		
	7	29				5	14	2	
	8	13					1	5	1
	9	6						1	1
	10	2							0
1	4	62	20	34	5				
	5	282		79	152	23			
	6	398			95	211	34		
	7	118				20	51	9	
	8	25					3	8	1
	9	8						1	1
	10	2							0
2	4	20	7	10	1				
	5	113		34	58	8			
	6	252			66	129	19		
	7	254				54	120	19	
	8	88					12	31	5
	9	18						3	3
	10	2							0
3	4	7	3	3	0				
	5	44		14	22	3			
	6	125			36	61	8		
	7	191				47	92	14	
	8	151					27	60	10
	9	53						11	10
	10	8							1
4	4	3	1	1	0				
	5	17		6	8	1			
	6	58			19	29	4		
	7	111				32	54	7	
	8	127					31	61	9
	9	85						27	21
	10	21							5
5+	4	1	0	0	0				
	5	11		4	4	0			
	6	47			16	22	3		
	7	118				34	56	7	
	8	194					46	94	13
	9	285						89	70
	10	156							38

Table 7 (continued)

k = 2			B(x + 1, i, j, k)						
x	j	B(x, j, k)	i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
0	5	198		59	120	18			
	6	658			162	398	62		
	7	122				24	73	12	
	8	37					6	19	3
	9	18						2	6
	10	6							1
1	5	59		18	33	4			
	6	282			73	156	23		
	7	440				94	243	37	
	8	141					22	62	10
	9	33						3	11
	10	10							1
2	5	18		6	9	1			
	6	106			29	55	7		
	7	254				58	131	19	
	8	288					55	136	21
	9	102						13	35
	10	22							3
3	5	6		2	3	0			
	6	38			11	19	2		
	7	114				29	56	7	
	8	193					43	92	13
	9	168						28	65
	10	59							12
4	5	2		1	1	0			
	6	14			5	7	1		
	7	48				14	24	3	
	8	101					27	50	7
	9	127						30	61
	10	90							28
5 +	5	1		0	0	0			
	6	9			3	5	0		
	7	37				11	20	2	
	8	99					27	48	6
	9	174						41	83
	10	266							81
k = 3									
0	6	192			55	117	17		
	7	706				159	438	66	
	8	151					26	94	15
	9	51						6	25
	10	22							2

Table 7 (continued)

k = 3 (cont.)			B(x + 1, i, j, k)						
x	j	B(x, j, k)	i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
1	6	55			16	31	4		
	7	276				66	157	22	
	8	481					89	277	41
	9	166						23	75
	10	42							4
2	6	16			5	8	1		
	7	97				24	50	6	
	8	250					49	130	18
	9	322						55	151
	10	120							14
3	6	5			2	2	0		
	7	32				9	16		
	8	100					22	49	6
	9	191						38	90
	10	183							28
4	6	2			1	1	0		
	7	11				4	6	1	
	8	38					10	20	2
	9	89						22	45
	10	124							28
5+	6	1			0	0	0		
	7	8				3	5	0	
	8	29					8	16	1
	9	76						17	39
	10	148							43
k = 4									
0	7	182				50	113	16	
	8	749					154	476	70
	9	185						26	120
	10	66							7
1	7	50				14	28	4	
	8	267					59	156	21
	9	518						83	309
	10	197							7
2	7	14				4	7	1	
	8	87					19	44	6
	9	243						41	127
	10	355							53
3	7	4				1	2	0	
	8	26					6	13	1
	9	86						16	43
	10	186							33

Table 7 (continued)

k = 4 (cont.)			B(x + 1, i, j, k)						
x	j	B(x, j, k)	i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
4	7	1				0	0	0	
	8	8					2	4	0
	9	29						7	16
	10	77							17
5+	7	0				0	0	0	
	8	3					1	2	0
	9	17						4	9
	10	53							11
k = 5									
0	8	168					44	105	15
	9	788						147	513
	10	222							25
1	8	44					12	25	3
	9	252						51	151
	10	553							74
2	8	12					3	6	1
	9	76						15	38
	10	228							32
3	8	3					1	1	0
	9	21						5	10
	10	71							11
4	8	1					0	0	0
	9	6						2	3
	10	21							4
5+	8	0					0	0	0
	9	3						1	2
	10	10							1
k = 6									
0	9	150						38	95
	10	824							137
1	9	38						10	22
	10	232							43
2	9	10						2	5
	10	65							11
3	9	2						1	1
	10	16							3
4	9	1						0	0
	10	4							1

Table 7 (continued)

k = 7			B(x + 1, i, j, k)						
x	j	B(x, j, k)	i = 4	i = 5	i = 6	i = 7	i = 8	i = 9	i = 10
0	10	128							31
1	10	31							8
2	10	8							2
3	10	2							0

With this tabular layout, one can also visualize the selection processes occurring throughout the reproductive span. First there is a selection from the marriages, differentially by period of occurrence, and thus by age at marriage, of the respective proportions who progress to parity one, together with the length of time that progression takes. Then there is a selection from the first births, again differentially by period of occurrence and thus by age at first birth, of the proportions who progress to parity two, and the length of time that progression takes. The longer the preceding interval, in each case, the higher the age of entry into the next parity; the higher the entry age, generally speaking, the lower the proportion progressing beyond that parity and the longer they take to make that progression. This is behavioural selection: it is a manifestation of the continuing implications of previous behaviour. Because of the implications of selection, it is important, in considering the phases of the reproductive process, to distinguish between the role played by the entry age distribution (a consequence of behaviour in lower parities) and the role played by the entry age specific behaviour (a consequence of what happens in the parity in question).

Within a parity, there is also selection by exit time, so that those with higher fertility are not only more likely to be represented in the next parity, but also more likely to be represented in the earlier entry ages for that parity. This holds whether the higher and lower fertility is based on a selection for fecundity or for reproductive intention or for contraceptive efficacy, and whether the effect of age is conceived as primarily biological or primarily intention-modifying.

On the assumption that there is individual continuity in the characteristics associated with instrumental variables like fecundability, lactation and contraceptive use, there is a negative correlation between the probability of closing an interval, and the lengths of preceding intervals (as summarized in the combination of parity and age at entry into that parity); there is a positive correlation between the length of the interval, if closed, and the lengths of preceding intervals. Notwithstanding this, there is also a strong chance element.

5.4 PARENTHESIS ON FIRST ORDER BIRTHS

The general procedure for x + 1th order births is keyed to the consideration that an xth birth is necessary for the occurrence of an x + 1th birth. This is not so for the relationship between first births and marriages. Cultures differ in the extent to which their members adhere to a normative prescription that exposure to risk of conception must not precede marriage. A suggestive but incomplete indication of the incidence of premarital exposure to risk is the occurrence of births before or shortly after marriage. In populations where that is frequent, the definition of the universe is often modified so that there is no marital status stipulation as a prerequisite for interview. In other cultural contexts, there may be general adherence to the rule against premarital exposure, but the marriage date may precede the beginning of exposure to risk; in other words the marriage date signals the time of contractual commitment. While there is no barrier to the measurement of the interval from marriage, however defined, to first birth, it is clear that the outcome must be interpreted in the light of the cultural context.

If the universe is defined as ever-married women, births which occur in the period prior to the period of first marriage may, as proposed in chapter 2, be coded as if they occurred in the period of first marriage. Should their incidence be appreciable, there is no problem about incorporating a negative length of interval $y = i - j$. Although marriage is not a necessary condition for the occurrence of birth, the relevant question is the extent to which there is a relationship between the distribution of marriages by age and the distribution of first births by age, and there can be little doubt that the relationship is strong everywhere.

The date of marriage may be subject to more severe recall problems than the dates of children's births. Not only is the event ordinarily more remote in time than parenthood, but it may be misstated to prevent the inference that the norms against premarital exposure have been violated. Moreover, a birth (provided the child survives) yields continuing tangible evidence in terms of the stage of growth of the child for the time of its occurrence, whereas there is no such clue with respect to marital duration, and especially if there is no annual recognition of the anniversary. Given these considerations, one should not expect the quality of information for the first stage in the progression sequence to be of as high quality as for subsequent stages.

If the dating of marriage should be highly suspect, or if the universe has been defined as all women, there is no obstacle to deleting the first panel from table 7 altogether. Entirely apart from the foregoing considerations, there will necessarily be a discontinuity between the data for the progression from marriage to first birth, and the data for all subsequent progressions since, unlike the latter, the first interval does not begin with post-partum amenorrhoea, and the possibility of lactation. Indeed, in view of the sharp contrast between the cultural distinctiveness of marriage in all its diverse meanings, and the universality of the reproductive experience, there is a good case for conducting progressive analysis as two separate issues: the logical progression from births of lower to births of higher order, on the one hand; and the much more complex and problematic relationship between nuptiality and first order births on the other.

5.5 SUMMARY INDICES BY PARITY, ENTRY AGE AND INTERVAL

To express the output in symbolic terms, we propose the following. For women of birth cohort k with a birth of order x in period j, the number of births of order x + 1 in period i is $B(x + 1, i, j, k)$. Where only three labels are used, say $B(x, j, k)$, the referent is to births of order x occurring in period j to members of birth cohort k. As in preceding chapters, the label i is used for the period of occurrence of the birth in question (the exit event), and the label j is used for the period of entry into exposure to risk (which, in chapter 4, was marriage). Marriages, for convenience, are symbolized as births of order zero, ie $B(0, j, k)$.

For purposes of comparative analysis, in order to ensure comparability of life-cycle identification, and stay clear of censoring problems, the following implicit coding is employed: entry age $e = j - k$; interval length $y = i - j$. The parallelism with the previous treatment is evident. Thus $B(x + 1, i, j, k)$ may be thought of as $B(x + 1, k + e + y, k + e, k)$, the formulation required to identify particular values for the life-cycle variables in searching the basic data file (which is organized on the basis of identification of the periods i, j, k).

In this particular empirical illustration, the interval lengths have been artificially restricted to the values y = 0, 1, 2 (in quinquennial units). In conventional terms, these represent distributions in the form of triangles of ten-year width, with the apex at 0, 5, and 10 respectively, where the distribution for y = 0 is restricted to the positive half of the triangle.

Several possibilities exist for analysis of the reproductive history of a cohort, based on the data set in table 7. For each combination of entry age and parity, one has the number of entry births, and the associated number of exit births, period by period. The obvious quantum parameter to employ for such data is the parity progression ratio (the ratio of exits to entries), in this situation on an entry age-specific basis.

$$R(x, e, k) = \sum_{y=0}^{2} B(x+1, k+e+y, k+e, k)/B(x, k+e, k)$$

An appropriate tempo parameter is somewhat less obvious. While it would be feasible to calculate the proportion of exits occurring within each interval length, that would not evoke the passage of time. The formula chosen is simply the mean, in decoded form.

$$Y(x, e, k) = 5 \cdot \sum_{y=0}^{2} (y \cdot B(x+1, k+e+y, k+e, k)) \Big/ \sum_{y=0}^{2} B(x+1, y, k+e, k)$$

This is a crude approximation. Thus no account is taken of the circumstance that there is a necessary minimum length of birth interval, because of gestation time. On the other hand, it is unlikely that the births in the longer intervals would be distributed uniformly within those periods. More likely would be a declining pattern, as the more fertile select themselves out by progressing to the next parity. These two considerations are counterbalancing. Several models of a more sophisticated kind were considered. The assumptions underlying any model are not only matters of judgement, but would need to be varied from one cultural context to another, depending, for example, on the extent of lactation, or of contraceptive use. In the circumstances, it seems best to stay with the simple index.

There is, of course, no barrier to a more precise calculation of interval length should that be considered worthwhile. Once births have been sorted out relative to their particular preceding birth, one can calculate the mean dates of occurrence of exit and of the corresponding entries, using the century-month codes for dating. Even if that is to be done, the proposed tabulation format may be helpful to ensure that censoring bias does not intrude into particular comparisons.

Rather than calculate the indices R and Y for every available combination of entry age and parity, the tabulation may usefully be collapsed on one or the other dimension, in effect to examine the marginals. The results for cohort k = 1 are shown in tables 8 and 9, with rows and columns re-oriented as life-cycle variables. In table 8, the number of entries in the first row (1000) is the number of members of the cohort, since that plays the role of entries for those exits which are marriages. The Y entry for that row is the mean age at marriage, itself a particular kind of interval, calculated in the conventional manner from the distribution of marriages by (implicit) age. The final row provides a compact summary of the cohort's experience. In fact, the progression ratio is not a new piece of information, since it is simply the ratio FC(k)/(FC(k) + 1). On the other hand, the overall interval length, 3.84, is an index of the tempo of marital fertility, inaccessible without a formulation of the kind employed here.

Table 8 Entry births, B(x, e, k), and associated exit births, B(x + 1, y, e, k), aggregated over entry age (e) and summarized by parity progression ratios (R) and mean interval lengths (Y), for cohort k = 1

x	B(x, e, k)	B(x + 1, y, e, k)				R	Y
		y = 0	y = 1	y = 2	Sum		
	1000				954	0.954	20.14
0	954	253	553	89	895	0.938	4.08
1	895	218	457	72	747	0.835	4.02
2	747	176	351	52	579	0.775	3.93
3	579	139	248	35	422	0.729	3.77
4	422	121	174	21	316	0.749	3.42
5+	812	227	246	23	496	0.611	2.94
Sum	4409	1134	2029	292	3455	0.784	3.84

Table 9 Entry births, B(x, e, k), and associated exit births, B(x + 1, y, e, k), aggregated over parity (x), and summarized by parity progression ratios (R) and mean interval lengths (Y), for cohort k = 1

e	B(x, e, k)	B(x + 1, y, e, k)				R	Y
		y = 0	y = 1	y = 2	Sum		
3	293	93	168	24	285	0.973	3.79
4	1075	299	602	93	994	0.925	3.96
5	976	254	507	78	839	0.860	3.95
6	821	192	387	58	637	0.776	3.95
7	598	120	259	39	418	0.669	4.03
8	455	132	106		238	0.523	
9	191	44			44	0.230	

Table 9 provides the same kind of information, but for entry age, with the parity dimension collapsed. Such calculations as these may be useful not only to display the main outlines of zero order relationships (ignoring the possible interactions between parity and entry age, with respect to interval statistics) but also as a useful resort in the eventuality that the denominators of the more highly specific measures are so small as to introduce intolerable sampling variability.

5.6 PROGRESSION PROBABILITIES

The basic measure underlying the tabulations in table 7 is the progression probability, specific for parity, entry age, and interval.

$$q(x, y, e, k) = \frac{B(x + 1, k + e + y, k + e, k)}{B(x, k + e, k) - \sum_{i=0}^{y-1} B(x + 1, k + e + i, k + e, k)}$$

This is the probability that a member of cohort k who entered parity x at entry age e (and has not yet had an x + 1th birth) will have an x + 1th birth in interval y. For the first row in the basic data column for cohort k = 1 (table 7), the calculations are 62/200 = 0.310, 120/(200 − 62) = 0.870, and 18/(200 − 62 − 120) = 1.000.

The complete array of probabilities for the survey is shown in table 10. Together with the table of celibate survival ratios (table 3, in the preceding chapter), they suffice to permit the reconstruction of the entire basic data table, table 7, given only the additional information, N(k), size of birth cohort.

Since this is the case, one is in a position to use the probabilities observed in any period to define the experience of a synthetic cohort, and proceed in reverse to produce its distribution of births by order, by interval length, and by entry age, in short, to provide the same information contained in the basic data table, but for a period rather than a cohort.

To exemplify the procedure, the tabulation for period 10 is based on the progression probabilities contained in the right-hand negatively sloping diagonal of each panel in tables 3 and 10. The calculation begins with an arbitrary radix (here 1000) and proceeds as in the accompanying table.

Construction of the basic data table for synthetic cohort i = 10 from progression probabilities, for marriages and first births

e	p(e, 10−e)	B′(0,e,10)	q(0,y,e,10−e)			B′(1, y, e, 10)						
			y = 0	y = 1	y = 2							
3	0.9200	80	0.242	0.848	0.789	19	51	7				
4	0.3896	562	0.166	0.800	0.588		93	375	55			
5	0.5000	179	0.113	0.755	0.254			20	120	10		
6	0.6413	64	0.106	0.566	0.250				7	32	6	
7	0.7800	25	0.091	0.375	0.143					2	8	2
8	0.9077	8	0.167	0.200							1	1
9	0.9583	4	0.000									0
						19	144	402	182	44	15	3

The tabulation provided to exemplify the first step yields the distribution of first births by entry age, B′(1, y, e, 10), the prime being employed to signify that this is an artificial construct rather than original data. With those first births by entry age, one uses the comparable progression probabilities for parity one to develop second births, and so forth.

If one has the values B′(x, e, i) and B′(x + 1, y, e, i) obtained in this way, then a comprehensive synthetic cohort record is available for summarization in the fashion indicated in the previous section – providing parity progression ratios and mean interval lengths specific for entry age and parity, as well as a distribution of fertility by parity and age. This puts one in a position to do temporal analysis by comparing the results for the real cohort (k = 1) and the synthetic cohort (i = 10).

However, there is nothing about the procedure of making the fertility measure specific for interval that reduces the force of previous arguments concerning the distributional distortion of the synthetic cohort parameters. At whatever level of specificity, long-term changes in the tempo of cohort fertility are manifest in shifts of the quantum of period fertility, and short-term changes in the quantum of period fertility are manifest in shifts of the tempo of cohort fertility.

Table 10 Progression probabilities specific for parity (x), entry age (e), and interval (y), for birth cohort k, q(x, y, e, k) · 1000

y	x = 0	e	x = 1
y = 0	000	9	000
	167 167	8	125 100
	077 111 091	7	120 091 095
	172 162 118 106	6	169 156 139 127
	229 197 172 141 113	5	239 214 185 165 134
	266 246 225 206 187 166	4	280 259 239 221 202 185
	310 298 286 275 262 253 242	3	323 305 291 280 273 263 258
	k = 1 … i = 10		k = 1 … i =10
y = 1	200	8	143
	417 375	7	364 367
	583 613 556	6	761 521 524
	743 745 752 755	5	696 702 707 710
	803 802 801 800 800	4	749 746 748 750 751
	870 863 854 856 847 848	3	810 805 795 778 781 786
	k = 1 … i = 10		k = 1 … i = 10
y = 2	143	7	071
	200 250	6	191 175
	526 480 484	5	370 359 357
	659 633 606 588	4	451 434 415 404
	1000 947 850 842 789	3	625 500 500 500 429
	k = 1 … i = 10		k = 1 … i = 10

y	x = 2	e	x = 3
y = 0	000	9	125
	167 136	8	208 203
	136 127 117	7	179 167 153
	213 191 171 149	6	246 223 199 177
	262 228 196 169 140	5	288 254 220 186 155
	301 274 247 218 197 169	4	318 289 281 231 238 188
	350 333 313 286 250 200 250	3	429 333 400 250 333 500 000
y = 1	200	8	238
	408 393	7	484 464
	600 584 566	6	639 613 588
	694 668 647 629	5	685 659 628 614
	734 714 685 647 623	4	733 704 696 650 625
	769 750 727 700 667 625	3	750 750 667 667 500 1000
y = 2	111	7	156
	238 216	6	269 224
	333 292 254	5	281 241 207
	381 318 261 250	4	375 250 286 143
	333 333 333 333 333	3	000 000 000 000 000

y	x = 4	e	x = 5 +
y = 0	238	9	244
	318 311	8	312 305
	244 236 226	7	237 236 291
	288 267 247 221	6	288 273 224 208
	328 292 263 241 190	5	340 297 276 235 100
	353 357 364 250 333 250	4	364 333 375 333 333 000
	333 500 500 000 000 000 000	3	000 000 000 000 000 000 000

Table 10 (continued)

y = 1	362	8	357
	635 629	7	635 624
	684 676 672	6	667 667 661
	744 706 714 727	5	710 769 762 692
	727 778 857 667 750	4	571 833 1000 1000 1000
	500 1000 1000 000 000 000	3	000 000 000 000 000 000 000
y = 2	257	7	241
	280 292	6	250 250
	400 300 250	5	333 333 200
	333 500 1000 000	4	000 000 000 000
	000 000 000 000 000	3	000 000 000 000 000

One can think of the fertility of a cohort in a period as the outcome of interaction between characteristics descriptive of the cohort and its constituents (the history and experience they bring into the period) and the environment for the ensuing behaviour, the circumstances peculiar to the period. The process of specification and control, exemplified above at a refined level, does perform the assignment of arithmetical removal of the influence of the specified demographic characteristics which have accrued in the cohort's history. Yet there remains an interdependency of successive phases of the cohort experience, not captured by the calculation of probabilities, no matter how conditional.

One particular kind of interdependency becomes visible with the calculations we have just utilized. Consider the values used to develop the progression of the synthetic cohort from parity zero to parity one in the first entry age, viz, 0.242, 0.848, and 0.789. These values were contributed by three different real cohorts. Now the set of three such values, for any one cohort, although arithmetically independent, represent three readings on a single survival function: the process of selection in the first interval is part of what makes the value for second interval what it is, and so forth. The organic relationship among the probabilities for successive intervals, for the same cohort, is destroyed by the synthetic cohort construction.

Beyond such demographic considerations, there are many other respects in which cohorts differ from one another – in their distribution by years of schooling completed, in their proportion born on a farm, in their attitutes toward contraception, and so forth. To use fertility rates from a sequence of cohorts as if they constituted a history is to deny the relevance for reproductive behaviour of all such cohort-differentiating characteristics. Beyond that, the experience of any real cohort in one period is unlikely to be independent of its experience in previous periods, either in the short term, when births are displaced from an earlier to a later period because the earlier period is unpropitious for fertility, or in the long term, because the temporal pattern of fertility across the life cycle manifests a reproductive strategy applicable to the entire experience. Synthetic cohort measures destroy the meaningful sequence of life; they are unreal.

5.7 PARITY COHORT ANALYSIS

The case for the synthetic cohort is that it permits a comprehensive statement about recent reproductive behaviour; the case against the synthetic cohort is that the sequential integrity of individual histories is lost in the process. Fortunately there is an alternative. Much can be learned from examining the separate phases of the reproductive sequence, each reasonably complete in itself, and each describing recent experience, without the price of unreality.

The proposal is to study the first order fertility of marriage cohorts, the second order fertility of parity one cohorts, and so forth, for those cohorts originating in periods 8 and 9, and thus for exits in periods 8 and 9, for the former, and 9 and 10, for the latter. The exercise is to be conducted on an entry age-specific basis, for the reasons presented for marriage cohorts in chapter 4, in brief to control birth cohort size, and to ensure that censoring does not introduce non-comparability of the cohorts being compared. Of course the introduction of entry age into the design has much more than accounting utility. Age has a potent influence on interval parameters, in part because of the correlation of age with physiological change, in part because of the role age plays in the formation of reproductive intentions, and other such substantive considerations, and in part for a simple demographic reason: the combination of parity and age at entry into that parity can be thought of as creating yet another implicit code, in this case an index of past fertility. Moreover, one wants to be able to draw a distinction between that part of the outcome reflective of the distribution with which the interval begins – the heritage of previous reproductive experience – and that part attributable to behaviour in the interval in question.

As a compromise between the interest in recent behaviour, and the concern for reasonably comprehensive coverage of the parity progression, we propose that the calculations be confined to births of the next higher order which occur either in the same period as, or in the following period to, that signifying the parity cohort's origin. This takes advantage of the convenient circumstance that the very considerable majority of all progression is limited to intervals $y = 0$ and $y = 1$ (or, in conventional terms, up to interval length 5/10).

The abbreviated parity progression ratio is defined as

$$R'(x, e, j) = \frac{B(x+1, j, j, j-e) + B(x+1, j+1, j, j-e)}{B(x, j, j-e)}$$

for the parity x cohort of period j, entry age e, and the abbreviated mean interval length similarly

$$Y'(x, e, j) = \frac{5 \cdot B(x+1, j+1, j, j-e)}{B(x+1, j, j, j-e) + B(x+1, j+1, j, j-e)}$$

These are shown, for $j = 8, 9$, for $e = 3, \ldots 7$, and for all parities, in table 11.

In table 11, an average value is shown for each parity, covering the range of entry ages. That is obtained in the following way. If B(x, e, j) entries occur in entry age e, and N(k), where $k = j + e$, is cohort size, then the distribution of the parity cohort by entry age, with cohort size controlled, is

$$d(x, e, j) = \frac{B(x, e, j)/N(j+e)}{\sum_{e=3}^{7} (B(x, e, j)/N(j, e))}$$

For purposes of comparison, the entry age distribution is standardized by averaging d(x, e, 8) and d(x, e, 9) to give d(x, e). Then the required indices are:

$$R'(x, j) = \sum_{e=3}^{7} (d(x, e) \cdot R'(x, e, j))$$

and

$$Y'(x, j) = \sum_{e=3}^{7} (d(x, e) \cdot Y'(x, e, j))$$

Table 11 Abbreviated parity progression ratios, $R'(x, e, j)$, per thousand, and mean interval lengths, $Y'(x, e, j)$, for parity cohorts $j = 8, 9$, specific for parity (x) and entry age (e)

$R'(x, e, j)$ e	j	x = 0	x = 1	x = 2	x = 3	x = 4	x = 5+
3	8	887	841	750	667		
	9	887	842	700	1000		
4	8	841	805	724	731	750	1000
	9	838	802	697	714	833	1000
5	8	795	761	716	710	789	828
	9	789	757	691	686	793	765
6	8	676	596	663	699	762	758
	9	608	590	640	670	753	737
7	8	462	440	489	576	724	722
	9	444	424	471	554	717	713
Average	8	826	735	662	657	743	741
	9	820	731	638	633	739	725
$Y'(x, e, j)$							
3	8	3.52	3.38	3.33	2.50		
	9	3.57	3.44	3.57	2.50		
4	8	3.78	3.63	3.49	3.42	3.33	3.33
	9	3.89	3.74	3.58	3.33	3.00	3.33
5	8	3.92	3.78	3.63	3.45	3.33	3.33
	9	4.11	3.94	3.78	3.64	3.48	3.46
6	8	3.80	3.69	3.56	3.41	3.25	3.20
	9	4.03	3.83	3.67	3.52	3.36	3.48
7	8	4.17	3.64	3.60	3.45	3.32	3.36
	9	3.75	3.93	3.65	3.49	3.35	3.35
Average	8	3.78	3.71	3.58	3.43	3.29	3.31
	9	3.89	3.85	3.69	3.52	3.35	3.39

It may have been noted that there is a strong resemblance between the parity cohort concept and the marriage cohort concept discussed in the preceding chapter. The conclusion of that account was a recommendation against the marriage cohort procedure. The present proposal does not contradict that position. The reason for the judgment about the marriage cohort procedure was that there is no satisfactory way of achieving an age at marriage distribution for a real marriage cohort and for a synthetic marriage cohort which were appropriately analogous, as well as the obligation to make deletions from the available data set to ensure temporal comparability. In the present parity cohort proposal, there is no synthetic cohort, and the deletion, although still required, is of less moment because most progression occurs soon after entry.

Admittedly there is some recent experience not encompassed by the parameters reported in table 11. It is feasible to proceed one step further in the comparison, by calculating the progression probabilities for interval $y = 0$, for periods 9 and 10, but there is evidently no way of disentangling the quantum and tempo components of a single number. The result of the calculation may be useful, but it is limited.

We have stressed throughout this account the use of quantum and tempo indices as separable measures of reproductive performance. While the formulation is attractive – since reproductive decisions can be viewed as divisible into the question of whether or not to progress and, if so, when – the outcome we are examining reflects other characteristics of reproductive behaviour, such as the effectiveness with which an intended delay is achieved, the extent to which non-users of contraception are physiologically capable of giving birth, and so forth. The basic evidence is in fact a distribution of intervals by length, and a comprehensive model to elicit the determinants of that distribution would certainly encompass more than quantum and tempo decisions. For example, the time pattern of exit is in part a function of the heterogeneity of the group on entry, and the pattern of subsequent selection implicit in that heterogeneity. Accordingly, it is discreet to regard the indices R and Y simply as two measures of a distribution, framed in terms of the reproductive outcome of what may be a rather complicated model. That is, of course, the justification for investigation of the instrumental variables, to identify what has gone into the creation of the observed interval distribution.

Parity-cohort analysis permits one to resolve one particular problem which has plagued alternative approaches to the study of recent changes in reproductive behaviour, viz, the extent to which one or other of the periods being compared may be manifesting the characteristics of a short-term fluctuation. Since that is directly delineated by the relative frequencies of exits in successive periods, the hypothesis is precisely tested by the proposed calculations. Despite the lack of detail of information about interval lengths (a lack that can, of course, be remedied by resort to the precise century-month codes available for each birth), the procedure is recommended because it is very simple, once the basic data table has been prepared, and yet sophisticated, in the sense that it is free of bias, and capable of distinguishing between the contribution of weights (the entry age distribution) and rates (the progression probabilities) to the observed outcome.

6 Extensions

6.1 REPRISE

The principal subject of this work is the measurement of temporal variations in fertility. Two modes of temporal aggregation have been contrasted. The first, in the registration style of measurement, is oriented to the time of occurrence of births; this is the period approach. The second is oriented to the time of entry into exposure to risk of fertility; this is the cohort approach. The choice of entry event in the latter is discretionary, and may be respondent's birth or first marriage or entry into parity x. For each type there is a parallel synthetic cohort (period) construction.

Underlying the development of suitable measures has been concern about the consequences of the configuration of data provided by a fertility survey, given the criteria stipulated in the definition of the universe. Individual histories are censored by time of interview; ages at entry into exposure to risk are censored by the upper age limit; age at marriage is censored by the marital status criterion.

Within these constraints, two systems of measurement have been proposed. The first system emphasizes the distinction between nuptiality and marital fertility as determinants of fertility, not otherwise specified, throughout the reproductive span. The minimum detail required to produce measures free of avoidable censoring bias is the set of celibate survival ratios, p(e, k), where e is age at marriage and k the designation of birth cohort, and fertility rates g(y, e, k) where y is marital duration. Summary measures for any real or synthetic cohort can be constructed from these rates and ratios, within the constraints on e and y established by the universe definition. In some respects the outcome is inherently unsatisfactory.

The second system has been called progressive fertility analysis. The dimensions of the first system are enlarged to display the structure of marital fertility by parity and interval, not because one is impelled to do so from consideration of the shape of the data set, but rather because the product illuminates important features of the reproductive process inaccessible at a less detailed level. The basic measure is the progression probability, q(x, y, e, k), where x is parity, y is generalized from marital duration to years since entry into parity x, e is generalized from age at marriage to age at entry into parity x, and k is birth cohort. Although real and synthetic cohort constructions are derived, for the purpose of comprehensive summary, as with the first system, the emphasis of the account is placed on indices for each separate parity in the progression.

6.2 AGGLOMERATE MEASURES

There is no obstacle in principle to the calculation of the proposed measures for any interesting subset of the population as well as for the total population. The most attactive candidates would be those characteristics which, for any individual, remain fixed over the relevant reproductive span, such as race, religion, ethnic group, rural or urban birthplace, and, to a tolerable approximation, level of education. These encompass most of the important sources of differentials in fertility. The practical constraint on this approach is, of course, the increasing magnitude of sampling error in the components of the calculation as the denominators are diminished.

One recourse in such a situation has already been suggested by the account in the preceding chapter. Given the detailed output of parameters specific for parity and for entry age, for each cohort (or period), one can collapse the parity dimension, or the entry age dimension, or both, in the eventuality that sample size has produced intolerable irregularity in the multidimensional product. It is recommended that the detail be retained in the calculations, and then the dimensions collapsed, rather than proceeding directly with cruder measures, in order to maintain control over censoring bias, preserve the virtues of the higher level of specificity, and maintain comparability with other analyses of the same kind.

A further alternative may be proposed. For many analytic purposes, one may be willing to forego the study of temporal variations in the interest of differential fertility analysis. Although one may be able to achieve this objective by confining the calculations to the subdivisions (on the variable of interest) of the first real cohort or the last synthetic cohort, for example, that would frequently entail the embarrassment of small sample size. To maximize the yield from the available data, and at the same time achieve the objective of analytic depth free from censoring bias, we proposed the calculation of what may be termed agglomerate measures of fertility.

The simple procedure is as follows. Given the data sets p(e, k) and g(y, e, k), for example, one can simply average the p(e, k) measures for each age at marriage (e) over the available birth cohorts (k), by summing them and dividing by the number of cohorts, to give a set of values p(e), and average the g(y, e, k) measures in the same way, for each combination of duration of marriage (y) and age at marriage (e), to give a set of values g(y, e). Then the requisite summarization is based on the sets p(e) and g(y, e). Even with this step, one may be inclined to pay less attention to the detailed variations by y and e respectively than to the overall indices.

Similarly, with progressive fertility analysis, the measures g(x, y, e, k) can be averaged over the available k's, for each (x, y, e) combination, and the remaining calculations proceed on the basis of p(e) and g(x, y, e). As before, it is important that the detailed procedure be followed from the outset, and the dimensions subsequently collapsed, rather than ignoring the strictures of censoring bias in an effort to achieve a simple shortcut.

It may be objected that the proposal to construct an agglomerate index by simply averaging the rates across available cohorts is a contrived solution. Yet in fact such a contrivance is routine in the calculation of life tables, for all kinds of survival function. In the life-table procedure, there is, as a consequence of the censoring with which the procedure is designed to cope, a plenitude of intervals of the shortest length, but progressively fewer intervals available as the span lengthens. The eventual product is constituted of the average experience at each interval, based on the available cases. The same principle is followed in the agglomerate measures.

One way of regarding the outcome of the agglomerate calculation is as a compromise between real cohort and synthetic cohort calculations, since the averaging proceeds, in a sense, over the first real and the last synthetic cohort, the second real and the second to last synthetic cohort, and so forth. The implication of this interpretation is that, since synthetic cohort results are distorted versions of real cohort results, the agglomerate measures share some of the unfortunate characteristics of synthetic cohort measures; they are semi-distorted. On the other hand, if one is conducting a comparative analysis in which temporal variation is suppressed, the relevant question is not the magnitude of distortion but the magnitude of differential distortion, since that is what would prejudice the comparison. The question is not to be ignored, but it may often turn out on investigation to be of little consequence, where patterns of change are similar in the subsamples being compared.

6.3 FERTILITY MEASUREMENT AT THE INDIVIDUAL LEVEL

The account to date has been concerned solely with the construction of indices for aggregates, temporal or substantive. Aggregate analysis of this kind is a demographic tradition. No doubt one source of its persistence is the circumstance that, until recently, demographers were obliged to rely on data from secondary sources, particularly the systems of enumeration and registration. The output available for research consisted of tabulations based on decisions of others about the number and detail of the dimensions provided; there was no access to the original information at the individual level.

With the advent of the fertility survey, the information obtained from each individual is accessible on tape in exactly that form, so that the number of variables to be considered, and the detail in which each is to be coded, is at the discretion of the analyst. Accordingly, one need not feel obliged to work within the tradition of aggregate analysis, but proceed directly to the study of covariation at the individual level, using the various techniques for multivariate analysis.

To conduct analysis comparable in detail to that proposed for the consideration of nuptiality and marital fertility, the information for an individual would consist of dates of birth and first marriage, current parity, and the randomly determined age at interview. Although it may appear superficially that one could devise a comparable index for two individuals with differing ages at interview by determining a functional form for the relationship between parity and interview age (leaving aside the further complication of marriage), the critical stumbling block in such a procedure is the circumstance that differing age at interview necessarily signifies different birth cohort membership, and there can be no assurance that the life-cycle pattern is fixed over time. In short, the kinds of consideration spelled out for aggregates in the body of this work do not vanish at the individual level.

Comparable considerations arise when parity and interval are analysed. For the individual, the relevant items of information (quite apart from explanatory variables) would be entry age, birth cohort, parity and interval length, together with time of interview. For any particular parity, the dependent variable for the individual would take two forms: (1) the dummy variable of whether or not the interval is closed; (2) the length of the (closed or open) interval. Each would be conditional on the time elapsed between entry and the random event of interview. Once again, one would be obliged to respond to the implications of the definition of the universe before indulging in individual comparisons.

Except for characteristics fixed over time for the individual (which can be studied by the procedures outlined in the preceding section), efforts to study the relationships between fertility and various presumably associated variables have met with little success. Three reasons may be suggested for that regrettable outcome. (1) The basic information available for an individual is unreliable. Studies of consistency of response, based on re-interviews, uniformly show high inconsistency of response at the individual level, but a reasonable approximation of consistency at the aggregate level (suggesting the large random component to individual inconsistency). (2) Chance plays a large role in the reproductive process. Two individuals with an identical probability of conception per month are unlikely to conceive after the same length of exposure to risk. This is a matter of considerable concern, given that the level of detail required for satisfactory analysis implies small frequencies in any one class of interest (as is evident from the illustrative data provided in the preceding chapter). (3) Explanatory variables which are not fixed over time are themselves exposed to risk of censoring bias – often ignored because the date of entry into the current status is unrecorded – and are also subject to the noise associated with stochastic variables. However desirable fertility research at the individual level may be considered to be, in the abstract, chance plays a sufficiently large role in the outcome that the success of such efforts is contingent on being able to afford a very large sample.

This may be a less discouraging conclusion than it sounds. Although the predominant orientation to the analysis of the determinants of fertility is probably framed in terms of individual characteristics, a legitimate alternative conceptualization is that fertility is a collective property, the statistical manifestation of a cultural design. From this viewpoint, the individual is a representative of a culture or subculture, the reproductive behaviour reported is a contribution to the description of the reproductive pattern characteristic of that culture or subculture, and the pattern in its totality is to be interpreted in relation to other characteristics of the culture or subculture. This is not the appropriate place for an extended account of the theoretical issues involved. Suffice it to say that one may regard the output of the procedures described in this work not merely as descriptive statements about individuals, in statistical summary form, but as measurements of fertility as a collective property, appropriately studied at the macro-analytic level.

6.4 INPUTS AND OUTPUTS

The criterion for evaluating a measure, throughout this account, has been its worth as a dependent variable in the analysis of the determinants of fertility. We think of the dated events constituting a respondent's history, and summarized for her cohort, as the outputs of a model of reproductive behaviour, outputs which serve as clues to the structure of the model. The distinctive contribution of the fertility survey to this work is the collection of detailed information about the acts, practices and conditions underlying the occurrence of births in various numbers and at various times – the array of instrumental variables. Almost without exception, these variables are investigated on an interval-specific basis, typically by partitioning the time elapsed between one birth and another into components of various types, in an elaboration of the concept of exposure to risk. The measures proposed for progressive fertility analysis have the considerable advantage over conventional devices that they dovetail with the format of inquiry into those instrumental variables which, in a proximate sense, are responsible for fertility.

An important subset of the instrumental variables is fertility regulation, intentional interference with the reproductive process. The category encompasses periodic or prolonged abstinence, suppression of ovulation temporarily (by oral contraception) or permanently (by female sterilization), male sterilization, the blocking of fertilization by chemical or mechanical barriers, the inhibition of implantation by intra-uterine devices, and the induction of abortion. The premise is that, whatever may be the explanation for intentional interference in some deeper sense, that intention is formulated at least in part by reference to the respondent's life-cycle location: the parity, the age at entry into that parity, and the length of time spent in that parity. This is a further argument for progressive fertility analysis.

But there is another criterion for evaluating a fertility measure, with the direction of attention precisely the opposite of the foregoing: fertility not as a dependent variable for which one seeks determinants but as an independent variable from which flow consequences. A strong source of support for the investment in fertility surveys is the belief or faith or hope that they will make a contribution to the important task of population projection, ie models of population change in which fertility constitutes the most important input.

The process of population projection begins with a classification of the population by birth cohort, *inter alia*, at the most recent date for which the requisite information is available. Two classes of input, on a period-specific basis, suffice for the projection. The first class of input, mortality and migration, modifies the size of each birth cohort in each subsequent period. The second class of input is the reproductive performance of each birth cohort in each successive period, since that, in cross-sectional summary, is what determines the sizes of the new birth cohorts which are continually being fed into the population. It is important to recognize that this process is completely isomorphic with that displayed in the basic data table (table 7), although there the process is the generation of subcohorts at a higher parity from subcohorts at a lower parity. Progressive fertility analysis provides the elements of a fertility projection in precisely the form required.

Projection techniques differ in their level of sophistication. The most common technique is extrapolation, essentially the assumption that what has been happening will continue to happen (without consideration of the causes of past behaviour or the likelihood that those causes will prevail). But what has been happening can be described with more or less refinement or, in the terms employed in preceding chapters, at one or another level of specificity. The most common level employed (and the minimum level required for a component projection) is cohort by period, ie extrapolation of the surface of rates like those shown in table 2 above. But the same rationale for preferring more highly specific rates in the conduct of analysis applies to the conduct of projection as well. The appropriate base for extrapolation, then, is the surface of celibate survival ratios (table 3 above) and the surfaces of progression probabilities (table 10 above).

The gain from this increase in level of specificity is twofold. On the one hand, the progression probabilities (of which the celibate survival ratios are a special subset) are immediately descriptive of the ways in which people conduct their reproductive lives, so that one is directly cognizant of the behavioural implications of the extrapolated values, whereas age-specific birth rates, on the contrary, are remote from comprehension in the same terms. On the other hand, the output of the projection is enriched in several dimensions, providing a picture of the evolving distribution of women by parity and interval, for consideration of their consequences, as well as those provided by total population size, distributed by age.

Beyond mere extrapolation are projections based on particular assumptions about change in reproductive behaviour. In the minds of those making the assumptions, they may be formulated in terms of reductions in mean parity, or lengthening of birth intervals, or later age at marriage, for example, but there is no way of translating such conceptualizations into their manifestation in the period-specific outputs of cohorts without resort to one or another form of progressive fertility analysis. Similarly, most projections are based on assumptions of continuous long-term change. If one were interested, on the other hand, in the demographic consequences of an anticipated postponement of fertility, perhaps in response to a predicted period of economic hardship, the model for accomplishing that objective would likewise entail progressive fertility analysis, since the phenomenon of postponement is precisely defined as a temporary increase in mean length of interval with no change in the parity progression ratio.

A third class of projections is based on assumptions about the instrumental variables, for example that there will be an increase in the use of contraception or in its effectiveness. This class of projections may be undertaken to see the consequences of a predicted change, or the consequences of a change that may be implemented if the outcome appears desirable (essentially as a simulation). Since assumptions about instrumental variables are translated into reproductive consequences in the form of parity progressions and interval distributions, progressive fertility analysis is once again the preferred approach.

6.5 CONCLUDING CONSIDERATIONS

We have advocated the adoption of a new style of measurement, progressive fertility analysis. Since there are costs associated with any change from an accepted mode of analysis to something new, the case against change deserves careful consideration.

1 The detail of information provided may go well beyond what can reasonably be expected in terms of either statistical or substantive reliability. With respect to statistical reliability, we have proposed a temporal grid broad enough that adequate subsample size is available for a large part of the interesting reproductive terrain. As for substantive reliability, our position would be that, in investigating the hypothesis that some dimensions of the data configuration reflect misstatement, the same care should be expended in handling the data, with respect to remediable error, as would be used with impeccable data. One is surely not in a better position to argue a case if one ignores some available and relevant evidence than if one at least attempts to take that evidence into account.

2 There is a large chance component in the length of a birth interval, which may make it fruitless for us to press our measurement demands so far. Yet the proposed temporal grid is merely quinquennial, in part in response to concern about chance. And there is little to be lost by proceeding with the specified tabulations and then collapsing one or another dimension if the evidence betrays an intolerable random component.

3 Partly because of concern about quality of data, there is a body of opinion that one should keep the analysis of data simple for the developing countries, and leave the refinements to those working with data for developed countries. There are several responses to this somewhat condescending attitude. First, some of the proposed refinements have not been considered in any survey. Secondly, the challenge to measurement is greatest in those countries in which change in reproductive behaviour is only just beginning. The use of biased measures tends to conceal or at least blur small temporal variations; the same is true of measures of lower levels of specificity. Thirdly, if one agrees that the course of fertility is of major importance for the future well-being of many developing countries, then its measurement deserves more rather than less concern in terms of quality of effort.

4 The one unarguable disadvantage of any new system of measurement is that it differs from the traditional system. With a mass of information available in the traditional form, and an entire profession trained in traditional methods, resistance to change is only to be expected. Indeed it has been very difficult to produce the present work, because of the strength of ingrained habits and ways of thinking.

Nevertheless, the fertility survey is a new source of data, with new opportunities (as well as difficulties) not found in the traditional data sources. When it is used as a kind of substitute for a registration system, and forced to yield registration-style measures, the result is wasteful, incomplete, biased and distorted. Progressive fertility analysis has been developed with an eye to simplicity, in two senses: first, with respect to the correspondence between the form of data provided by the survey and the procedures for handling those data, and secondly, with respect to the correspondence between the form of parameter devised and the ways in which reproductive life in fact evolves.

Over the next few years it is our intention to conduct investigations in various populations, using progressive fertility analysis. Only through such work can the proposition be tested that the new measurement system increases our understanding of fertility.