Contraceptive Use and Birth Intervals

Gauthier Tshiswaka-Kashalala and Steven F. Koch

ERSA working paper 563

December 2015

Contraceptive Use and Birth Intervals*

Gauthier Tshiswaka-Kashalala†   Steven F. Koch‡

Abstract

We develop a model linking contraceptive efficiency to birth spacing decisions that incorporates the costs and benefits of child-rearing on the potential mother, as well as the stochastic process surrounding human reproduction. The model fits within the realm of optimal stopping-time problems, which naturally leads to the development of a First Hit Time duration model that we estimate using data from the Democratic Republic of Congo. Increased contraceptive efficacy is found to increase time to first birth. Furthermore, the results are consistent with the hypothesis that children are normal goods, in that both income and child-related benefits are associated with decreased durations to childbirth.

J.E.L: C41, J13

---

†Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, Republic of South Africa.

‡Department of Economics, University of Pretoria, Pretoria, Republic of South Africa; (0) 27-12-420-5285, (F) 27-86-691-2749. `steve.koch@up.ac.za`

1 Introduction

In family planning, optimal family size is achieved mainly through the use of contraceptives. As such contraceptive use has two main goals: delay the onset of childbearing and/or increase the duration of the intervals between births. The impact of contraception in delaying childbirth, stopping unwanted childbearing and postponing childbirth is an important contributor to fertility reduction (see Goldin and Katz, 2002; Moultrie et al., 2012).

Contraception derives its importance from the uncertainty at the center of the human reproduction process. In particular, when looking for some pattern in any birth history data, one quickly realizes that the number and timing of births is naturally uncertain. One way of dealing with uncertainty in birth spacing would be to assume that for every fecundable woman there is an underlying stochastic process leading to childbirth. Then, borrowing from the current literature on event history analysis, two types of stochastic fertility models can be formulated. A stochastic hazard fertility model where the childbirth hazard rate is some suitable function of the underlying stochastic process (see Woodbury and Manton, 1977; Yashin and Manton, 1997), and a birth interval model, which assumes that a woman becomes pregnant when the underlying stochastic process first satisfies a specified condition; thus, the time until birth of the next child is a *first hitting time* (see Aalen and Gjessing, 2001; Abbring, 2012).

We are interested in both types of models. In particular, we present a stochastic childbirth hazard model, as well as a first hitting time (FHT) fertility model. Both models use the diffusion process describing the dynamics of the continuation value of contraception as the main building block, and, therefore, provide similar predictions. While the stochastic childbirth hazard directly models a woman's childbirth hazard rate, the FHT fertility model arrives at the childbirth hazard by estimating the density and survival function of the time-to-birth, computing the hazard as a ratio of the two. However, directly modeling the hazard is limited, because closed-form solutions require specific assumptions about the underlying stochastic process. The more flexible FHT model is more easily extended to account for practical empirical problems related to censoring, which we account for through the application of a propensity score.

2

The duration of the interbirth interval can be affected by a variety of factors, other than contraception (and its efficacy) that are more parity specific, such as breastfeeding duration, temporary postpartum infecundity and the characteristics and survival of the preceding child. In this research our goal is to explicitly address the contribution of contraceptive efficiency. To assess the net effect of contraception, we focus on its efficacy and the relation to first birth timing. We choose to focus on the duration to first childbirth, because it has the advantage of not depending on parity-specific factors. Moreover, the age at which childbearing begins is not only an important determinant of the overall level of fertility, it is a key factor in the realized level of a woman's human capital investment, *i.e.* education and work experience (see Klepinger et al., 1995; Upchurch and McCarthy, 1990; Fitzenberger et al., 2013). It is generally assumed that increased age at first birth (see te Velde et al., 2012) and longer birth intervals reduce the number of children a woman can have, although Bongaarts and Casterline (2013) suggest that these intervals are naturally longer in Africa than in other regions. Unfortunately, not so much attention has been directed towards explicitly understanding the behavioral pathways linking contraceptive efficiency to birth spacing and timing (see Yeakey et al., 2009), which influences the fertility transition, and underpins the contribution of this research.

For our analysis we use information on the timing of a woman's first birth in an attempt to link contraceptive efficiency to birth spacing and contribute to the debate over why the fertility transition has so far eluded the Democratic Republic of Congo. According to recent research, while most countries have completed or are well advanced in the transition to low fertility, DR Congo is still far from meeting conditions for a sustained fertility transition (see Romaniuk, 2011). In the empirical analysis we focus on the timing of the first birth, given the high fertility rate among Congolese teenagers and to mitigate worries surrounding the effects of parity on birth spacing. Defining duration as the time between first intercourse and first birth, we find that the efficiency of contraception plays an important role in increasing the observed durations, as predicted by our model.

Furthermore, since our model explicitly incorporates income and child-related ben-

efits, our results suggest a partial explanation for the income-fertility puzzle. At the aggregate, empirical evidence invariably suggests that, within a given society, fertility is often higher in poorer families (see Becker, 1960; Jones and Tertilt, 2008). Further evidence suggests that countries with higher average fertility levels have lower average levels of industrialization (see Galor and Zang, 1997; Bloom et al., 2009). Despite the aggregate evidence, our empirical results support the hypothesis that children are normal goods, since both income and child-related benefits are associated with reduced childbirth 'intervals'. Because effective contraception is costly, access to it is limited to those in better economic circumstances. Such an explanation does not deny that child-drearing costs may be socio-economic status expenditures directly related to parents' income or that there are implicit costs associated with parent's time spent looking after a child that are linked to wage rates in the labour market (see Becker, 1965; Mincer and Polachek, 1974) – in fact, our model includes such costs – or that these costs could result in a quality-quantity tradeoff (see Becker and Lewis, 1973; Leibenstein, 1975; Caldwell, 1976). Instead, such an explanation provides a means for understanding how parents might achieve their objectives.

2   Continuation Value of Contraception

Assume that from first intercourse and for the rest of her sexually active life, a woman makes decisions about the level e_t of contraception efficiency at each time t, where fertility decisions are based on the reward she expects to derive from using contraception. If the contracepting woman is forward looking at parity P, her reward from contraception at this parity is the full flow of resources plus the expected discounted future earnings (*i.e.*, the continuation value), less contraception costs $h(e_t)$.

Consider a woman at parity P, who experiences natural fecundity p_t, has access to I_t resources per unit of time, and b_t net child related resource benefits per child. Further, assume that if she decides to move to parity $P + 1$, she will have to take some time off to care for the newborn child, thus losing a possible α_t percent of her resources I_t. If t is a *stopping time*, and V denotes the total expected discounted reward from following an optimal contraception strategy over a finite horizon $[0, T]$, where T is the

time to next childbirth or the advent of menopause. Given information at the beginning of a birth interval, $t = 0$, the woman's total expected reward from contracepting with efficiency e_t can be thought of as the sum of two terms: expected reward over $[0, t)$ and the continuation value over $[t, T)$ (see Stokey, 2009)

$$V_0 = \text{Expected reward over } [0, t) + \text{Expected reward over } [t, T)$$

$$= \int_0^t e^{-\rho s} \Big[(1 - \pi_s) (R_{1s} - h(e_s)) + \pi_s (R_{2s} - h(e_s)) \Big] ds + e^{-\rho t} W_t(e, P),$$

where $\pi_t[\equiv (1 - e_t)p_t]$ is the probability of falling pregnant, ρ is the rate of time preference, and $R_{1t}[\equiv I_t + b_t P]$ and $R_{2t}[\equiv (1 - \alpha_t)I_t + b_t(P + 1)]$ are the streams of the woman resources at parity P and $P + 1$ respectively. After multiplying through and collecting terms, the total expected reward becomes

$$V_0 = \int_0^t e^{-\rho s} [I_s + b_s P + \pi_s (b_s - \alpha_s I_s) - h(e_s)] ds + e^{-\rho t} W_t(e, P). \tag{1}$$

By the Martingale Representation Theorem (see Gawarecki and Mandrekar, 2011, pg. 49), the continuation value of contraception, $W_t(e, P)$ in (1), is described by a diffusion process which solves the following linear stochastic differential equation (see Sannikov, 2008)

$$dW_t = (\rho(t)W_t - \mu(t)) \, dt + \phi(t) dB_t, \tag{2}$$

which, provided that $\sup_{0 \leq t \leq T} \Big[|\mu_t| + |\rho| + |\phi_t| \Big] < \infty$ and $E(|w_0|^2) < \infty$, has a unique solution,

$$W_t = e^{\rho t} \left[w_0 - \int_0^t e^{-\rho s} \left(\mu_s ds - \phi_s dB_s \right) \right], \tag{3}$$

where B_t is a Gaussian Brownian motion, $\mu_t \equiv \Big[\Big(I_t + n_t N - h(e_t) \Big) - \pi_t(\alpha_t I_t - n_t) \Big]$ is the expected current net rewards from contracepting with efficiency level e_t, and $\phi_t \equiv \pi_t \sigma(p_t)$ are the diffusion coefficients. Assume for illustration purposes that the childbirth risk process is time homogeneous. Then, the expected current net flow of rewards μ and the diffusion coefficient ϕ do not depend explicitly on time t, and the

diffusion process in (2) can be reduced to

$$dW_t = [\rho W_t - \mu] \, dt + \phi dB_t \tag{4}$$

with realization

$$W_t = \frac{\mu}{\rho} + e^{\rho t} \left[w_0 - \frac{\mu}{\rho} + \phi \int_0^t e^{-\rho s} dB_s \right], \tag{5}$$

where ρ is assumed to be negative, $\rho < 0$. A negative rate of time preference implies that if a contracepting woman's objective is to maximize the expected utility of the continuation value W_t at a stopping time t, when she decides to stop contraception, then, holding total value constant, she would prefer an increasing series of continuation values to a declining one; thus, there is a negative rate of time preference (*i.e.*, $\rho < 0$) for choices among continuation value sequences (see Loewenstein and Prelec, 1991). In statistical terms, assuming $\rho < 0$ makes the process W_t *mean-reverting* (see Aalen et al., 2008), which means that the continuation value of contraception will tend to oscillate around some equilibrium state, and ρ is the speed of reversion.

However, the expected net benefit of childbearing can be negative or positive, $\mu \in \mathbf{R}$. If the net flow of resources at parity P is bigger than the expected loss in earnings due to childbearing, μ will be positive, and vice-versa if the expected loss in earnings is higher than the net flow of resources at parity P. If $\mu > 0$, the continuation value W_t is attracted to negative values and the birth interval is likely to be shorter, as illustrated in the left panel of Figure 1. On the other hand, the right panel of the same figure shows that if $\mu < 0$, the process W_t generally stays positive and away from zero, such that giving birth is an accidental or unwanted event.

For fixed t and $W_0 = w_0$, with probability 1, W_t has a *mean-reverting* Gaussian distribution

$$W_t \sim \mathcal{N} \left[\frac{\mu}{\rho} + e^{\rho t} \left(w_0 - \frac{\mu}{\rho} \right), -\frac{\phi^2}{2\rho} \left(1 - e^{2\rho t} \right) \right], \tag{6}$$

where $\rho < 0$ and $\mu \in \mathbf{R}$. The ratio μ/ρ is the intertemporal equilibrium level of the continuation value of delaying child birth, and $w_0 - \mu/\rho$ is the deviation from that equilibrium and the force pulling W_t back toward μ/ρ (see Aalen et al., 2008, p. 421).

The continuation value of contraception becomes stationary over time. As time
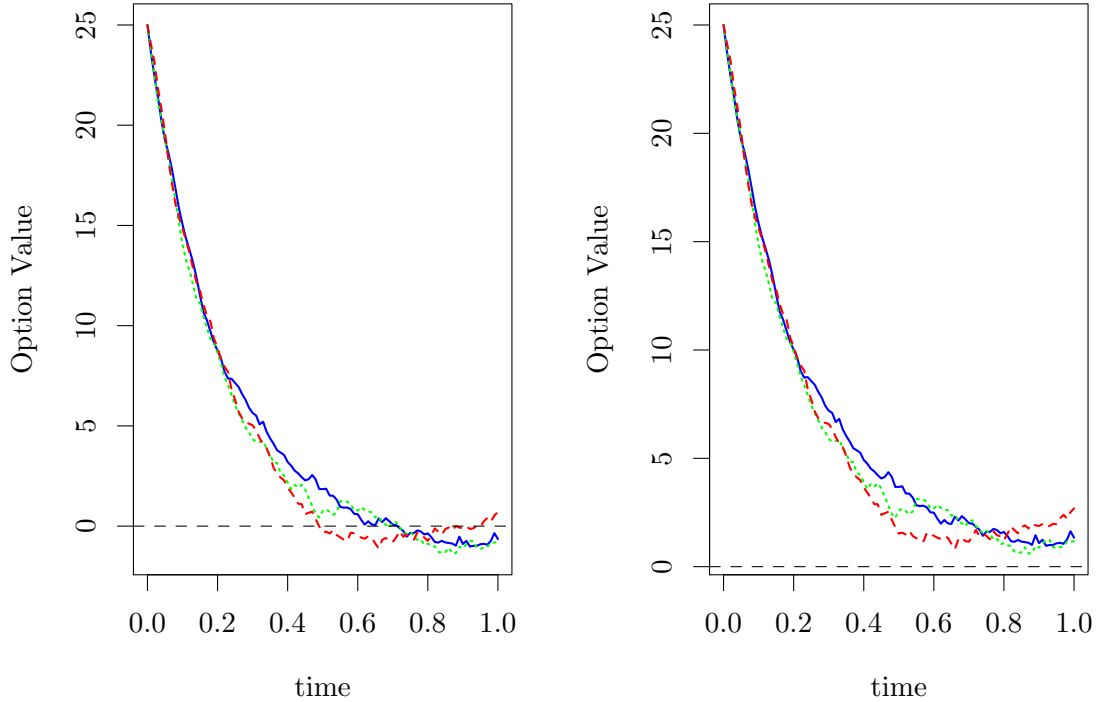
6

Figure 1: Three simulated paths for the continuation value W_t for parameter values $\rho = -5$, $\phi = 2$, $w_0 = 25$, when $\mu = 5$ (left panel) and $\mu = -5$ (right panel). Differences within each panel driven by realizations of the simulated stochastic process.

passes W_t itself will not converge, because of the noise dB_t, but its expected value moves from its starting point w_0 and converges to the mean reversion point μ/ρ. From (6), letting $t \to \infty$ suggests that the process is ergodic, and its invariant law is the Gaussian density with mean μ/ρ and variance $-\phi^2/2\rho$

$$W_\infty \sim \mathcal{N}\left(\frac{\mu}{\rho}, -\frac{\phi^2}{2\rho}\right). \tag{7}$$

3 Childbirth Hazard Rates

As implied above, human reproduction is a stochastic dynamic process, which naturally implies that the risk of falling pregnant is a stochastic process. In survival analysis, the goal is to evaluate the rate at which a stochastic dynamic process fails, and the relation

between that rate and observed and unobserved covariates. Unfortunately, standard survival models rarely take into account the dynamics of unobserved covariates. Thus, we underpin our analysis using a stochastic hazard model developed by Yashin (1984), which has the necessary properties to characterize a woman's childbirth hazard as a stochastic process. Unlike models of fixed unobserved heterogeneity used in standard survival analysis, Yashin's (1984) approach has the advantage of explicitly taking into account the dynamics of the underlying stochastic covariates leading up to the event, and gives rise to non-proportional childbirth hazards.

Initially, Woodbury and Manton (1977) proposed a random-walk model of the mechanics of human physiological aging and mortality. In their model, the hazard rate is a quadratic function of an Ornstein-Uhlenbeck diffusion process. Their model has been extended to deal with the combination of observed and unobserved state variables by using equations similar to the Kalman filter equations developed to estimate signals (see Yashin, 1984). However, given the current state of knowledge, obtaining such analytical results for hazard functions other than quadratic functions, or for more complex diffusion processes, is not yet feasible.

3.1 Marginal versus Conditional Birth Hazard Rates

Given the stochastic nature of the human reproduction process, assume that W_t is an unobserved process as described in (2), which defines a woman's randomly changing willingness to conceive and encompasses changes in her immediate environment and influences her individual childbirth hazard through time, and leads (potentially) to childbirth.

Let $k(W_t)$ be a nonnegative stochastic process on a continuous time and state space, and assume that it describes the childbirth hazard at time t of a woman, with unobserved continuation value W_t, randomly selected from a population of fecundable women. Since W_t is a developing stochastic process, it follows that the lifetime variable T, which denotes the time until the occurrence of a childbirth, can be interpreted as the time at which her cumulative childbirth hazard strikes a randomized barrier H. If the latter is exponentially distributed with mean 1, and is stochastically independent of W_t (see

Aalen and Gjessing, 2004) then

$$T = \inf \left\{ t : \int_0^t k(W_s)ds = H \right\}, \tag{8}$$

where the distribution of T depends on the path of W_t on the same interval. In particular conditional on W_0^t and the trajectory of W_t for the interval $[0, t]$, the distribution of the woman's time to childbirth T is determined by $k(W_t)$ such that

$$S(t|W_0^t) = P(T > t|W_0^t) = \exp\left[-\int_0^t k(W_s)ds \right], \tag{9}$$

where $S(t|W_0^t)$ denotes the conditional survival function, and $\mathrm{E}[k(W_s)] < \infty$ for any $s > 0$. Since individuals are naturally dissimilar and face, at any time t, a different set of factors, every woman will have her own level of risk of falling pregnant, which we assume to be driven by the benefits she expects to derive from contraception, *i.e.* her continuation value of contraception. Assuming the probability distribution of W_t and the functional form of $k(W_s)$ are known, T is measured for N women and it is the only observable information available, then the marginal (*i.e.* population) survival function, averaged over W_s, is

$$S(t) = P(T > t) = \mathrm{E}\left[\exp\left(-\int_0^t k(W_s)ds \right) \right] = \exp\left(-\int_0^t \mathrm{E}[k(W_s)|T > s]ds \right), \tag{10}$$

where T is related to W_s by

$$\mathrm{P}(T > t|W_s, s \leq t) = \exp\left(-\int_0^t W_s ds \right). \tag{11}$$

Thus, the relation between the observed and the conditional hazards is

$$\theta_t = \mathrm{E}\left[k(W_t)|T > t \right], \tag{12}$$

where θ_t is the marginal, or observed, childbirth hazard (Yashin, 1984). The observed childbirth hazard rate is the risk a woman to becomes pregnant at time t, such that $\theta(t)dt$ represents the instantaneous probability that a woman will become pregnant in

the interval $(t, t + dt)$ given that this woman is not pregnant at t. The hazard function can be defined as

$$\theta(t) = -\mathrm{P}(T > t)^{-1} \frac{\mathrm{d}P(T > t)}{\mathrm{d}t}. \tag{13}$$

3.2 Individual Childbirth Hazard as the Square of W_t

In applied work, the computation of the observed hazard θ_t is simplified if averaging over the unobserved variable can be expressed in an analytic form. Therefore, tet us assume that the individual childbirth hazard rate is a quadratic function $Z_t = k(W_t) = W_t^2$ (see Yashin, 1984). For a stochastic process W_t, Yashin (1984) proved the Gaussian property of the distribution of survivors at any time t, $P(W_t \leq w | T > t)$, as well as the following proposition in the case where the distribution of the initial position of W_t is a normal distribution with mean m_0 and variance v_0.

Proposition 1. *Let the stochastic process W_t satisfy*

$$dW_t = [\rho W_t - \mu] \, dt + \phi dB_t. \tag{14}$$

Then the average survival function is given by

$$\mathrm{E}\left[\exp\left(- \int_0^t W_s^2 ds \right) \right] = \exp\left\{ - \int_0^t \left[m^2(s) + v(s) \right] ds \right\}, \tag{15}$$

where the mean $m(s) = \mathrm{E}[W_t | T > s]$ and variance $\nu(s) = \mathrm{Var}[W_t | T > s]$ of the conditional distribution of W_t given survival are the unique solutions to the Riccati equations

$$\frac{dm(t)}{dt} = \rho m(t) - 2m(t)v(t) - \mu, \tag{16}$$

$$\frac{dv(t)}{dt} = \phi^2 + 2\rho v(t) - 2v^2(t), \tag{17}$$

with initial conditions $m(0)$ and $v(0)$.

If μ, ρ, and ϕ are known scalars, analytic form solutions to the system of nonlinear differential equations in (16) and (17) can be obtained. Thus a complete description of the hazard process conditional on survival is available. Assuming, for simplicity, that all coefficients are constant, solving first for the variance of the normal distribution of

survivors (*i.e.* non pregnant fecundable women), $v(t)$ in the separable (17), yields [1]

$$v(t) = \frac{1}{2}\rho + \frac{1}{2}\sqrt{\rho^2 + 2\phi^2} \tanh\left[t\sqrt{\rho^2 + 2\phi^2} - \operatorname{arctanh}\left(\frac{-2v(0) + \rho}{\sqrt{\rho^2 + 2\phi^2}}\right)\right]. \quad (18)$$

Following Aalen et al. (2008), define the constant $C = \sqrt{\rho^2 + 2\phi^2}$ and let

$$h = \operatorname{arctanh}\left(\frac{-2v(0) + \rho}{\sqrt{\rho^2 + 2\phi^2}}\right)$$

and then solve for $m(t)$ to get

$$m(t) = \frac{\sqrt{\sinh^2(tC - h) - \cosh^2(tC - h)}}{\cosh(tC - h)}\left[\frac{\cosh(h)}{\sqrt{\sinh^2(h) - \cosh^2(h)}}m(0)\right.$$
$$\left. - \mu \int_0^t \frac{\cosh(sC - h)}{\sqrt{\sinh^2(sC - h) - \cosh^2(sC - h)}}ds\right]. \quad (19)$$

Using the identity $\cosh^2(x) - \sinh^2(x) = 1$ reduces (19) to

$$m(t) = \frac{1}{\cosh(tC - h)}\left[\cosh(h)m(0) - \mu \int_0^t \cosh(sC - h)ds\right], \quad (20)$$

and since we know that $\cosh(ax+b) = \cosh(ax)\cosh(b) + \sinh(ax)\sinh(b)$ and $\sinh(-x) = -\sinh(x)$, the integral in (20) can be computed as

$$\int_0^t \cosh(sC - h)ds = \frac{1}{C}\left[\sinh(tC - h) + \sinh(h)\right]. \quad (21)$$

Substituting (21) into (20) yields the following solution for the mean of the distribution of survivors

$$m(t) = \frac{1}{\cosh(tC - h)}\left\{\cosh(h)m(0) - \frac{\mu}{C}\left[\sinh(tC - h) + \sinh(h)\right]\right\}. \quad (22)$$

Similar and related results have been reported in Wenocur (1990) and Aalen and Gjessing (2004).

---

[1] All analytical solutions are obtained using the symbolic computation software Maple 13 from Maplesoft.

3.3 Quasi-stationarity

In what follows we focus on the properties of the childbirth hazard rates explicitly reviewing the effects of the latent stochastic process W_t. Recall that the continuation value of contraception W_t is the latent underlying risk process leading to childbirth. As such, it represents an unobserved stochastic development that influences the individual hazard of becoming pregnant as time evolves. The solutions for the conditional moments m and v ()(22) and (18), respectively) suggest that the distribution of women who are not yet pregnant may stabilize over time, and has a normal limiting distribution $\mathcal{N}\left(-\mu/C, 1/2(\rho+C)\right)$. This quasi-stationary distribution is useful in understanding the shape of the hazard. In particular, it tends to produce non-proportional hazards, which characteristically stabilize at a positive value over time (see Aalen et al., 2008, chap. 10). In our case, the childbirth hazard converges to a limit

$$\lim_{t\to\infty} \theta_t = \left(-\frac{\mu}{C}\right)^2 + \frac{1}{2}(\rho+C). \tag{23}$$

As an illustration, we simulated the shape of the observed childbirth hazard assuming the hidden heterogeneity is a square of the stochastic process W_t. Figure 2 illustrates the different values of the population childbirth hazard, $\theta_t = m^2(t) + v(t)$, associated with different values of the expected current net reward of contraception μ_t.

If we compare the group of women with positive expected gain from childbearing ($\mu > 0$) to those with a negative expected gain from childbearing ($\mu < 0$), the quadratic model for the individual childbirth hazard produces non-proportional childbirth hazards between the two groups. Over time, these steadily converge towards a limiting value. Women who expect a loss of due to childbearing have childbirth hazards that are essentially declining, while the opposite is true for women expecting to gain.

For all its advantages, the stochastic quadratic hazard model has some serious shortcomings. Though mathematically convenient, the assumption that the unobserved heterogeneity is a quadratic function of the underlying stochastic process is somewhat arbitrary and may not be well-suited for survival analysis in population studies. In particular, it is hard to believe that all women with $\mu > 0$ have quasi-proportional in-
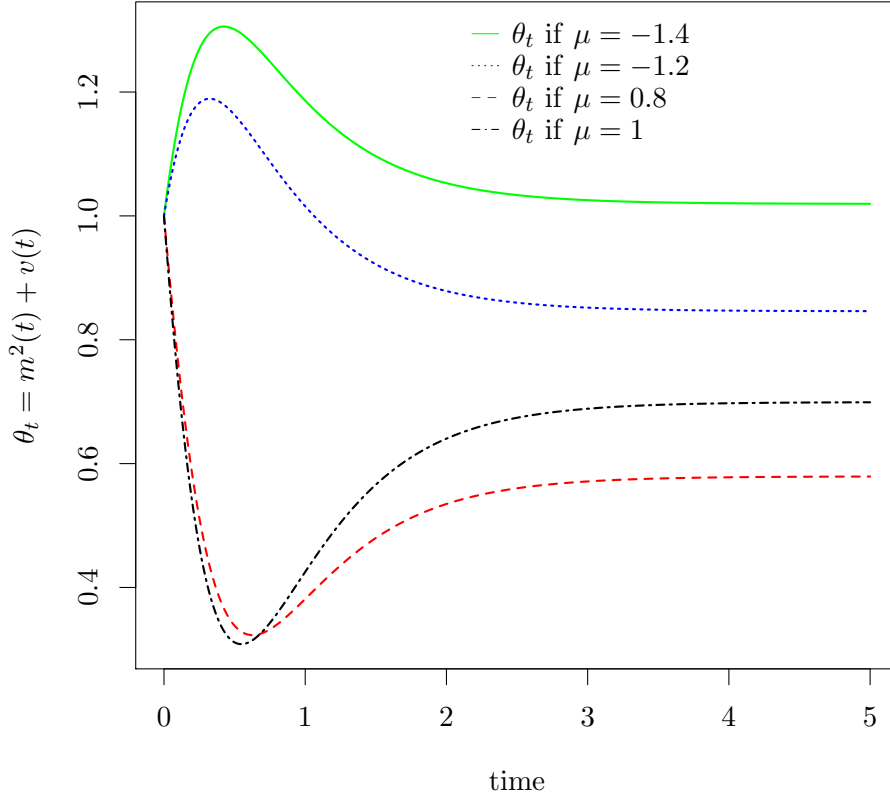
Figure 2: Population childbirth hazard θ_t for different values of μ, when $\rho = -1$ and $\phi = 1$, along with starting values $m(0) = 1$ and $v(0) = 0$.

creasing hazards, and vice versa. As seen in Section 3, for a negative value of μ, there is no guarantee that the childbirth risk process will reach the threshold in zero. It follows that it is practically impossible for the stochastic quadratic hazard model to account for *cure effects*, the fact that a certain portion of women will remain childless be it naturally or by choice.

4   FHT Fertility Model

Rather than assuming a stochastic quadratic hazard model, we interpret reproductive history data as first passage times of a threshold by sample paths of a stochastic process, and present an FHT model as a model for the onset of childbearing. Instead of directly modeling the population childbirth hazard, through the individual childbirth hazard

rates as shown in Section 3, one might first choose to model the birth interval before deriving the childbirth hazard. We define birth interval as the length of time between two successive live births, or between a woman's date of first intercourse and the birth of her first child; empirically, however, we focus on the latter.

FHT models are threshold models with regression structures that accommodate the effects of observed covariates and unobserved heterogeneity in duration data analysis. An FHT model is a useful alternative to the Cox proportional hazards model, and such models are gradually finding broad application, due to their conceptual appeal and flexibility (see Lee and Whitmore, 2006, 2010). In economics, first hitting times arise in structural models in which agents are assumed to solve an optimal-stopping problem with related rewards described by stochastic processes (see Stokey, 2009). Economic applications have so far been confined to labour economics, including Lancaster's (1972) strike duration study, the analysis of labour turnover (Whitmore, 1979), and the analysis of unemployment spells (Shimer, 2008).

To the best of our knowledge, the FHT model has not been previously applied to the study of the onset of childbearing using large household level survey data, especially in the context of a high fertility environment. The FHT fertility model we present here is driven by the latent stochastic process as defined in Section 2, based on the continuation value of delaying conception, which is assumed to influence women's birth spacing practices. We assume that a woman optimally times the decision to become pregnant, and the dynamics surrounding the decision are largely explained by this latent stochastic process. Applying the FHT framework to the field of population economics is inspired by the sequential fertility model introduced in Heckman and Willis (1976) and the discrete-time mixture duration model based on a latent process that crosses thresholds developed by Heckman and Vytlacil (2007).

4.1 A Basic First Hitting Time Fertility Model

The FHT framework provides a modeling structure that is both flexible and realistic enough to incorporate the impact of a variety of observed and unobserved demographic and socio-economic characteristics on the parameters of the latent stochastic process.

One of the main covariates in our analysis is contraceptive efficiency. We assume that at each point in time, the woman weighs the direct rewards of stopping contraception, against the value of retaining the option of postponing childbearing, given the primitive parameters and the history of the continuation value. In this case, she maximizes her expected discounted rewards by becoming pregnant when the continuation value of contraception hits a time-invariant threshold for the first time; implying that the optimal decision rule involves a threshold. Thus, our fertility model has two basic components: (1) a parent *latent stochastic process* in time $\{W_t, t \geq 0\}$ which describes the dynamics of the continuation value of postponing childbearing, with initial value $\{W_0 = w_0\}$ and (2) an *absorbing set* $\mathcal{B}$ in the state space of the unobserved stochastic parent process and defines its stopping condition.

For simplicity, we assume that the woman places equal value on both present and future utility, meaning that the rate of time preference ρ is equal to zero(see Ramsey, 1928). As a result, the diffusion process in (14) reduces to a Brownian motion with drift of the form

$$dW_t = -\mu dt + \phi dB_t. \tag{24}$$

This assumption is underpinned by the observation that the population hazards resulting from an Ornstein-Uhlenbeck process, as the one in (14), and the Brownian motion in (24) are almost identical (see Aalen et al., 2008). Also, working with Brownian motion with drift has the advantage of allowing us to use the well-known analytical result that the first hit time of a Brownian motion has an inverse Gaussian probability distribution.

Since the childbirth risk process is unobservable for the econometrician, the only observable effect of W_t is through the individual event time $T > 0$, when the woman gives birth. Let the process W_t start in a positive initial value $W(0) = w_0 > 0$, and assume that the timing of the birth coincides with the time when the process is absorbed in the absorbing boundary zero. Thus, the random variable T is defined as

$$T = \inf_{t \geq 0}\{t : W_t = 0\}, \tag{25}$$

Consider one fecundable woman for a moment. As she postpones childbearing

through birth control, the rewards she expects from postponing childbearing fluctuate. Over time, since first intercourse or her last live birth, she might experience a relatively steady decline in this contraception continuation value, and, eventually, it hits zero, the level at which we assume she gives birth. The first hitting time to zero is the woman's duration of the birth interval. On the other hand, the woman may have experienced a relatively steady increase in her level of the continuation value. In this case, she may never give birth. If the latter woman does give birth, it will be unplanned.

If we assume that a woman's level of the continuation value of contraception as a function of time is described by the Brownian motion in (24), and that all coefficients are constant, then the first hitting time T from the initial level of the continuation value w_0 to the threshold set at $W_t = 0$ has an inverse Gaussian probability distribution which has the following probability density function (*i.e.* p.d.f) (see Chhikara and Folks, 1989)

$$ f(t) = w_0 (2\pi\phi^2 t^3)^{-\frac{1}{2}} \exp\left\{ -\frac{(w_0 + \mu t)^2}{2\phi^2 t} \right\}, \quad \text{for} \quad \phi^2 > 0, \ w_0 > 0, \qquad (26) $$

and the associated survival function (*i.e.* c.d.f)

$$ S(t) = \Phi\left(\frac{w_0 + \mu t}{\sqrt{\phi^2 t}} \right) - \exp(-2w_0\mu/\phi^2)\Phi\left(-\frac{(w_0 - \mu t)}{\sqrt{\phi^2 t}} \right), \qquad (27) $$

where $\Phi(\cdot)$ is the c.d.f of the standard normal distribution, and μ and ϕ are the drift and volatility of the process, respectively.

The theory-based hazard rate for time to childbirth is found from $\theta(t) = f(t)/S(t)$ and its shape generally exhibits the same stability phenomenon as for the stochastic hazard derived in Section 3. Figure 3 illustrates the fact that regardless of initial value of the continuation value of contraception, all hazards for the Brownian motion converge to the same limiting hazard. Furthermore, as expected, the shape of the hazard rate is associated with the distance between the starting point and the point of absorption. At time $t = 0$, if the process W_t starts at a level close to zero relative to the quasi-stationary distribution, the childbirth hazard rate is essentially decreasing; if it starts at an intermediate value of w_0 the childbirth hazard first increases and then decreases; while if it starts at a value of w_0 far from zero the childbirth hazard rate is essentially
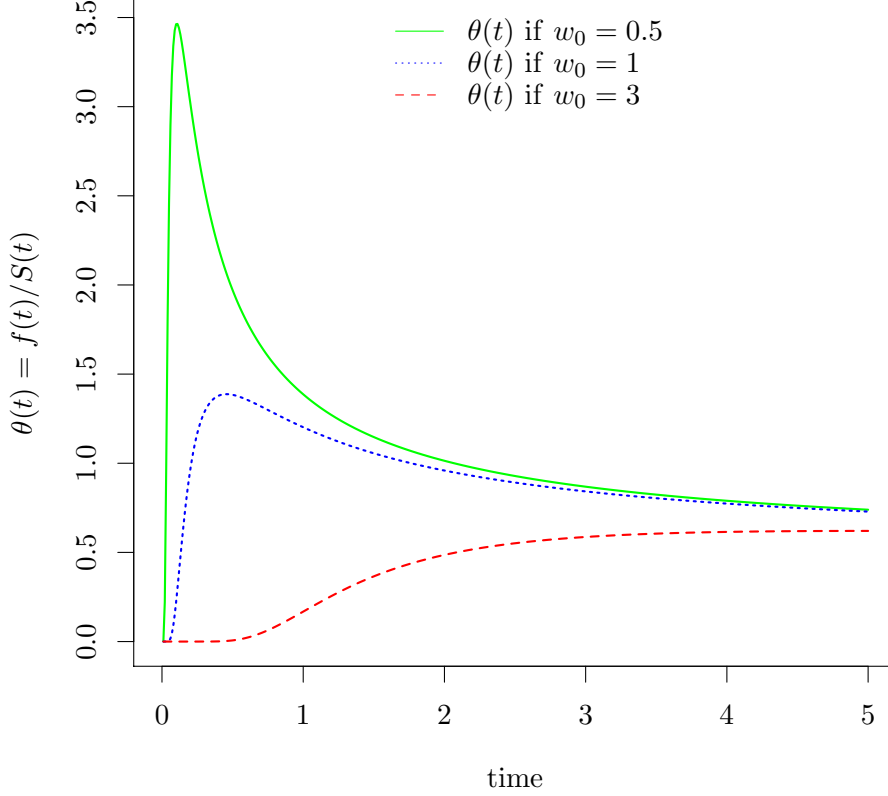
increasing (see Aalen et al., 2008).

Figure 3: Hazard rates for time to childbirth $\theta(t)$ when the process starts in different values of w_0, when $\mu = 1$ and $\phi^2 = 1$.

The drift parameter quantifies the rate at which the woman approaches childbirth. However, there is no guarantee that the process will reach the boundary set $\mathcal{B}$. We recognize the fact that for some women, the childbirth risk process W_t may diffuse away from the childbirth threshold for a long time, and diffuse almost directly toward it for others. This would be the case when some women are temporarily infertile or choose not to have a child, such that $T = \infty$. To use the terminology in Abbring (2012), infertile women make up an unobserved subpopulation that may be described as *stayers*, while those women who might choose not to have a child are *defecting movers*.

If $\mu \geq 0$, meaning that the expected loss from childbearing is less than or equal to current income, then there is a tendency to drift towards the childbirth threshold

17

zero. In this case, childbirth is a certain event, which will occur in some finite time with probability one. The mean survival time conditional on the event that the childbirth threshold is eventually reached is

$$E(t) = \frac{w_0}{|\mu|}, \qquad \text{for} \qquad \mu \neq 0.$$

To complete the model, we now turn our attention to the issue of estimating the model from data as a practical application of our FHT fertility model. So far, taking advantage of the fact that the probabilistic specification of the parent stochastic process in FHT models is usually explicit, parameter estimation for FHT models have been conducted mostly through maximum likelihood methods (see Lee and Whitmore, 2006).

In a total sample of $N = N_A + N_B$ women, each woman i who has given birth contributes probability density $f(t_i|w_{i0}, \mu_i)$ to the sample likelihood function, where t_i is the observed time of childbirth for $i = 1, \cdots, N_A$, while a woman j in the sample dataset who stays childless to the end of the study contributes the survival probability $S(t_j|w_{j0}, \mu_j) = 1 - F(t_j|w_{i0}, \mu_j)$, where t_j is the right-censored survival time of the woman for $j = N_A + 1, \cdots, N_A + N_B$. Then the sample likelihood function to be maximized should be of the form

$$\mathcal{L}(\theta|t) = \prod_{i=1}^{N_A} \left[f(t_i|w_{i0}, \mu_i) \right] \prod_{i=N_A+1}^{N_A+N_B} \left[S(t_i|w_{i0}, \mu_i) \right]. \tag{28}$$

However, if the expected loss from childbearing is higher than current income (*i.e* $\mu < 0$), a woman may never fall pregnant, because she is "cured". Let T be a random variable representing survival time to next child birth. Let $S(t)$ represent the survivor function for the population of woman of reproductive age where, like before, $S(t) = P(T > t)$. The cure rate is defined as $\lim_{t \longrightarrow \infty} S(t)$. In order to explicitly accommodate the *cure effect*, we enrich the FHT model with a new parameter c representing the propensity score of giving birth. In fact, for values of $\mu < 0$, the distribution is defective, that is $T(\infty) > 0$ with a probability density mass concentrated at $S = \infty$ and the expected

proportion of *cured women* given by

$$1 - c = P(T = \infty) = 1 - \exp\left(\frac{2w_0\mu}{\phi^2}\right),$$

which implies that the probability of childbirth $P(T < \infty) = 1 - P(T = \infty)$ may be less than 1. We denote the proportion of women who will eventually give birth to a child if given enough time by c, the *propensity rate*. The propensity rate may either be determined by the parameter values of the latent stochastic process when $\mu < 0$ or be a free parameter that is independently linked to covariates in the FHT regression model. It follows that the modified likelihood function for the *cure-rate* FHT model becomes

$$\mathcal{L}(c, \theta | t) = \prod_{i=1}^{N_A} c_i \left[f(t_i | w_{i0}, \mu_i) \right] \prod_{i=N_A+1}^{N_A+N_B} \left[1 - c_j F(t_i | w_{i0}, \mu_i) \right]. \tag{29}$$

Although the model has three parameters, namely w_0, μ, and ϕ, there are, statistically speaking, only two free parameters. The distribution only depends on the three parameters through two functions: μ/ϕ and w_0/ϕ. Thus, the variance ϕ^2 may be set to one without loss of generality, when considering time to childbirth (see Aalen et al., 2008).

5 Pathways to a First Child in DR Congo

In what follows we apply the FHT fertility model to data mothers' first births from the DR Congo's 2007 Demographic and Health Survey (DHS). Our analysis considers a sample of individuals, $i = 1, \cdots, n$ and the individual latent childbirth risk process $W_i(t)$ as defined in (24). For every fecundable woman i, if ϕ_i^2 the variance of $W_i(t)$ is set to one, the density of the first-hitting time T is inverse Gaussian distributed as in (26) with a vector of free parameters $\theta = (c_i, w_{0i}, \mu_i)$ representing the propensity rate, the initial value and drift of the childbirth risk process, respectively.

5.1 Incorporating Covariate Information

Focussing on a woman who has just had her first intercourse, some personal characteristics are expected to influence the underlying childbirth risk process W_t. In particular, the values of her parameters c, w_{0i}, and μ_i are linked to k covariates that are represented by the vectors $\mathbf{x}_i = (1, x_{i1}, \cdots, x_{ip})$, $\mathbf{y}_i = (1, y_{i1}, \cdots, y_{ir})$ and $\mathbf{z}_i = (1, z_{i1}, \cdots, z_{iq})$.[2]

We follow Xiao et al. (2012) and link the log-odds ratio of c to a linear combination of covariates as follows

$$\text{logit}(c_i) = \log\left(\frac{c_i}{1-c_i}\right) = \lambda_{i0} + \lambda_1 x_{i1} + \cdots + \lambda_p x_{ip} = \mathbf{X}'\lambda. \qquad (30)$$

Let us further assume that $\ln(w_{0i})$ and μ_i are linear in regression coefficients, and use an identity function of the form

$$\mu_i = \mathbf{Y}'\beta \qquad (31)$$

to link the parameter μ_i to the covariates, and the following logarithmic function to link the parameter w_{0i} to covariates

$$\ln(w_{0i}) = \mathbf{Z}'\gamma, \qquad (32)$$

where $\lambda = (\lambda_0, \lambda_1, \cdots, \lambda_p)'$, $\beta = (\beta_0, \beta_1, \cdots, \beta_q)'$ and $\gamma = (\gamma_0, \gamma_1, \cdots, \gamma_r)'$ are the respective covariate effects (see Aalen and Gjessing, 2001; Aalen et al., 2008).

We use the theoretical model presented in Section 4 as a guide in choosing the covariates to include in $\mathbf{X}'$, $\mathbf{Y}'$ and $\mathbf{Z}'$. According to Aalen et al. (2008), one of the major advantages of the threshold regression framework is its ability to differentiate between the effects of covariates on how far the risk process has advanced prior to the study (*i.e* the effects on the initial level w_0) and the causal effects on the dynamics of the risk process (*i.e* the effects on the drift μ_i), although some variables may have both effects.

Regarding the covariates to include in $\mathbf{X}'$, we assume that the propensity to give birth to a first child is determined by physiological and environmental factors; however, we only have information related to her age at first marriage, which we use. As for

---

[2]In the most general FHT model, the parameters of the process, threshold state and time scale may also depend on covariates (see Abbring, 2012).

w_{0i}, the initial value of the childbirth risk process, we assume that it will depend on external factors related to the young woman's socio-economic background at the time of her first sexual intercourse. These factors may include, among others, the woman's taste for risks, her general childhood environment and the age at which she first got married. With regard to the drift μ_i we can deduce, from the analysis in Section 2, that for a young woman with no children we have

$$\mu_i \equiv I_i(1 - \alpha_i \pi_i) + b_i \pi_i - h_i(e),$$

which suggests that values of the woman's income I_i induce higher values of μ_i, the rate at which the process moves towards the threshold, since $(1 - \alpha_i \pi_i)$ will always be positive. This is the same for the child related benefits b_i which is positively related to μ. Note that higher values of μ_i mean shorter intervals between first sexual intercourse and first birth. In the same vein, the prospect of higher income loss due to childbearing will increase this interval, since α_i is negatively related to μ, and the woman's probability of conception $\pi_i \equiv (1 - e)p_i$ is always positive.

Let us fix, for simplicity, the parameters representing the potential loss in income, α_i, and the one representing the woman's probability of giving birth, π_i. The equation for μ_i becomes

$$\mu_i = aI_i + b_i \pi - h_i(e), \tag{33}$$

where it is clear that, increased affordability through the rise in income or the availability of child related transfers will shorten the waiting time to first motherhood. Furthermore, assuming that the cost of contraception $h(\cdot)$ is a positive monotonic transformation of the efficiency of the contraception method e, our model suggests that higher values of e give rise to lower values μ, meaning that the use of a more efficient contraceptive method will increase the interval between first intercourse and first birth.

5.2   Data

The 2007 DHS for the DRC is a nationally representative survey for urban and rural residence, which provides information mainly on reproductive behaviour and reproduc-

tive health for $9,995$ women aged $15-49$, as well as $4,757$ men aged $15-59$. The choice of the Congo is dictated by the fact that, despite its size and a large population, very little is known about this country, and its fertility level remains among the highest in the world. With a total population estimated at around 70 million people unevenly distributed on a $2,344,858$ km^2 surface area, the Congo's fertility rate is estimated at 6.3. In the sample, some women were still virgins and were not included, while some had given birth to their first child at the time of the interview; still, others were childless and are right-censored observations in our analysis.

Deciding which covariates influence the two parameters of the underlying Brownian process, w_0 and μ, is facilitated by our theoretical model. We assume that baseline characteristics only influence the initial childbirth risk process state w_0, and include the woman's place of childhood residence and her age at first marriage. On the other hand, according to our theoretical model, other covariates, such as income, contraception efficiency, and child related benefits influence the drift in the childbirth risk process. In the absence of income information, a common DHS limitation, we assume that the woman's current wealth represents her long-term income. We define a woman's contraception efficiency as the interaction between the woman's literacy and her choice of contraceptive method.[3] Also, since Congo lacks a formal system of child related benefits, we let the number of the young woman's older siblings represent the value of the child related support she can expect to access at childbirth.

Summary statistics in Table 1 suggest that the majority of young women in our study are poor, live in small towns, have difficulties reading and have never used modern contraception in their lives. They further indicate that the average woman has a duration of 2 years and 9 months from the time of first intercourse to first childbirth, has 2 older siblings and was first married at the age of 18.

---

[3]In the data set contraception methods are classified as *traditional*, *folklore* and *modern*. Traditional methods include Periodic Abstinence (also known as Rhythm), Withdrawal and Abstinence. Modern methods include Pill, IUD, Injections, Diaphragm, Condom, Sterilization, Implants, Foam/Jelly and Lactational amenorrhea. Folkloric methods include all the methods not specifically mentioned, but believed to be less efficient than the traditional methods. If a woman has used both a traditional method and modern method then the latter takes priority. In the same vein, a woman who is recorded to have used a traditional method if she has used both a traditional method and folkloric method.

Table 1: Summary Statistics

Nominal Variables				
	Levels	%	$\sum$%	*Obs.*
Childhood Residence	City	34.25	34.25	2,213
	Town	65.75	100.00	4,248
Literacy	Cannot Read	54.02	54.02	3,490
	Can Read	45.98	100.00	2,971
Wealth	Non Rich	60.08	60.08	3,882
	Rich	39.92	100.00	2,579
Contraception Use	Never	45.95	45.95	2,969
	Traditional Only	31.92	77.87	2,062
	Modern Only	22.13	100.00	1,430
Continuous Variables				
	Min.	Median	Mean	Max.
Survival Time (Years)	0.58	2.75	4.09	35.08
Age at first Marriage	10.00	18.00	18.29	44.00
Number of Older Siblings	0.00	2.00	2.77	15.00

Sample size, $N = 6{,}461$.

5.3 Parameter Estimates

In order to show the importance of the concept of contraceptive efficiency, rather than the simple choice of contraceptive method, we start by fitting a model where the choice of contraceptive method is not interacted with the level of literacy. We then add the interaction between contraception choice and the level of literacy to the μ-term. Using the Akaike information criterion, it is clear that adding the interaction term improves the fit (lower AIC), suggesting that it is not enough to include the choice of a contraceptive method in the analysis, one must also look at how efficiently it is being used.

The estimated parameters for both models are shown in Table 1. Results are consistent with our theoretical model and suggest that contraception efficiency is negatively correlated with the rate at which the expected benefits from contraception approaches zero, which means higher contraceptive efficiency leads to a longer waiting period from the first intercourse to the birth of the first child. In particular, we note that the use of *modern* contraceptive methods by women who can read has a very strong negative effect on the onset of motherhood, as compared to those who have never used contraception. For those women who can read but have only used *traditional* contraception methods, the effect is statistically insignificant, but still has the expected negative effect on the drift.

Furthermore, the parameter estimates tell us something more consistent with the theoretical model. They suggest that the level of wealth, which we assume to be a good approximation of the woman's long-term income, and the presence of older siblings in the family have a statistically significant 'income effect' on children that implies children are normal goods. In particular, higher levels of wealth and the presence of older siblings motivate young women to have their first child earlier; they are positively related to the drift, μ, of the stochastic process. On the other hand, as expressed by the positive sign of the estimates linked to age at first marriage and childhood residence, young women who spent their childhood years in a small town or those who were married at a relatively advanced age start with a lower risk of giving birth to the first child.

24

Table 2: Threshold Regression Cure Rate Model Estimates

	Model 1		Model 2	
AIC	27,450.97		27,446.39	
	Estimate	Std. Err.	Estimate	Std. Err.
$\ln w_0$				
Intercept	.13398	.03484	.13519	.03480
Age at First Marriage	.03152	.00171	.03160	.00171
Childhood Residence	.04020	.01607	.03752	.01611
μ				
Intercept	.46178	.01518	.44445	.01621
Wealth	.04712	.01563	.03664	.01687
Number of older Siblings	.00597	.00279	.00593	.00280
Contraception (Traditional)	.02299	.01608	.03410	.02040
Contraception (Modern)	-.01424	.01893	.03137	.02968
Literacy			.07153	.02402
Contraception (Trad.)$\times$ Lit.			-.04876	.03308
Contraception (Mod.)$\times$ Lit.			-.10028	.03921
logit c				
Intercept	8.94160	.63608	8.97492	.64436
Age at First Marriage	-.21834	.02247	-.21944	.02266

Estimates of FHT model, based on the application of Xiao et al.'s (2012) STATA package `sttthreg`. Separate estimates presented for equations (30), (31) and (32). Recall, positive estimates for (31) imply reduced durations, while positive values for estimates in (32) relate to higher initial conditions, and, thus, imply longer durations.

5.4 Estimated Hazards and Probabilities of First Childbirth

As expected from our theoretical model, the shapes of the estimated hazards and cumulative hazards suggest nonproportionality. To show this, we group the women into three different risk groups: *High*, *Moderate* and *Low*-risk. We define a high-risk group to include women that are closer to the point of absorption than the moderate, who are closer to the point of absorption than the low-risk group. In other words, a woman belongs to the *high* risk group if her estimate of the initial value of the stochastic process is less than the *lower* quartile of the estimated values of w_0, *moderate* risk if it is between the lower and the upper quartiles, and *low* risk if it is *above* the upper quartile.
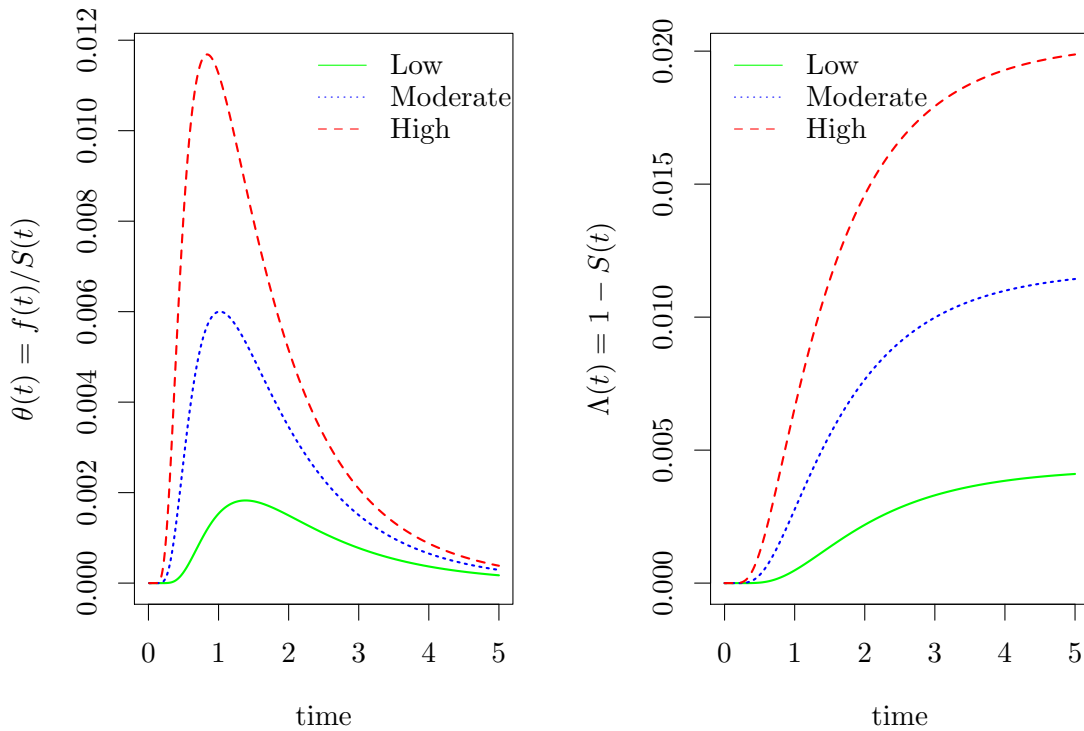

Figure 4: *Hazards* (left panel) and the corresponding *Probability of Giving Birth by Time t* (right panel) by *Risk Group* according to the estimated *Model 2*.

It is natural to think that women in the low-risk group start their reproductive history with a "wait and see" attitude and have higher expectations of the contraception related benefits than the other groups. Using the median (by risk group) of the individual

estimated values of the parameters w_0 and μ, we show (see Figure 4) that the median low-risk woman is characterized by a delay in her childbirth hazard, before catching up with those of the median woman in other groups with higher risk. This confirms a stylized fact related to the delay in childbirth hazard for low-risk groups that has been reported by other scholars (see Aalen et al., 2008, p. 414). As a consequence, the probability of the onset of motherhood by any given time is clearly lower for those women who start their active sexual life with higher expectations of contraception related benefits, than for those women who start with lower initial values of the expected future benefits linked to contraception.
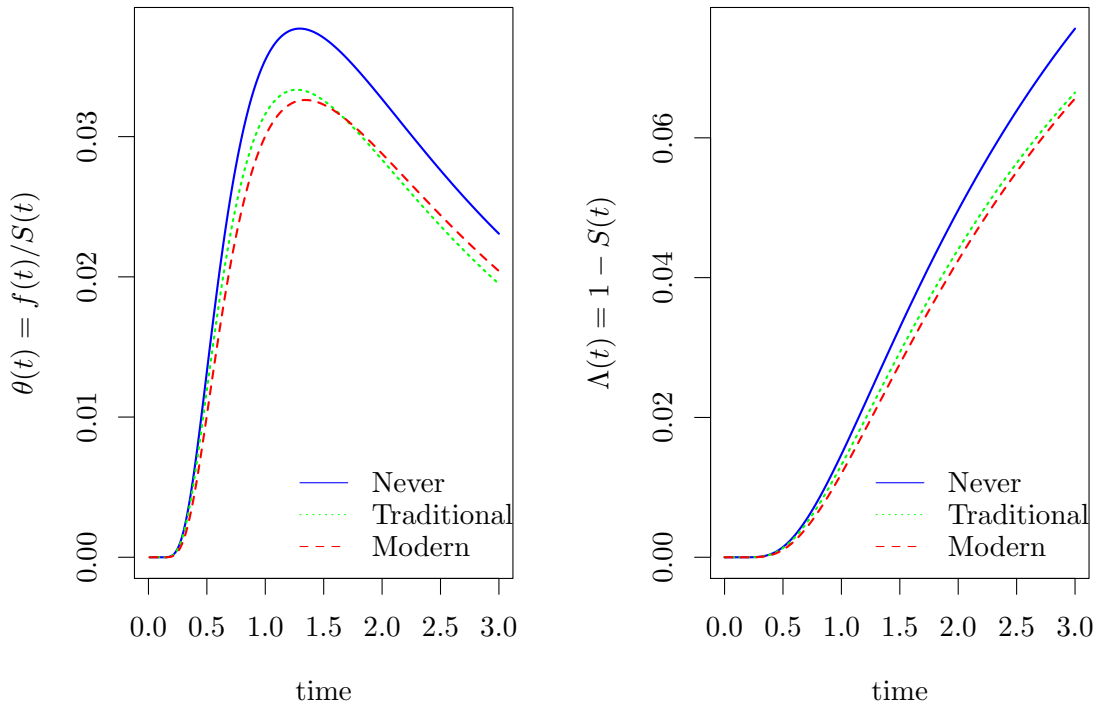

Figure 5: *Hazards* (left panel) and the corresponding *Probability of Giving Birth by Time t* (right panel) by *Method of Birth Control Used according to the estimated* Model 2 .

Moreover, if we analyze the median values of w_0 and μ by the most efficient contraception method ever used, we find that the median woman who has never used contraception and the one who has only used traditional contraception have failry similar initial values

of the expected contraception related benefits; the estimated value of $\hat{w}_0 = 2.03$, while those whose most efficient contraception method ever used is the modern method start their active sexual life at a relatively lower risk level ($\hat{w}_0 \approx 2.09$). However, the median woman who has never used contraception has the lowest drift towards zero, $\hat{\mu} = 0.47$, as compared to her counterpart who has only used the traditional method ($\hat{\mu} = 0.52$), or the one who used a modern method ($\hat{\mu} = 0.50$).

The combined effect of the estimated values of w_0 and μ for a median woman (by risk group), as expected, produces non-proportional hazards (see Figure 5). A median woman who has never used contraception has the highest risk of giving birth to the first child and, consequently, the highest probability of first child birth by any time compared to the one who uses any contraception. However, estimation results suggest that modern contraception methods are only slightly better than traditional methods in postponing the timing of the first child birth. There appears to be no difference between traditional and modern contraception methods with respect to risk and probability of first child birth. The probability of giving birth to the first child by any time for an average woman who uses only traditional contraception methods is just marginally higher than the probability of her counterpart who uses modern contraception methods. Also note that at approximately 18 months from first intercourse, the hazards for a median woman who has only used a traditional method of contraception and the one who has used a modern method cross each other, further suggesting non-proportionality.

6   Conclusion

In this paper, we presented a model for birth timing that was based upon the stochastic nature of the human reproductive process and focused on contraception decisions. An empirical counterpart, based on first hit times, was estimated using threshold regressions. It focussed on the duration to first childbirth using data from a high fertility country in Africa, where there is evidence that birth intervals and the fertility transition is rather different than in other parts of the world (see Romaniuk, 2011; Moultrie et al., 2012; Bongaarts and Casterline, 2013). Our empirical results suggest a negative effect of contraception efficiency on the duration from first intercourse to first childbirth.

As expected, the use of higher efficiency modern contraceptive methods result in the postponement of the onset of motherhood.

The aim of this paper was to develop a theoretical model consistent with that observation and present empirical evidence on that negative relationship. The question is of importance, because optimal birth timing, and ultimately optimal family size, is achieved through the practice of birth control. Our model and findings can be used as an additional building block in explaning the puzzling negative relationship between income and family size.

Our analysis did not explicitly account for a quality-quantity trade-off that women might take into account in their optimal family size decision calculus, which suggests a direction for future research. In particular, one could allow child related benefits in the model to depend on parity or on the expected costs of child quality. Furthermore, our analysis does not consider higher parity birth intervals, primilary because of data limitations. Extending the model to account for these additional considerations could provide further insights into family formation, and the fertility transition.

References

Aalen, O. O., Borgan, O. and Gjessing, H. K. (2008), *Survival and event history analysis: A process point view*, Springer.

Aalen, O. O. and Gjessing, H. K. (2001), 'Understanding the shape of the hazard rate: A process point view', *Statistical Sciences* **16**(1), 1–22.

Aalen, O. O. and Gjessing, H. K. (2004), 'Survival models based on the ornstein-uhlenbeck process', *Lifetime Data Analysis* **10**, 407–423.

Abbring, J. H. (2012), 'Mixed hitting-time models', *Econometrica* **80**(2), 783–819.

Becker, G. S. (1960), An economic analysis of fertility, *in* 'Demographics and Economic Change in Developed Countries', Princeton: National Bureau of Economic Research, pp. 209–231.

Becker, G. S. (1965), 'A theory of the allocation of time', *The Economic Journal* **75**(299), 493–517.

Becker, G. S. and Lewis, H. G. (1973), 'On the interaction between the quantity and quality of children', *Journal of Political Economy* **81**(2), S279–S288.

Bloom, D. E., Canning, D., Fink, G. and Finlay, J. E. (2009), 'Fertility, female labor force participation, and the demographic dividend', *Journal of Economic Growth* **14**, 79–101.

Bongaarts, J. and Casterline, J. (2013), 'Fertility transition: Is sub-Saharan Africa different?', *Population and Development Review* **38**(s1), 153–168.

Caldwell, J. C. (1976), 'Toward the restatement of demographic transition theory', *Population and Development Review* **2**(3/4), 321–366.

Chhikara, R. S. and Folks, L. (1989), *The Inverse Gaussian Distribution: Theory, Methodology, and Applications*, Dekker.

Fitzenberger, B., Sommerfeld, K. and Steffes, S. (2013), 'Causal effects on employment after first birth - A dynamic treatment approach', *Labour Economics* **25**, 49–62.

Galor, O. and Zang, H. (1997), 'Fertility, income distribution, and economic growth: Theory and cross-country evidence', *Japan and the World Economy* **9**(2), 197–229.

Gawarecki, L. and Mandrekar, V. (2011), *Stochastic Differential Equations in Infinite Dimensions: with Applications to Stochastic Partial Differential Equations*, Springer, Berlin.

Goldin, C. and Katz, L. F. (2002), 'The power of the pill: Oral contraceptives and women's career and marriage decisions', *Journal of Political Economy* **110**(4), 730–770.

Heckman, J. J. and Vytlacil, E. J. (2007), *Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation*, Vol. 6B, Elsevier, chapter 70, pp. 4779–4874.

Heckman, J. J. and Willis, R. J. (1976), Estimation of a stochastic model of reproduction: An econometric approach, *in* N. E. Terleckyj, ed., 'Household Production and Consumption', Colummbia University Press (for NBER), pp. 99–146.

Jones, L. E. and Tertilt, M. (2008), An economic history of fertility in the United States: 1826-1960, *in* P. Rupert, ed., 'Frontiers of Family Economics', Vol. 1, Emerald, pp. 165–230.

Klepinger, D. H., Lundberg, S. and Plotnick, R. D. (1995), 'Adolescent fertility and the educational attainment of young women', *Family Planning Perspectives* **27**(1), 23–28.

Lancaster, T. (1972), 'A stochastic model for the duration of a strike', *Journal of the Royal Statistical Society. Series A* **135**(2), 257–271.

Lee, M.-L. T. and Whitmore, G. (2006), 'Threshold regression for survival analysis: Modeling event times by a stochastic process reaching a boundary', *Statistical Sciences* **21**(2), 501–513.

Lee, M.-L. T. and Whitmore, G. (2010), 'Proportional hazards and threshold regression: their theoretical and practical connections', *Lifetime Data Analysis* **16**, 196–214.

Leibenstein, H. (1975), 'The economic theory of fertility', *The Quarterly Journal of Economics* **89**(1), 1–31.

Loewenstein, G. and Prelec, D. (1991), 'Negative time preference', *The American Economic Review* **81**(2), 347–352.

Mincer, J. and Polachek, S. (1974), 'Family investments in human capital: Earnings of women', *The Journal of Political Economy* **82**(2), 76–108.

Moultrie, T. A., Sayi, T. S. and Timaeus, I. M. (2012), 'Birth intervals, postponement, and fertility decline in Africa: A new type of transition?', *Population Studies: A Journal of Demography* **66**(3), 241–258.

Ramsey, F. P. (1928), 'A mathematical theory of saving', *The Economic Journal* **38**, 543–559.

Romaniuk, A. (2011), 'Persistence of high fertility in tropical Africa: The case of the democratic republic of the congo', *Population and Development Review* **37**(1), 1–28.

Sannikov, Y. (2008), 'A continuous-time version of the principal-agent problem', *The Review of Economics Studies* **75**(3), 957–984.

Shimer, R. (2008), 'The probability of finding a job', *American Economic Review* **98**(2), 268–273.

Stokey, N. L. (2009), *The economics of inaction: Stochastic control models with fixed costs*, Princeton University Press.

te Velde, E., Habbema, D., Leridon, H. and Eijkemans, M. (2012), 'The effect of postponement of first motherhood on permanent involuntary childlessness and total fertility rate in six European countries since the 1970's', *Human Reproduction* **27**(4), 1179–1183.

Upchurch, D. M. and McCarthy, J. (1990), 'The timing of a first birth and high school completion', *American Sociological Review* **55**(2), 224–234.

Wenocur, M. L. (1990), 'Ornstein-uhlenbeck process with quadratic Killing', *Journal of Applied Probability* **27**(3), 707–712.

Whitmore, G. A. (1979), 'An inverse gaussian model for labour turnover', *Journal of the Royal Statistical Society. Series A (General)* **142**(4), 468–478.

Woodbury, M. A. and Manton, K. G. (1977), 'A random-walk model of human mortality and aging', *Theoretical Population Biology* **11**, 37–58.

Xiao, T., Whitmore, G. A. and He, X. (2012), 'Threshold regression for time-to-event analysis: The sttthreg package', *The Stata Journal* **12**(2), 257–283.

Yashin, A. (1984), *Dynamics in Survival Analysis: Conditional Gaussian Property Versus Cameron-Martin Formula*, Working paper (International Institute for Applied Systems Analysis), International Institute for Applied Systems Analysis.
URL: *https://books.google.co.za/books?id=IRVkHQAACAAJ*

Yashin, A. I. and Manton, K. G. (1997), 'Effects of unobserved and partially observed covariate processes on system failure: A review of models and estimation strategies', *Statistical Science* **12**(1), 20–34.

Yeakey, M. P., Muntifering, C. J., Ramachandran, D. V., Myint, Y., Creanga, A. A., and Tsui, A. O. (2009), 'How contraceptive use affects birth intervals: Results of a literature review', *Studies in Family Planning* **40**(3), 205–214.