WORKING PAPER

Are deliberate birth spacing effects in fact statistical and bio-demographic artifacts?

A critical study with data from the Netherlands

Jan Van Bavel & Jan Kok

Draft version

Paper to be presented at the XXVth IUSSP International Population Conference in Tours, France, 18-23 July 2005. Session 161: *Fertility control in history*, sat. July 23, 8:30-10:00 AM Salle Turgot.

Abstract

Our paper will take a critical stance towards the recently revived thesis that fertility might have been controlled before the fertility transition in Europe. To this end, we will use as guidelines critiques from both a methodological-statistical as well as a bio-demographic point of view. First, we review the kind of regression models that have been used in order to detect non-parity specific fertility control with individual level data. Particular attention will go to unobserved sources of heterogeneity and possible statistical solutions. Secondly, we discuss how birth spacing may be influenced by malnutrition and what the consequences for our statistical models may be. We apply our models to individual-level data about pre-transition populations in the western-central part of the Netherlands, including both rural municipalities and two major cities (Rotterdam and Utrecht). The fertility histories for first-married women we use are uncensored in that we have followed them, as they migrate, through different localities.

The dominant interpretation of the decline of fertility in Europe since the nineteenth century is that it was a transition from natural to controlled fertility. A fundamental historical-sociological implication of this view is a lack of agency with respect to reproduction in past times. Yet, there is ample literary and other qualitative historical evidence that aspirations to control fertility can be found in all times and social strata. Furthermore, the means employed to effectively limit fertility during the fertility transition (withdrawal, abstention, abortion) were not new at all. Therefore, the natural fertility view of pre-transitional reproduction has always been contested. However, it appeared to be very difficult to find statistical proofs of widespread fertility control by married couples before the secular decline started. Most evidence of fertility limitation could possibly just as well be explained by differential fecundity or some other component of natural fertility.

Yet, since the 1990s, with the advent of longitudinal and individual-level datasets and the application of event history techniques in historical analysis, more advanced multivariate techniques have allowed to control for a number of disturbing factors, allowing to some extent to test alternative explanations for observed fertility patterns. In earlier papers (Van Bavel 2004; Van Bavel and Kok 2004) we have applied simple hazard rate models in order to assess in new ways the extent of deliberate birth spacing in Belgian and Dutch populations that supposedly, according to conventional assessment techniques, exhibited natural fertility. In these papers, we have rejected the hypothesis of natural fertility: birth intervals seem to have been deliberately manipulated by men and women in response to household-demographic and –economic conditions (see also Van Bavel 2003).

In this new paper, we take a critical stance towards the recently revived thesis that fertility might have been controlled before the fertility transition in general, and towards our own previous findings in particular. To this end, we use as guidelines critiques from both a methodological-statistical as well as a bio-demographic point of view. In this paper we analyze individual-level data about pre-transition populations in the western-central part of the Netherlands, including both rural municipalities and two major cities (Rotterdam and Utrecht). The fertility histories for first-married women we use are uncensored in that we have followed them, as they migrate, through different localities.

First, we review the kind of simple hazard rate models that have been used in order to detect fertility control with individual level data. In particular, we focus on the effect of biological and behavioral heterogeneity within the population. On the one hand, we argue that our way of including proxies for differential fecundability in the regression model is doing a better job than earlier attempts in capturing an important source of heterogeneity. On the other hand, we are aware that our model has not been dealing well with (other) unobserved sources of heterogeneity (like differential breastfeeding practices; family size preferences; contraceptive skill), nor with the fact that our data have a multilevel structure (were birth intervals are nested within families). Therefore, we build and fit a new, mixed model, including random coefficients. Second, we discuss how birth spacing may be influenced by malnutrition and what the consequences for our conclusions may be. In order to construct a new empirical test, we include lagged time series data about food prices in our birth spacing model, inspired by the method proposed by Bengtsson and Dribe (2004).

1. Controlling fertility by spacing births

The classic definition of fertility control by Louis Henry refers to one specific form of deliberate fertility limitation, i.e. parity-dependent stopping behavior. Fertility control is said to exist when couples modify their behavior when the number of children reaches the maximum they do not wish to exceed (Henry 1961). Parity-dependence is implied by the reference to a maximum number of children. If a couple does not want to exceed a maximum, it means that they want to stop at the latest when the desired maximal parity is reached.

Yet, as pointed out by many authors – and we bet that Louis Henry would agree with them – other forms of fertility control are thinkable and even plausible in theory and in practice. Couples may prefer that births do not come too close to each other for health and household budget reasons. This preference for birth spacing may depend on the number of children already born and/or alive or not; i.e. birth spacing may be parity-dependent or -independent. An example of parity-dependent birth spacing would be when a couple would delay a third child (but not a first or a second) in order to devote more time and attention to the children already born, even if they do want to have more children eventually. An example of parity-independent spacing would be when a couple would try to delay any next birth. Theoretically, it makes maybe more sense to talk about forms of birth spacing that are parity-*aimed* or not, i.e. that are aimed at reaching some parity or not. Indeed, fertility control that does not have a specific target family size in mind may still *de facto* be depending on current parity (statistically).

2. Modeling closed birth intervals

In earlier articles, we have tried to detect two forms of deliberate birth spacing: one supposedly parity-aimed and one supposedly not. We did this by estimating the effect of net parity (i.e.

number of children currently alive) and the effect of the proportion among children alive who are still young and economically dependent on the speed of parity progression, i.e. birth spacing. Unusual in historical demography is that we applied our regression model only to closed birth intervals by conditioning on the *a posteriori* outcome that a next birth did indeed take place. We will now discuss this earlier birth spacing model in detail, defending what we think was right about it and revising what could be better.

2.1 Revision of our previous birth spacing model

Closed birth intervals

Like any hazard rate, a fertility rate has two components: on the one hand a likelihood that a birth eventually occurs anyway and on the other how long it takes before that birth occurs, if it does. With respect to mortality, the likelihood of the event eventually occurring is equal to one. With respect to fertility, however, this probability is lower than one because some couples are sterile. Therefore, it makes sense to distinguish between the probability and the speed of parity progression.

Unfortunately, we don't know who is currently sterile and who is still able to give birth. Therefore, the common approach in calculating fertility rates is to assume that everyone remains in the risk set until the wife reaches age 45 or 50. Of course, this causes a serious overestimation of the population at risk of giving birth: at age 40, for example, a substantial proportion of married couples are already sterile (about a third according to Wood 1994).

This is not really a problem when we want to look at unconditional fertility rates on the population level, because after all we are interested in the actual outcome of both the likelihood and the speed of giving birth. When we are particularly interested in the speed of parity progression, like in an analysis of birth spacing, this is a serious problem. Ideally, we would want to exclude all sterile couples from the analysis, as well as the couples who stopped their fertility careers deliberately.

Like Yamaguchi and Ferguson (1995) and Clegg (2001), the solution we have opted for in our birth spacing model is to look at conditional birth intervals, i.e. birth intervals that did end with a new birth, proving the fecundity of the couple as well as the absence of successful stopping.

We are convinced that this is the optimal solution in the sense of the best possibility. Yet, it is not an ideal solution, because, again, we don't know who stopped deliberately and who was in fact just postponing the next birth but ended up as a *de facto* stopper due to the onset of sterility before the birth of the next birth. So we surely lose some potential birth spacers by limiting our analysis to closed birth intervals, but the bias this potentially causes is much less severe, we think, than when we would include the high number of person-years for people who are in fact not at risk of giving birth anymore.

Time-dependence of the hazard function

In our previous birth spacing model, we applied Cox regression to assess the impact of a number of natural and social covariates on parity progression. A major advantage of Cox' partial likelihood approach is that there is no need to specify how the hazard depends on time, because the baseline hazard is not estimated and is allowed to go up and down with time freely. In this paper we have abandoned this semi-parametric approach for three reasons.

1. First, partial likelihood estimates do not allow to calculate in a straightforward way predicted failure times, although there are a number of pseudo-solutions suggested in the literature. As a result, only rather crude measures of goodness-of-fit, based on the partial likelihood are available. It is not easy to assess goodness of fit in an intuitive way, for example by comparing the model predicted frequency distribution of birth intervals with the observed frequency distribution.
2. Second, we know quite a lot in advance about how the hazard function depends on time when we limit ourselves to closed birth intervals. We know that it will probably first increase, then reach a plateau and finally increase ever more rapidly towards one, because we have arranged things in advance so that everyone in the risk set eventually gives birth again. As a result of this a priori knowledge about the time-dependence of the hazard function, we may apply more powerful parametric methods.
3. Third, we would like to try out a hazard regression model with random slopes as well as random intercepts. Whereas estimation of Cox regression parameters with random intercepts is technically possible nowadays, it is very difficult, if at all possible, to get good estimates for a random slopes model. Multilevel models with random components are important to account for unobserved heterogeneity among families (for an early application of this kind of models in historical demography, see David and Mroz 1989a; 1989b; Mroz and Weir 1990; Mroz and Weir 2003).

For these three reasons, we will adopt a parametric approach in discrete time in this paper. More specifically, we switch over to discrete time and use logistic regression to model the monthly probability of conception. A logistic regression model with random slopes as well as random intercepts can be estimated with up-to-date standard statistical software, like the SAS-procedure

NLMIXED. An additional advantage is that it is easier to include time-varying covariates in the regression.

A life table analysis reveals that the discrete-time hazard depends on time as expected, i.e. it first increases, then reaches a plateau and then rapidly soars to reach one (see figure 1).

Figure 1. Estimates of the discrete time, conditional hazard rate of having a next birth (calculated per five months)

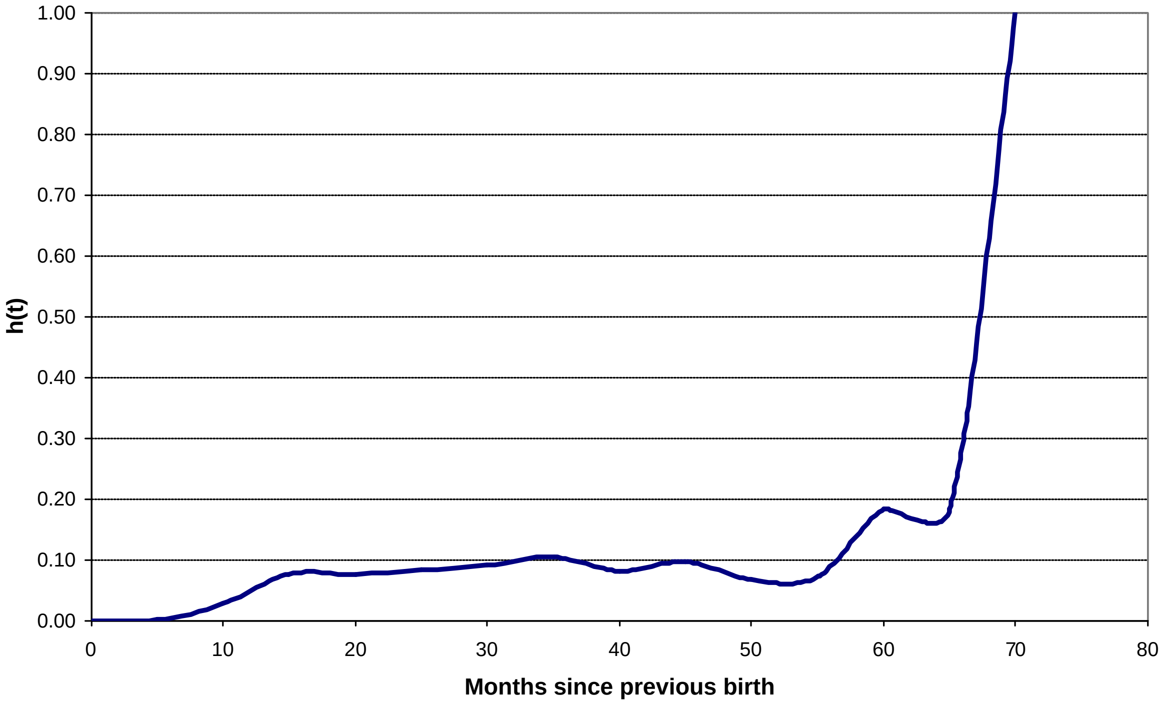

Taking the logit transformation of these discrete hazards rates, which are in fact conditional probabilities, yield the following picture (figure 2):

Figure 2. Logit of the life table estimates of the conditional hazard function, and third order polynomial regression on time

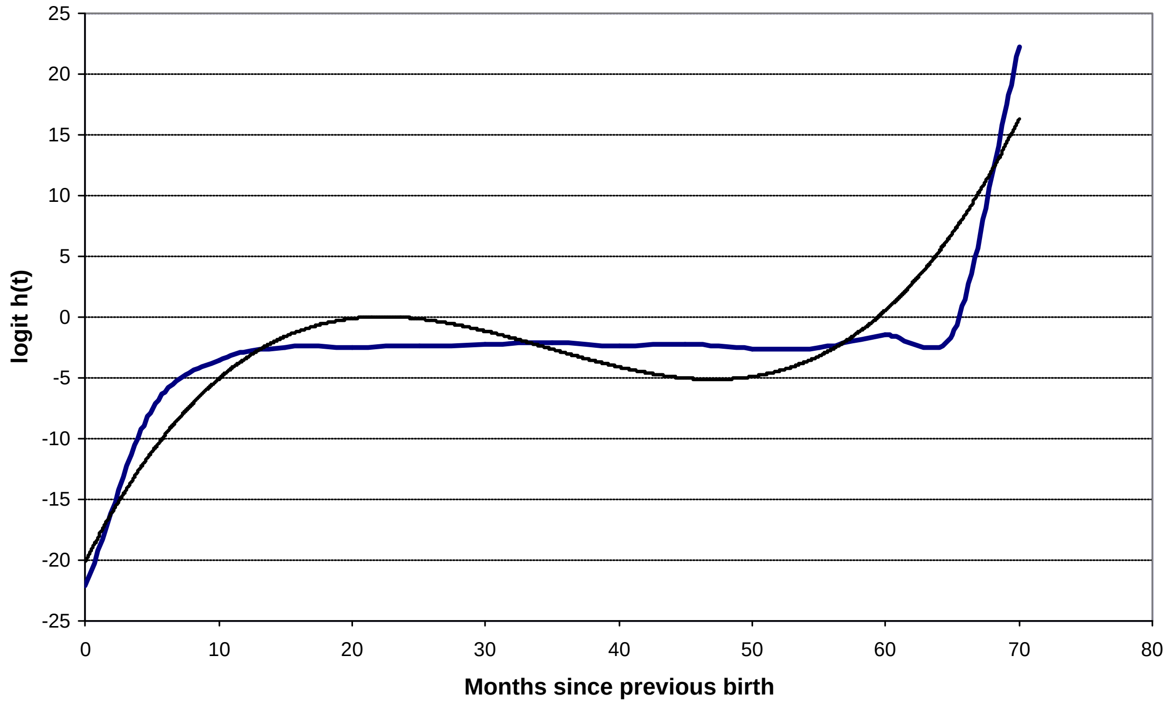

As can be seen in the figure 2, the time-function of the logit can quite well be approximated by a third order polynomial. Therefore, we will include time, time² as well as time³ in our logistic regression model to describe time-dependence in a parsimonious way.

Note that there are other time-dependent covariates in our model, like age and marriage duration. Later on we will show that the polynomial time-terms combined with age and marriage duration yields a remarkably good fit between non-parametric estimates of the hazard function and the results from our logistic regression.

A number of covariates have been included in our birth pacing model because they are important proxies for natural fertility determinants.

Age and marriage duration

There is quite general consensus that the age pattern of natural fertility is basically a function of the onset of sterility (Wilson, Oeppen and Pardoe 1988). Indeed, in our earlier analysis we found that the age of the wife did not significantly affect the speed of parity progression among non-sterile couples, i.e. among couples who went on to have another birth. In our earlier analyses, we included four dummies corresponding to different age-of-the-wife-groups. In our current model, we simplify by just including age in years as a covariate. Tests showed that including a second order polynomial or dummies does not significantly improve the fit of our models.

In our view, marriage duration, i.e. number of years elapsed since the start of marriage, is the most important observed determinant of the speed of parity progression, probably due to its link to coital frequency. Preliminary analysis revealed that categorizing marriage duration or including polynomial terms, in order to allow for a non-linear relationship between marriage duration and the logit of the hazard, does not significantly improve the fit of the conditional hazard model. In fact, including just marriage duration in years in a logistic hazard regression model yields the following functional form of the relationship between marriage duration and the monthly probability of giving birth (see figure 3):

Figure 3. Functional form of the modeled relationship between marriage duration and the discrete time, conditional hazard rate of giving birth

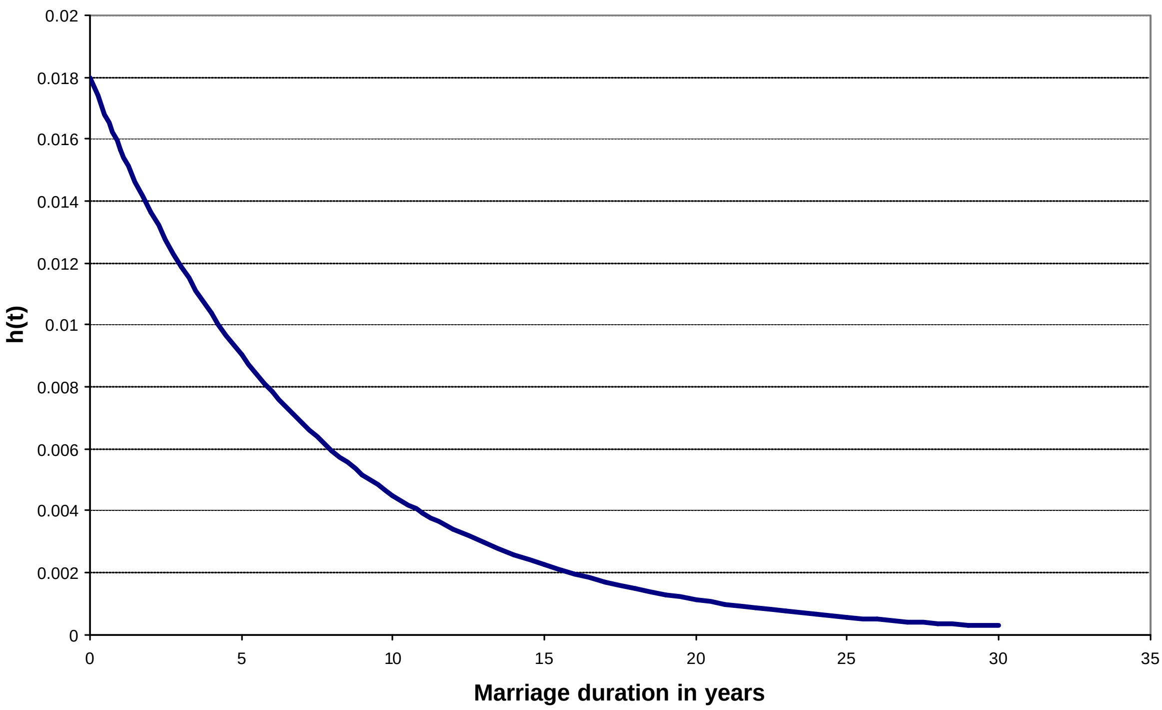

This looks like a good approximation of "the saddest curve" in the social sciences (to quote Westoff):

Figure 4.

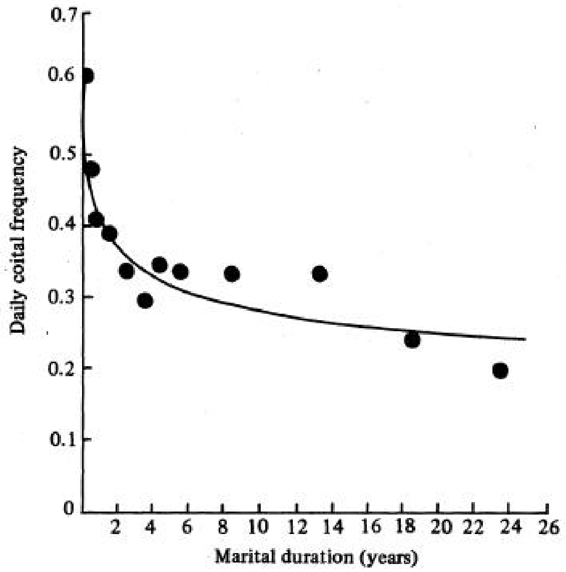

Source: Wood and Weinstein (1988)

With age of the wife, duration of marriage and time since the previous birth, we have three time variables in the model which are running jointly:

$$\log\left(\frac{h(t)}{1-h(t)}\right) = b_1 M + b_2 M^2 + b_3 M^3 + b_4 D + b_5 A + \ldots \qquad (1)$$

The difference between these covariates is that they have different starting times and they are measured on a different time scale: age A is elapsed time since the birth of the wife in years, duration D is time since the start of the current marriage in years, and M is the number of months elapsed since the previous birth. The latter time covariate is included as a third-order polynomial.

Maximum likelihood estimation of the beta's for these covariates in a logistic regression and calculating the predicted discrete time hazard rates from this as a function of months since the previous birth, yields the following result (see figure 5):

Figure 5. Model predicted discrete time conditional hazard rate of next conception as a function of elapsed time since the previous birth

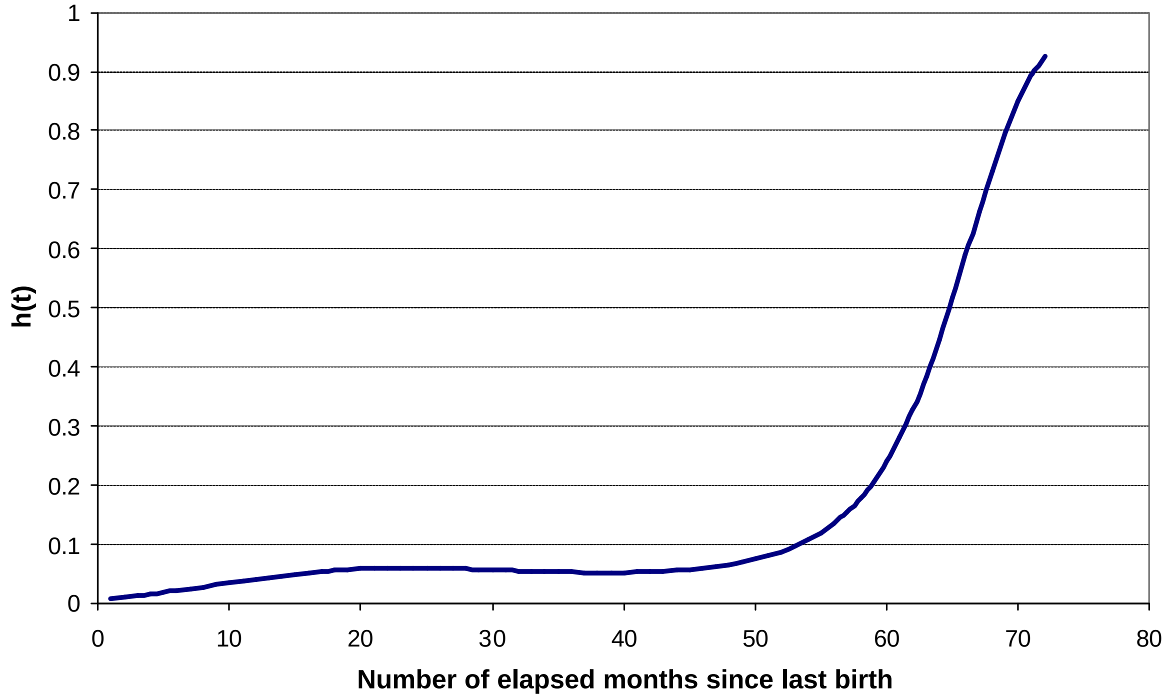

The overlap with the nonparametric estimates of the discrete time hazard rate function by time is striking (see figure 6). We conclude that we have found a realistic parametric specification of how the hazard depends on three dimensions of time.

Figure 6.

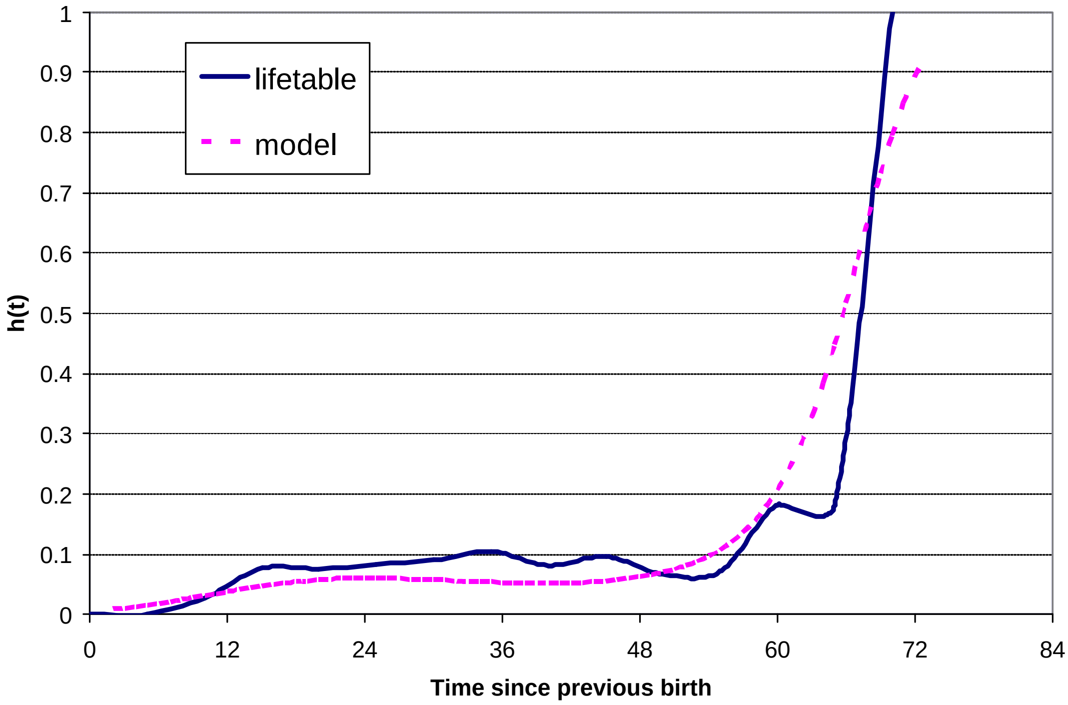

CEB as an indirect indicator of fecundability and postpartum amenorrhea

An important but unobserved source of heterogeneity with respect to the speed of parity progression follows from differences in fecundability and postpartum amenorrhea. To some extent we can control for differences in coital frequency, which is a major determinant of fecundability, by taking marriage duration into account. Yet, coital frequency is not the only determinant of the monthly probability of conception: physiological differences also play an important role. And even if coital frequency would be the only determinant, marriage duration is still just a crude proxy. Also the length of the postpartum infertile period, to a large extent a function of breastfeeding, is not observed directly.

Yet what we do observe is the result of the joint influence of fecundability and postpartum amenorrhea, i.e. the number of births that have already occurred in the past, given the current marriage duration. Therefore, after controlling for marriage duration, we consider the number of children ever born (CEB) within that marriage as a good indicator of the level of fecundability and postpartum amenorrhea combined. The higher the fecundability characteristic of a given marriage, the higher the number of legitimate births at a given duration of marriage. The longer the postpartum infertile period lasts, the lower the number of children already born within that marriage.

If this argument is correct, we would expect that marriages with a higher number of children already born will have their next birth quicker, given that they do go on to have a next birth, than marriages with a lower CEB at a given duration. Kaplan-Meier estimates of the survival distribution function, applied to the Dutch data we have used in previous articles, prove that this is indeed the case. The pattern displayed in figure 7 for couples who have been married for four to six years holds for other marriage durations just as well: the more children already born, the faster a next birth occurs. We take this as evidence that CEB within marriage, controlling for the duration of that marriage, is a good indicator for the combined influence of fecundability and postpartum amenorrhea.

Marriage duration [4,6[

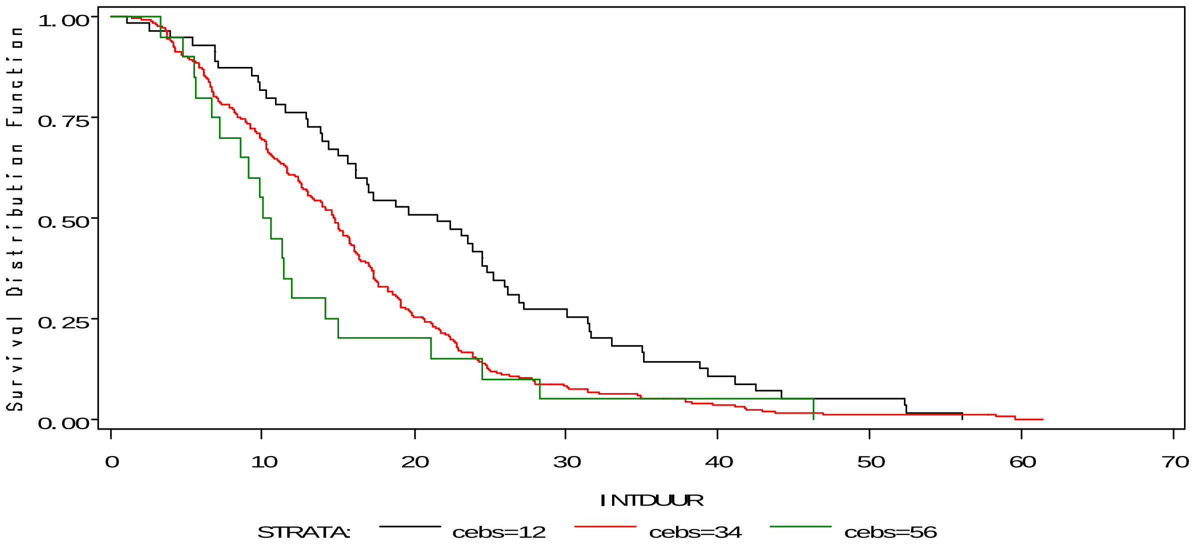

The detection of fertility control

After taking into account the proxies for natural determinants of the speed of parity progression, including the survival of the lastborn child, the source of heterogeneity at issue is fertility control. In our earlier contributions, we argued to have indicators for two kinds of controlled birth spacing (Van Bavel 2003; 2004a; Van Bavel and Kok 2004): one aimed at keeping the household economic pressure under control by keeping down the proportion of young, dependent children,

and one aimed at the number of surviving children as such. Therefore, we included the proportion of children under age nine as well as the number of children currently alive in our regression models. There is, however, a problem with the interpretation of the former covariate.[1]
To make clear what the problem is, we first assume that there are no children born before marriage. Under these circumstances, the proportion of children under age nine will be equal to one during more or less the first nine or ten years of marriage, as long as at least one child is born and still alive. Hence, the meaning of this covariate during about the first ten years of marriage is that it indicates whether or not the marriage has produced at least one child that is still alive, irrespective of whether just one child is alive or four or six. When the oldest child reaches age nine, the meaning of the covariate changes completely. Only at higher durations of marriage, it makes sense to interpret it as a genuine proportion among the children alive who are still under age nine.

This suggests that if it is really the proportion that matters and not just the question whether or not at least one child is alive, that this covariate should exert its strongest effect after nine or ten years of marriage. Yet, tests showed that there was no effect at all of the proportion of young children after ten years of marriage. Instead, the effect appeared to be limited to the first ten years of marriage. This makes sense in terms of birth spacing, but not in terms of the interpretation we originally gave to it. It makes sense in terms of birth spacing because deliberate birth spacing, if present, is most likely to occur during the first years of marriage, when there is still room for flexibility and adaptive behavior with respect to speeding up or slowing down parity progression, depending on the household-economic circumstances. At higher marriage durations, associated with higher biological ages and lower fecundability, there is less room for such adaptive behavior. But our original interpretation does not make sense because the covariate we used just jumps up and down from zero to one, depending on whether or not at least one child is alive.

If children are born before marriage and if they are included in the calculation of the proportion of children under age nine, like in Van Bavel (2003 and 2004a), then the dependent children covariate may or may not include children over age nine. In the beginning of marriage, it may just indicate whether or not there are children alive that were already born before marriage. Anyway, its meaning is clearly ambiguous.

In sum, taking into account that the proportion of dependent children has no effect at higher marriage durations (>9 years), and given that the "proportion"-covariate in fact indicates whether there are any children or not during the first years of marriage (<10 years), we drop this ambiguous variable from the regression equation. Instead, taking into account that survival of at

[1] We thank George Alter and Martin Dribe for their useful criticism of our previous model.

least one child has a strong negative effect on the speed of parity progression during the first decade of marriage, we introduce an interaction effect between marriage duration and number of children currently alive to see whether birth spacing was parity-dependent during the first part of marriage. Recall that we also have a dummy in our regression that indicates whether or not the previous child is still alive, so the net parity covariate is capturing the survival of children that were born before the previous one. The results of the new specification are presented and discussed in the findings section.

Unobserved heterogeneity

We have argued that marriage duration and the number of children already born are good proxies for important sources of unobserved heterogeneity among marriages, i.e. fecundability and postpartum amenorrhea. Yet, the observed covariates in our model are not direct measures of these proximate determinants of birth intervals. Therefore, we want to try out what happens when unobserved heterogeneity is taken into account. First, we see whether the intercept of our logistic regression model has significant variance, reflecting differences between marriages in the overall speed of parity progression. Hence, our regression model looks like this:

$$\log\left(\frac{h_{im}(t)}{1-h_{im}(t)}\right) = (a + u_m) + bX_{im}' + gZ_m' + e_{im} \qquad (2)$$

where a is the overall intercept and u_m the family-specific deviation from the overall intercept. In order to make estimation possible, we assume that u_m has a normal distribution with mean 0; to be estimated is the variance of u_m. X_{im} is a vector of interval *and* marriage specific covariates (like CEB or marriage duration) and Z_m is a vector with characteristics typical of the marriage but not the interval (like occupation of the husband, for example). For ease of notation, the time-index for time-varying covariates is omitted.

Second, we want to test whether there is also heterogeneity with respect to the effect of our indicator covariate for parity-dependent fertility control. Maybe some marriages were adapting their behavior as a function of the number of children alive and some were not displaying this behavior. To test this, we also estimate a variance component for the slope of number of children alive. The regression model with random intercepts and random slopes looks like:

$$\log\left(\frac{h_{im}(t)}{1-h_{im}(t)}\right) = (a + u_m) + bX_{im}' + (f + v_m)N_{im} + gZ_m' + e_{im} \qquad (3)$$

where, in addition to the parameters in model (3), N_{im} is the observed number of children currently alive, F is the overall slope for that covariate and v_m the family-specific deviation from the average effect. Again, v_m is assumed to have a normal distribution with mean 0 to make estimation of the model possible. The variance components we estimate are var(u_m), var(v_m) as well as the covariance between u_m and v_m. (Note that the X-matrix also includes the product of marriage duration and number of children alive, in order to estimate the interaction effect discussed in the previous paragraph.)

2.2 Economic stress and postponing births

In a recent contribution about living standards and economic stress, Bengtsson (2004) has suggested that there exists a social ladder of economic and demographic responses to short-term economic stress (see p.35). The better-off overcome short-term economic stress by the spending of savings or by borrowing from a bank, for example. People without much movable or immovable property could maybe bridge hard times by some form of relief, possibly poor relief, if available and eligible. Families could send out some of their children to look for work elsewhere. Furthermore, married couples could try to control the consumption load in their families by postponing births. The ultimate response mentioned by Bengtsson (2004: 35-36) is, instead of postponing consumption by postponing births, reallocating consumption due to mortality of dependent family members (old or young).

Food, fecundity and fertility control

If we would find evidence of the postponement of births in reaction to economic stress, as measured by exceptionally high food prices, this may be evidence of deliberate birth spacing. The problem is that it is difficult to distinguish between purely physiological responses to malnutrition in years of economic crisis on the one hand, and deliberate behavior on the other. The issue of the possible effect of malnutrition on fertility is still a controversial one, although demographers tend to cite Gray (1983) and conclude from studies carried out in developing countries in the 1970s and 1980s that chronic malnutrition below starvation level has little effect on fertility; and that only in case of actual starvation fertility is greatly reduced, presumably because menstruation ceases (Campbell et al. 2002: 743-745).

Yet, claims to the contrary have to be taken seriously. Peter T. Ellison (2001; 2003) has recently argued that chronic malnutrition, even below the level of starvation, may influence fertility significantly through its inhibiting effect on the resumption of ovulation during breastfeeding.

The exact mechanism connecting breastfeeding with the infertile period is still a controversial issue but the nursing frequency hypothesis has been most widely accepted. In a nutshell, according to that hypothesis, more frequent nursing is associated with longer postpartum amenorrhea. Yet, Ellison (2001) argues that the evidence does not really support that hypothesis. He puts forward the thesis that the length of the postpartum amenorrhea depends highly on the relative metabolic load of breastfeeding: a given breastfeeding pattern will imply a shorter infertile period for mothers with a high energy budget than for mothers who have a low energy budget because the female ovarian function is highly sensitive to energy balance and energy flux. The energy budget depends both on food intake and on available fat reserves.

Without wanting to go in detail into this debate, it is clear that distinguishing between possibly physiological effects of malnutrition on fecundity and fertility control as a deliberate response to high food prices is no straightforward task. Yet, in a recent paper, Bengtsson and Dribe (2004) have proposed a method that may be a solution to this problem. In a nutshell, they argue that if fertility is lowered very soon after an economic downturn, say within six months, it would be difficult to explain this by physiological effects as a result of malnutrition. Instead, deliberate control as a result of families foreseeing bad times would be a more plausible conclusion.

We will check whether early responses to economic crisis are also present in our Dutch population. To this end, we will include de-trended information on past rye prices as well as current and future rye prices in our regression.

Food Prices

In anticipation of economic hard times, people may choose to delay the arrival of a next child, by lowering their coital frequency or applying other methods of birth control. To distinguish between physical subfecundity and deliberate spacing we need food prices that are dated as accurately as possible. For various markets in the Netherlands monthly prices have been recorded, but few have been printed. After 1855, when the excises on grain milling were lifted, the figures become even rarer. In addition, the available series show many gaps that cannot always be filled by interpolation. For our purposes, we have selected the most complete series: the monthly price of rye on the market in Groningen in the northeastern part of the Netherlands, covering the period between 1785 and 1913.[2] In addition, we use the price of Prussian rye on the Amsterdam stock exchange (Posthumus 1943) to impute missing values. The correlation between the Hodrick-Prescott de-trended Groninger and Amsterdam series was very high (Pearson's r 0.88, p<0.0001), indicative of the high level of market integration in The Netherlands. A regression analysis allows us to predict the value of the Groningen price for a small number of missing months on the basis of the Amsterdam series. However, this was not feasible for the period between September 1867 and July 1869, so we miss price information for this period (see figures 8 and 9). All prices are converted so that they are in Dutch Guilders per last (i.e. 30 hectoliter) of rye.

Figure 8.

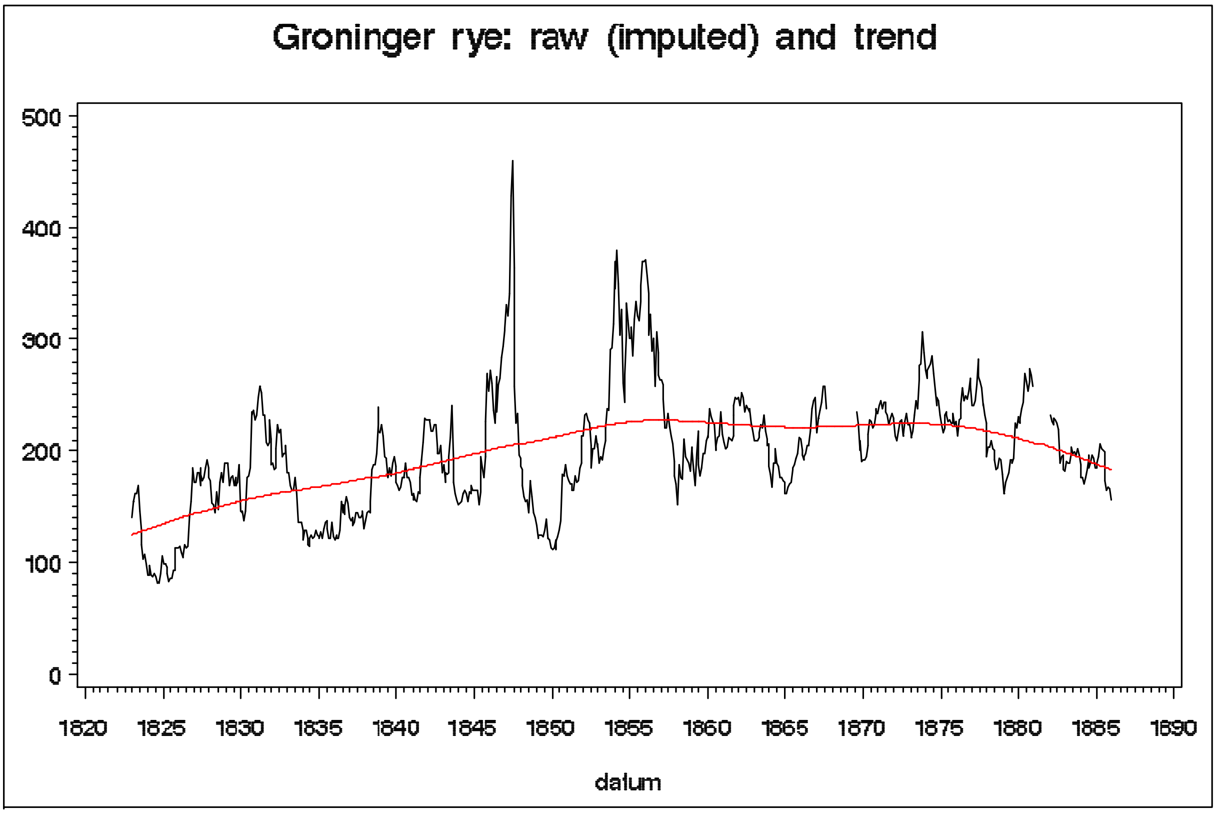

[2] The Groninger prices were downloaded from http://odur.let.rug.nl/~nahi (W. Tijms 1823-1853); and http://www.iisg.nl/hpw/data.html#netherlands (A. van Riel 1861-1913). We converted the different regional measures into a single denominator.

Figure 9.

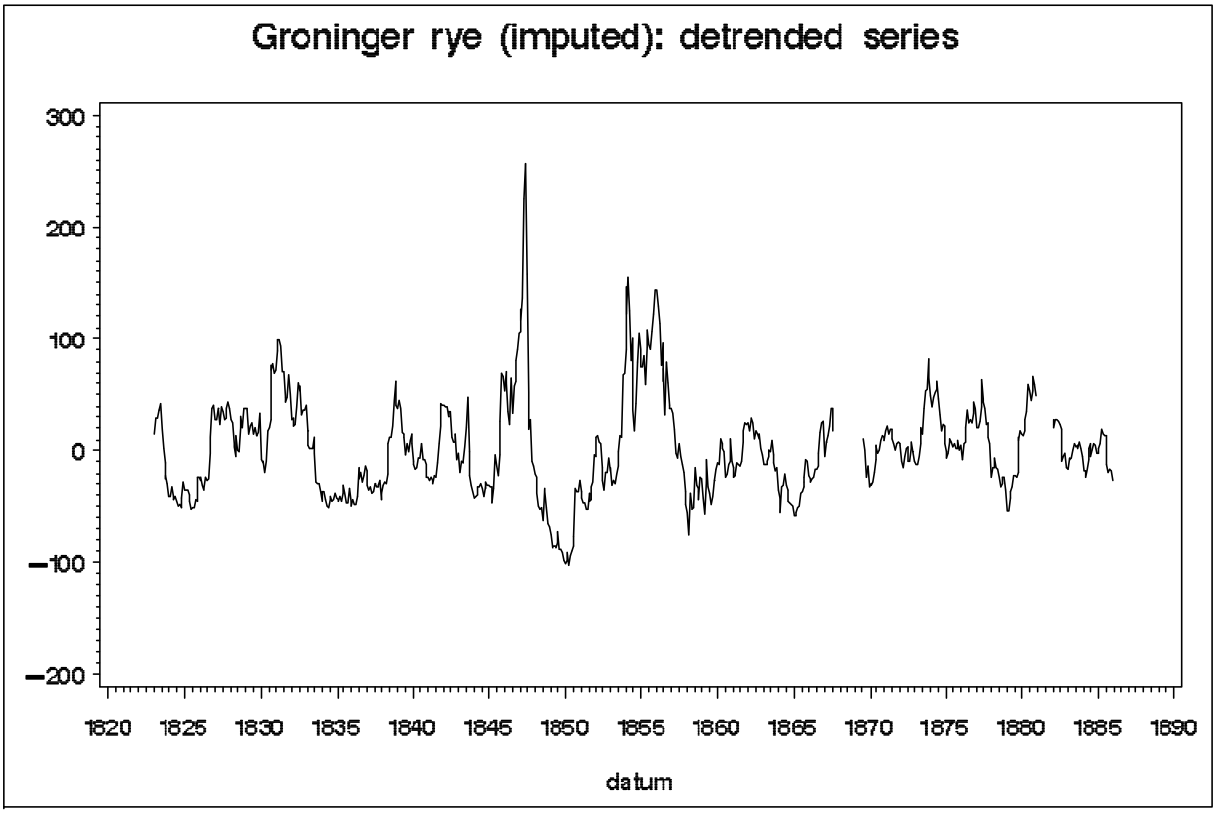

After de-trending (using the Hodrick-Prescott filter), we first convert monthly prices into quarterly series. We first experimented with quarterly de-trended prices in the regression model using several lags. These experiments were not successful: the estimated effects were pointing in inconsistent, mostly non-significant directions.

In our effort to find an explanation and a solution, we came up with the following argument: if we just include de-trended series, we assume that the effect of price fluctuation on birth pacing is symmetric and continuous. Symmetric: low prices are assumed to have exactly the opposite effect of high prices. And continuous: there are no thresholds in the price level to be crossed before any effect of prices can be expected. Experiments showed this to be not true: when we estimated the effect of *positive* prices only (prices above the trend-line), the results became consistent. When we estimated the effect of *high* prices only (more than one standard deviation above the trend-line), these became significant as well. In sum, in the regression models fitted below, “high prices” means prices that are at least one standard deviation above the time trend in the price series, i.e. more than one standard deviation above zero (which is the mean for the de-trended series).

Price effects in different socio-economic groups

To what extent were the various socio-economic groups in the central and northwestern part of the country susceptible to fluctuations in the price of basic foodstuffs? Following Bengtsson (2004), we can hypothesize the reactions of diverging social groups to increases in food prices. Since our dataset covers both rural and urban areas in a country already highly commercialized and specialized, there are many occupational titles in our dataset. We use only the occupation declared at marriage, because although there is some amount of occupational mobility, it is not possible to date changes. To categorize occupations, we employ a classification widely used in Dutch historical demography (Giele and Van Oenen 1974, 1976, see for instance Van Poppel, Jonker & Mandemakers 2005). Our first two groups, the elite and the officials were probably least affected by price fluctuations. The first group consists of people with well-paid public or clerical functions, such as mayors and church ministers. Members of the second group are mainly lower military, lower civil servants, school teachers etcetera. Their stable positions and relatively good salaries are assumed to have offered them some protection during economic crises.

Farmers in North-Holland and the western part of Utrecht were rather prosperous. The farms were small in size and workforce but highly oriented towards producing quality products for export. Around the cities, horticulture expanded whereas in the coastal parts of North-Holland bulb farming was of importance as well. Contrary to Bengtsson (2004), we do not expect farmers to benefit from an increase in grain prices, since they hardly produced grain themselves. In addition, some farmers were crofters, who often could not sustain their holding and declared themselves day-laborers within a few years after the wedding. In our North-Holland sample, about one-third of the farmers had very small holdings, no employees and was exempted from taxation (Damsma and Kok 2005).

The lower middle classes in our sample are mainly shopkeepers, traders and artisans. They must have reacted rather strongly to rising food prices, since the purchasing power of most of their clientele would drop sharply. In mid-nineteenth century, the economic situation of this group was very similar to the working-class. In fact, many shopkeepers were combining wage work with running a small shop at home (De Jonge 1968: 18; De Meere 1979: 254-254).

We divide the working class in two groups; on the one hand the skilled workers and those with fixed labor contracts, on the other hand the casual and unskilled laborers. Clearly, the latter and most numerous group was the most vulnerable for sudden rises in the food prices. For agricultural day-laborers it was vital to rent a small plot of land to grow some foodstuff or have some poultry or a pig (Kok, 2004). In budgets of working class families in the middle of the nineteenth century, primary food stuffs (bread and potatoes) took up about 30-40 per cent

(Horlings and Smits, 1996). The working classes provided the bulk of the conscripted soldiers. The development in stature of nineteen-year old conscripts reflects consumption patterns. During the second quarter of the nineteenth century the average height of the conscripts declined. The period 1855-1860 stood out with the largest shares of conscripts declared unfit for service because of insufficient height (De Meere, 1982). Around 1860 the downward trends in stature and real wages were reversed for good.

Although the large and growing percentages of rejected conscripts indicate chronic malnourishment among significant parts of the working class, there is little evidence of true subsistence crises before 1860. Detailed analysis of food prices and causes of death registration suggests that, in the silted western part of the country, endemic malaria was primarily responsible for high mortality rates and for lowering general resistance against infectious diseases. Malnourishment certainly contributed to the excessive mortality of 1846 and 1847, but was of secondary importance (De Meere 1982: 80-92). The reactions of the Dutch population to fluctuations in the food prices was probably mitigated by poor relief, which seems to have reached far larger proportions than the few percentages mentioned for Southern Sweden (Bengtsson 2004). In fact, around 1850 23% of the whole North-Holland population was on poor relief. In Utrecht province, this figure was 15% and in the country as a whole 14% (Van Leeuwen, 2003).

To sum up, we do expect reactions to food prices among the casual and unskilled workers and to a lesser extent among the skilled workers and lower middle class as well. However, the food supplements offered by poor relief and the knowledge that poor relief would be available will have diminished the effect of rising prices on fertility.

3. Data and context

In this article we analyze elaborate family reconstructions covering two provinces of the Netherlands in the period 1820-1900. These provinces, Utrecht and North-Holland, are located in the central and northwestern part of the country. The first reconstruction is drawn from the Historical Sample of the Netherlands, a large database that will contain more than 70.000 life courses. The database is built from a random sample (0.5%) in the Dutch birth certificates of 1812-1922, linking and entering all information in both the civil registers (birth, marriage and death certificates) and the continuous population registers which started in 1850 (Mandemakers, 2000). We use the first, more or less completed, part of this database that covers the province of Utrecht. We have restricted the analysis to first marriages in which the wife was born before 1861

(N=302). The second dataset is built from a marriage cohort in one North-Holland village (1830-1879). All first-marrying couples in Akersloot, an agricultural community about thirty kilometers to the northwest of Amsterdam, were traced in their migration trajectories. However, most of them remained in or near the village of Akersloot. This was an ordinary North-Holland village, except for the fact that its continuous population registers started already in 1830. For our analysis, we once again select women born before 1861 (N=280). Finally, we make use of a small family reconstruction from the city of Amsterdam (1820-1850). This database, originally intended as a three-generation study of poor relief recipients, consists of 84 families.[3]

During the 17th and early 18th centuries, North-Holland and Utrecht had benefited greatly from the commercial successes of the Dutch Republic. However, the second half of the 18th century was characterized by economic decline. The Napoleonic Wars sealed the fate of the Netherlands as a leading seafaring nation. In the first half of the 19th century, the coastal province of North-Holland was only slowly recovering from this crisis. The size of the population in cities like Amsterdam stagnated until 1850 (Kok, 2003). In the second half of the 19th century, the provinces of North- and South-Holland regained their former dominance, due to a redirection of the trade streams towards Germany. Amsterdam became a world centre of financial services. Also, the area to the north of Amsterdam industrialized rapidly. The small, inland province of Utrecht had suffered less from the collapse of the sea trade. Its population grew strongly, in particular in the eastern part of the province. From around 1850 onwards, the number of factories increased with chemical, textile and cigar making industries predominating. The province, especially its major cities Utrecht and Amersfoort, profited strongly from is central location at the nexus of Dutch railway lines. This attracted railroad offices and workshops, metallurgical industries and a host of commercial service companies (Knippenberg, 1995).

4. Results

First, we will present the results from our basic discrete time hazard model, without random components or price information. Second, we introduce random intercepts and random slopes. Third, we try to assess the influence of high prices on the speed of parity progression.

[3] We only use the second generation here. We thank Marco van Leeuwen for his kind permission to use the data files.

4.1 Results without price information

The estimates for the regressions without price information are in table 1. The left panel of that table displays the results without family-level variance components.

Table 1. Logistic discrete time hazard model of effective conception in closed birth intervals, western-Netherlands, 1825-1885

	Without random components			*With family-level random components*		
	b	*s.e.*	*p*	*b*	*s.e.*	*p*
Intercept	-4.0418	0.2080	<.0001	-4.2123	0.3143	<.0001
Age of mother in years	-0.0114	0.0067	0.0912	-0.0148	0.0112	0.1866
Marriage duration in years	-0.1416	0.0158	<.0001	-0.0322	0.0275	0.2429
Months since previous birth	0.2521	0.0146	<.0001	0.2810	0.0159	<.0001
Month²	-0.0082	0.0007	<.0001	-0.0087	0.0007	<.0001
Month³	0.0001	0.0000	<.0001	0.0001	0.0000	<.0001
CEB within marriage	0.2176	0.0261	<.0001	-0.0176	0.0493	0.7210
Death of previous child	0.5779	0.0737	<.0001	0.6319	0.0852	<.0001
Occupation of father						
Unskilled worker (ref)	1.0000	/	/	1.0000	/	/
Elite	0.0233	0.2565	0.9278	0.0246	0.4430	0.9557
Farmer	0.2440	0.0622	<.0001	0.2470	0.1070	0.0214
Official/white collar	0.3643	0.1312	0.0055	0.4721	0.2311	0.0415
Shopkeeper/artisan	0.1758	0.0821	0.0323	0.2498	0.1445	0.0845
Skilled worker	0.0478	0.0639	0.4546	0.1060	0.1067	0.3207
Religion of parents						
Both liberal Protestant (ref)	1.0000	/	/	1.0000	/	/
Both Catholic	0.1942	0.0532	0.0003	0.2820	0.0936	0.0027
Orthodox Protestant	0.1952	0.0923	0.0345	0.4199	0.1507	0.0055
Mixed Catholic-Protestant	0.1101	0.1309	0.4001	0.1632	0.2239	0.4665
Both Jewish	-0.4748	0.3083	0.1235	-0.6314	0.5469	0.2488
Both unknown	-0.1593	0.1068	0.1359	-0.0836	0.1792	0.6409
Other	0.5875	0.2223	0.0082	0.6588	0.3897	0.0916
Number of children alive	-0.1249	0.0378	0.0009	-0.1821	0.0493	0.0002
x marriage duration	0.0084	0.0025	0.0007	0.0095	0.0030	0.0016
Marriage-level variance components						
(1) Variance of intercept				0.5544	0.1186	<.0001
(2) Variance of slope of number of children alive				0.0071	0.0084	0.3993
(3) Covariance (1)x(2)				-0.0282	0.0279	0.3123
Number of marriages			516			516
Number of births			2206			2206
-2 log likelihood			15800			15705
AIC			15846			15753

First, the small influence of age of the mother, after controlling for marriage duration, on the conditional speed of parity progression is confirmed, reflecting the fact that we are restricting the analysis to non-sterile couples. The importance of marriage duration, presumably due to its relationship to coital frequency, stands out again. The same holds for the number of children already born within the current marriage: as found earlier, people who have already had many children have their next child sooner, if they have one, than people with fewer children already born. This reflects differences in fecundability and postpartum amenorrhea. The results from our earlier analysis of occupational and denominational differences are confirmed as well (Van Bavel and Kok 2004): farmers, shopkeepers, artisans, officials and other white collar workers have their next child, if any, quicker than unskilled workers. Catholics and orthodox Protestants had their next child sooner than liberal Protestants.

These are interesting and important differences, but at issue here is the effect of the number of children alive. After dropping the proportion of dependent children from the analysis, for reasons discussed above, and after including an interaction term with marriage duration, the speed of parity progression appears to depend significantly on the number of children alive. We take this as strong evidence of deliberate, parity-dependent birth spacing.

The interaction of number of children alive with marriage duration is significant as well: the longer the marriage has already lasted, the smaller the negative effect of the number of children alive on the speed of parity progression. The functional form of the hazard rate by number of children alive and marriage duration implied by these estimates is illustrated in figure 10.

The right panel of table 1 gives the estimates from the multilevel model that includes family-level variance components. In general, the conclusions to be drawn from this regression are the same as for the model without these components: the estimated slopes for occupational groups have the same direction and order of magnitude. The estimated effects of denominations as well as of the number of children alive are even somewhat stronger in the multilevel version. The most interesting, additional finding is that, after taking into account unobserved heterogeneity on the marriage level, the effects of marriage duration as well as CEB almost disappear and become non-significant. This strengthens our argument that these two variables are indeed picking up important sources of heterogeneity between families, i.e. differential fecundability and postpartum amenorrhea. Apparently, the family-level variance component is doing a still better job in taking these sources of heterogeneity into account. The variance of the intercept (allowing every family to have its own overall pace of parity progression) is highly significant. The variance of the slope for net parity, however, is very small and not significant. Our interpretation of this is that, after taking into account characteristics like denomination and occupational group, parity-dependent birth spacing was a quite general phenomenon in our population.

Figure 10. Functional form of the discrete time hazard by marriage duration and number of children alive

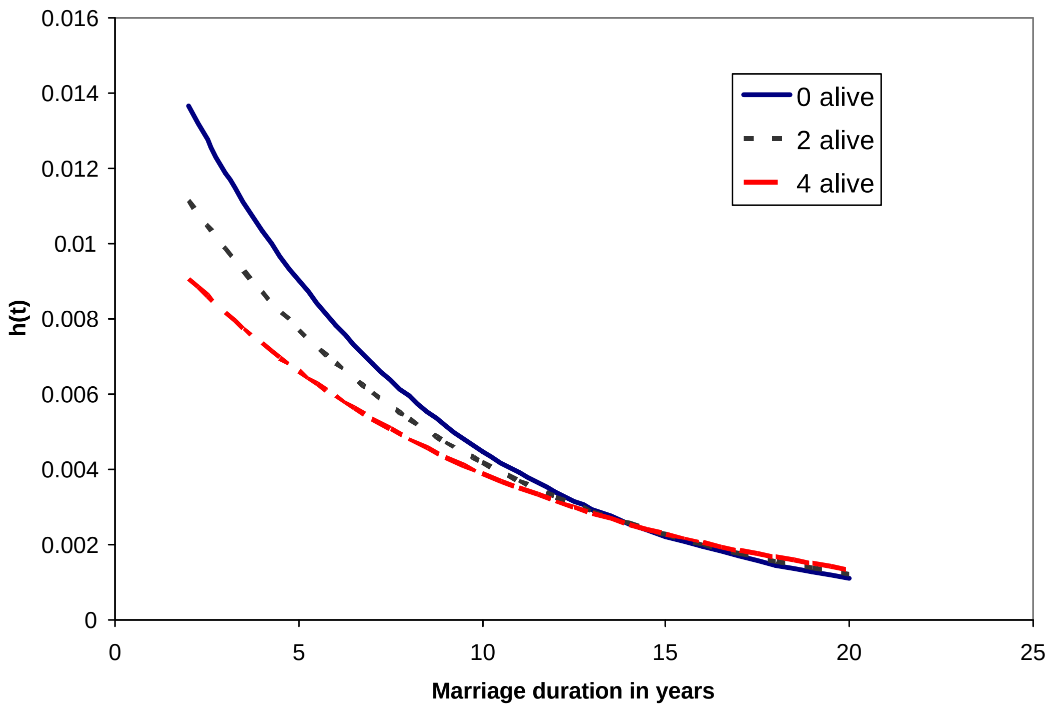

4.2 The influence of high prices

As explained above, our attempts to find price effects were not successful as long as we assumed these effects to be symmetrical (low prices having the opposite effect as high prices) and continuous (no thresholds to be crossed before effects can be expected): the estimates seemed to point quite erratically in changing directions and where not significant. Things improved a lot when we specified price effects to be discontinuous and asymmetric. More specifically, the "high price" covariates we eventually used in our regression are defined as follows: when a de-trended price level was below the threshold of one standard deviation, we set the high price variable to zero. For prices above that level, we subtracted one standard deviation and defined the high prices level to be that difference. So the estimated slopes of high prices have to be interpreted as the effect of one extra Dutch Guilder above the threshold of one standard deviation.

As explained above, following Bengtsson (2004) and Bengtsson and Dribe (2004), we expect that price effects differ between socio-economic groups. Yet, in order to estimate robust interaction effects, we had to combine some occupational groups in broader categories. We did this as follows:

- white collar: elite plus officials and other white collar occupations;
- working class: skilled and unskilled workers as well as artisans and shopkeepers (the latter because socio-economically they resembled the working class more than the white collar group, see section 2.2);
- farmers.

As to lags in the price effects, the following specification yielded the most consistent and robust results:

- past high prices: the de-trended price level of the four past semesters was on average at least one standard deviation above the mean (i.e. above zero);
- current and future high prices: the de-trended price level of the current and next two quarters was on average at least one standard deviation above mean zero.

The estimates for our hazard model with price effects are in table 2. (Note that we have only fitted the version without the family-level variance components. In the future, we will try to fit the multilevel version).

In order to test whether different socio-economic groups reacted differently to high rye prices, we specified our price effects in the same way as Bengtsson and Dribe (2004), i.e. we estimate the effect of high prices in a reference category and then see whether the other occupational groups reacted differently to high prices in comparison with the reference group. So for the reference group the p-values for the effects of prices correspond to the null-hypotheses that there is no effect, while the p-values for price effects on other categories show the probability associated with the null-hypothesis that there is no difference with the reference category. Like Bengtsson and Dribe, we selected the well-off as the reference group. In our case, we argue that this is the white collar group, consisting of people with the most favorable and stable incomes.

Table 2. Logistic discrete time hazard model of effective conception in closed birth intervals, western-Netherlands, 1825-1885; with estimates for the effects of high prices in different socio-economic groups

	b	*s.e.*	*p*
Intercept	-3.6784	0.2335	<.0001
Age of mother	-0.0122	0.0067	0.0689
Marriage duration	-0.1406	0.0158	<.0001
Months since previous birth	0.2510	0.0146	<.0001
Month²	-0.0081	0.0007	<.0001
Month³	0.0001	0.0000	<.0001
CEB within marriage	0.2190	0.0260	<.0001
Death of previous child	0.5793	0.0737	<.0001
Occupation of father			
White collar (ref)			
Farmer	-0.0772	0.1257	0.5391
Working class	-0.2892	0.1186	0.0147
Religion of parents			
Both liberal Protestant (ref)			
Both Catholic	0.2013	0.0531	0.0002
Orthodox Protestant	0.1951	0.0923	0.0345
Mixed Catholic-Protestant	0.0914	0.1309	0.4850
Both Jewish	-0.3794	0.3073	0.2169
Both unknown	-0.1616	0.1067	0.1299
Other	0.6077	0.2217	0.0061
Number of children alive	-0.1262	0.0376	0.0008
x marriage duration	0.0083	0.0025	0.0008
Past high prices			
White collar (ref)	0.0202	0.0105	0.0545
Farmers	-0.0236	0.0109	0.0312
Working class	-0.0217	0.0107	0.0429
Current and future high prices			
White collar (ref)	-0.0518	0.0219	0.0179
Farmers	0.0485	0.0221	0.0279
Working class	0.0485	0.0220	0.0273
Sample of origin			
Utrecht (ref)			
Akersloot	-0.1350	0.0574	0.0187
Amsterdam	-0.0032	0.0809	0.9684
Number of births			2206
Number of marriages			516
-2 Log Likelihood			15786
AIC			15838

The estimates indicate that when *past* prices had been high the most favorable socio-economic group was not slowing down parity progression at all. Rather on the contrary: if anything they did even speed up their family building process. Maybe this was a response to a slowing down of parity progression during times of crisis. Indeed, when the *current* and *expected* future prices where exceptionally high, the white collar group did slow down parity progression. Yet, in the less favorable occupational groups this providence was much smaller, if present at all. The differences with the white collar group are significant, but not their slowing down of parity progression in response to high current and expected prices as such (see figure 11). The estimated effects of past prices are small and nonsignificant: the working class and the farmers where anything but speeding up their pace of parity progression if past prices were high. Again, the difference with the white collar group is significant, but not the slowing down as such.

Figure 11. Estimated effects of past and current plus future high prices on the speed of parity progression

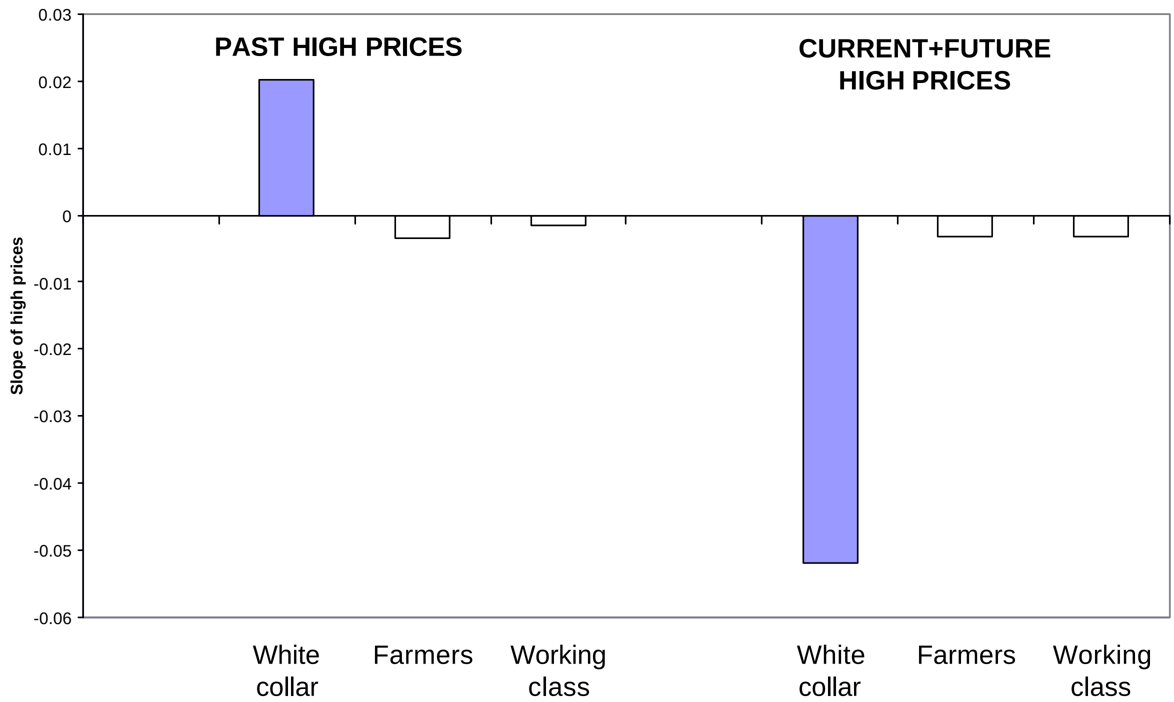

Note: white blocks are non-significant effects

In sum, the price effects suggest that the behavior of the elite may be characterized by prudence if prices were expected to be high. This providence in response to the expected economic situation appears to have been largely absent in the other occupational groups. The working class and farmers, i.e. the least well-off groups, appear to have been basically irresponsive to high prices.

5. Discussion

We are looking forward to your reactions in Tours, Saturday morning 8.30h, 23 July 2005.

References

Bengtsson, T. 2004. " Living Standards and Economic Stress." Pp. 27-59 in *Life Under Pressure. Mortality and Living Standards in Europe and Asia, 1700-1900*, eds. T. Bengtsson, C. Campbell, J. Z. Lee, and et al. Cambridge (Massachusetts): The MIT Press.

Bengtsson, T. and M. Dribe. 2004. "Deliberate Control in a Natural Fertility Population: Southern Sweden 1766-1865." Paper presented at *the Annual Meeting of the Population Association of America, 1-3 April 2004*, Boston, MA.

Clegg, E. J. 2001. "Starting, spacing and stopping in the reproductive histories of outer Hebridean families." *Journal of Biosocial Science* 33(3):405-26.

Damsma, D. and J. Kok. 2005. 'Ingedroogde harten? Partnerkeuze en sociale reproductie van de Noord-Hollandse boerenstand in de negentiende en vroeg-twintigste eeuw', pp. 285-307. in: *Genegenheid en gelegenheid. Twee eeuwen partnerkeuze en huwelijk., eds* J. Kok and M.H.D. van Leeuwen , Amsterdam: Aksant. ,

David, P. A. and T. A. Mroz. 1989a. "Evidence of fertility regulation among rural French villagers, 1749-1789. A sequential econometric model of birth-spacing behavior (Part 1)." *European Journal of Population* 5(1):1-26.

David, P. A. and T. A. Mroz. 1989b. "Evidence of fertility regulation among rural French villagers, 1749-1789. A sequential econometric model of birth-spacing behavior (Part 2)." *European Journal of Population* 5(2):173-206.

Ellison, P. T. 2001. *On Fertile Ground. A Natural History of Human Reproduction*. Cambridge (Massachusetts): Harvard University Press.

Ellison, P. T. 2003. "Energetics and reproductive effort." *American Journal of Human Biology* 15(3):342-51.

Giele, J.J. and G.J. van Oenen. 1974, 'De sociale struktuur van de Nederlandse samenleving rond 1850', *Mededelingenblad van de Nederlandse Vereniging voor Sociale Geschiedenis* 45, 2-32.

Giele, J.J. and G.J. van Oenen. 1976, Theorie en praktijk van onderzoek naar de sociale struktuur, *Tijdschrift voor Sociale Geschiedenis* 2, 167-185

Gray, R. 1983. "The Impact of Health and Nutrition on Natural Fertility." Pp. 139-62 in *Determinants of Fertility in Developing Countries, Vol. 1*, eds. R. A. Bulatao and R. D. Lee. New York: Academic Press.

Henry, L. 1961. "Some data on natural fertility." *Eugenics Quarterly* 18:81-91.

Horlings, E., J.-P. Smits. 1996. 'Private consumer expenditure in The Netherlands, 1800-1913', *Economic and Social History in the Netherlands* 7: 15-40.

Jonge, J.A. de. 1968 *De industrialisatie in Nederland tussen 1850 en 1914.* (Amsterdam.

Knippenberg, H. 1995. "Het hart van Nederland. De provincie Utrecht in de periode 1800-1940." Pp. 7-36 in *De levensloop van de Utrechtse bevolking,* edited by K. Mandemakers and O. Boonstra.Assen: Van Gorcum.

Kok, J. 2003'Van achterblijver tot koploper. Bevolkingsgroei, gezondheid en gezin', Pp 261-298 in *Geschiedenis van Holland. IIIa 1795 tot 2000 (* eds. Thimo de Nijs and Eelco Beukers . Hilversum: Verloren, |

Kok, J. 2004 .'Choices and Constraints in the Migration of Families. Central Netherlands 1850-1940, *History of the Family. An International Quarterly* 9 (2): 137-158.

Mandemakers, K., 2000, 'The Netherlands. Historical Sample of the Netherlands', Pp 149-177 in: eds P. K. Hall, R. McCaa and G. Thorvaldsen , *Handbook of International Historical Microdata for Population Research* Minneapolisi: Minnesota Population Center..

Meere, J.M.M. de. 1979 'Standen en klassen in een Noordhollands dorp. De gemeenter Wormer in 1844: de visie van een tijdgenoot', *Tijdschrift voor Sociale Geschiedenis* 5 : 245-263.

Meere, J.M.M. de. 1982. *Economische ontwikkeling en levensstandaard in Nederland gedurende de eerste helft van de negentiende eeuw.* 's-Gravenhage: Nijhoff.

Mroz, T. A. and D. R. Weir. 1990. "Structural Change in Life Cycle Fertility During the Fertility Transition: France Before and After the Revolution of 1789." *Population Studies* 44(1):61-87.

Mroz, Th. A. and D. R. Weir. 2003. "Stochastic Dynamic Optimization Models With Random Effects in Parameters: an Application to Age at Marriage and Life-Cycle Fertility Control in France Under the Old Regime." Pp. 483-510 in *History Matters: Essays in Honor of Paul A. David*, edited by T. Guinnane, W. Sundstrom, and W. Whatley. Palo Alto (CA): Stanford University Press.

Posthumus, N.W..1943. *Nederlandse prijsgeschiedenis. Part 1. Goederenprijzen op de beurs van Amsterdam 1585-1914, wisselkoersen te Amsterdam 1609-1914*(Leiden: Brill.

Tijms, W. 2000. “De prijzen van agrarische producten tussen 1546 en 1990”, *Historia Agriculturae* 31(Groningen/Wageningen: Nederlands Agronomisch Historisch Instituut .

Van Bavel, J. 2003. "Birth spacing as a family strategy. Evidence from 19th century Leuven, Belgium." *The History of the Family: An International Quarterly* 8(4):585-604.

Van Bavel, J. 2004a. "Deliberate birth spacing before the fertility transition in Europe: evidence from 19th century Belgium." *Population Studies* 58(1):95-107.

Van Bavel, J. and J. Kok. 2004. "Birth spacing in the Netherlands. The effects of family composition, occupation and religion on birth intervals, 1820-1885." *European Journal of Population* 20(2): 119-40.

Van Poppel, F., M. Jonker and K. Mandemakers. 2005. ‘Differential infant and child mortality in three Dutch regions, 1812-1909’, *Economic History Review* LVIII (2): 272-309.

Wilson, C., J. Oeppen, and M. Pardoe. 1988. "What is natural fertility? The modelling of a concept." *Population Index* 54(1):4-20.

Wood, J. W. and M. Weinstein. 1988. "A model of age-specific fecundability." *Population Studies* 42(1):85-113.

Wood, J. W. 1994. *Dynamics of Human Reproduction. Biology, Biometry, Demography*. New York: Aldine De Gruyter.

Yamaguchi, K and L. R. Ferguson. 1995. "The stopping and spacing of childbirths and their birth-history predictors: rational-choice theory and event-history analysis." *American Sociological Review* 60(2):272-99.
