Routledge
Taylor & Francis Group

Population Studies
A Journal of Demography

ISSN: 0032-4728 (Print) 1477-4747 (Online) Journal homepage: https://www.tandfonline.com/loi/rpst20

Worldwide fertility declines do not rely on stopping at ideal parities

Daniel J. Hruschka, Rebecca Sear, Joseph Hackman & Alexandria Drake

To cite this article: Daniel J. Hruschka, Rebecca Sear, Joseph Hackman & Alexandria Drake (2019) Worldwide fertility declines do not rely on stopping at ideal parities, Population Studies, 73:1, 1-17, DOI: 10.1080/00324728.2018.1513164

To link to this article: https://doi.org/10.1080/00324728.2018.1513164

Population Studies, 2019
Vol. 73, No. 1, 1–17, https://doi.org/10.1080/00324728.2018.1513164

Worldwide fertility declines do not rely on stopping at ideal parities

Daniel J. Hruschka[1], Rebecca Sear[2], Joseph Hackman[1] and Alexandria Drake[1]
[1]Arizona State University, [2]London School of Hygiene and Tropical Medicine

A key demographic hypothesis has been that fertility declines rely on stopping at target parities, but emerging evidence suggests that women frequently reduce fertility without specific numeric targets. To assess the relative importance of these two paths to fertility decline, we develop a novel mixture model to estimate: (1) the proportion of women who stop at a target parity; and (2) mean completed fertility among those who do not. Applied to Demographic and Health Survey data from women aged 45–49 in 84 low- and middle-income countries, and to United States Census cohorts, the model shows considerable variation in the proportion stopping at specific parities (1–84 per cent). The estimates also show that declines in completed fertility are largely attributable to women who do not stop at target parities, suggesting that stopping at ideal parities may be less important than parity-independent decisions for a wide range of fertility transitions.

Supplementary material for this article is available at: https://doi.org/10.1080/00324728.2018.1513164

Keywords: fertility transition; demographic transition; mixture model; parity-dependent; Gamma–Poisson; parity-independent

[Submitted December 2016; Final version accepted July 2018]

Introduction

In the last few centuries, human populations worldwide have witnessed dramatic reductions in fertility. While this decline is well established empirically, there is still considerable debate about the specific cultural, social, and psychological factors driving the change. A key point of controversy surrounds the role played by women or couples aiming for specific ideal parities. In an influential formulation of preconditions for the fertility transition, the economist and demographer A. J. Coale (1973) proposed that one of the requirements for fertility decline is a cognitive shift whereby reproduction moves into the realm of 'conscious calculation'. Demographers have often interpreted this shift as a cognitive change towards specific numeric goals or targets (e.g., two children) accompanied by limiting fertility after an ideal completed fertility has been achieved (Henry 1961; Knodel 1983; Van de Walle 1992; Reher and Sanz-Gimeno 2007). According to this perspective, the major force behind fertility decline is an increasing focus on achieving one of a limited number of low fertility targets (e.g., two, three, or four children). Indeed, a long-stated assumption in the demographic literature is that controlling fertility through parity-specific stopping has a greater impact on fertility rates than other means of regulating reproduction, such as spacing or postponement (Van de Walle 1992; Westoff and Bankole 2000; Cleland et al. 2006; Van Lith et al. 2013).

Numerous studies of historical Europe have been interpreted to support this view (Knodel 1979, 1987; Hionidou 1998; Reher and Sanz-Gimeno 2007). However, in the last two decades, converging lines of evidence have challenged the primacy of parity-specific stopping in driving down fertility. Ethnographic research, based largely in sub-Saharan Africa, suggests that a great deal of strategic decision-making focuses on maintaining appropriate timing of births and ensuring future reproduction is possible, rather than on the absolute number of children (Mason 1997; Bledsoe et al. 1998; Bledsoe and Banja 2002; Johnson-Hanks 2002, 2007). If this is the case, the existence of non-numeric responses to ideal family size (e.g., 'Only God knows' or 'Don't

know') in many low- and middle-income countries may reflect not a lack of concern about regulating fertility, but rather a lack of concern about specific numeric targets (Knodel and Van de Walle 1979; Hayford and Agadjanian 2011). In such situations, we might expect decisions about other aspects of reproduction—such as spacing, postponement, and starting later—to play a greater role in fertility decline than stopping based on specific achieved parities (Moultrie et al. 2012; Timæus and Moultrie 2013; Casterline and Odden 2016a, 2016b).

As Van Bavel notes, quantitative demography has traditionally focused on parity-specific stopping, in part due to the early development of analytical tools for assessing parity-specific stopping (Coale and Trussell 1974; David et al. 1988; Van Bavel 2004b). However, in the last two decades, quantitative demographers have also developed a number of tools to demonstrate that fertility declines in a number of specific settings have involved a combination of parity-specific stopping and increases in spacing and postponement (Anderton and Bean 1985; Knodel 1987; Feng et al. 1995; Hionidou 1998; Szreter and Garrett 2000; Van Bavel 2004a; Timæus and Moultrie 2008, 2013; Van Bavel and Kok 2010). Changes in the timing of births can arise for a number of reasons, including spacing births (focused on the interval between one birth and the next), starting reproduction later, stopping for reasons other than current parity, and postponing for reasons unrelated to age of the last child or total number of children (Timæus and Moultrie 2013; Casterline and Odden 2016b; Towner et al. 2016; Mattison et al. 2018).

Current approaches to estimating parity-independent strategies, however, have two significant limitations. First, most current models demand data that are not always available across a wide range of settings, making it difficult to generalize or to examine broad ecological influences on these fertility patterns. Although a number of ethnographic, historical, and quantitative case studies have demonstrated the potential importance of fertility control without ideal family sizes (Knodel 1987; Van Bavel 2004a), it is not clear how widespread the phenomenon is or how important it is for recent fertility declines compared with parity-specific stopping. Methods that can be used across a broad range of data sets are crucial for direct comparisons of, for example, how the recent (and relatively fast) fertility transitions in Asia and Africa might differ from the slower transitions observed in historical Europe. Second, current methods estimate the degree of parity-specific stopping as well as parity-independent spacing and postponement for entire populations. However, it is likely that any population consists of a mixture of some women who have stopped at a specific target and others who are regulating fertility based on other criteria. If this is the case, estimates from current methods potentially confound two very different kinds of decision-making. Thus, to model gradual change within populations, as seen in many fertility transitions, we need to track the changing mixtures of strategies.

To estimate the importance of an increasing focus on numeric targets relative to other strategies for regulating reproduction in modern fertility declines, we analyse completed fertility data from two sources: (1) 301 Demographic and Health Surveys (DHS) taken between 1985 and 2016 across 84 low- and middle-income countries at varying stages of fertility transitions (ICF International 1985–2016); and (2) twentieth-century historical census data from the United States (US) (Ruggles et al. 2015). The DHS data sets provide global, but temporally shallow, data on recent variation in fertility outcomes. These represent populations in Europe and Central Asia ($n = 20$), South Asia ($n = 23$), sub-Saharan Africa ($n = 163$), Middle East and North Africa ($n = 21$), Southeast Asia ($n = 22$), and Latin America and the Caribbean ($n = 52$). Most DHS data sets also include detailed information on the timing of births for each woman. This permits comparison of our population-based model estimates with micro-level data on birth timing. The US data complement the temporally shallow DHS data sets by illustrating how the model detects changes over long-term historical time without the benefit of such micro-level data.

The analyses presented here rely on a novel extension of zero-inflated regressions to decompose sample distributions of completed fertility into: (1) the proportion of the population that ends fertility at specific low numeric targets; and (2) lifetime fertility (i.e., mean completed fertility) among the remaining proportion that does not stop at a specific numeric target. The model is an effort to estimate the approximate proportion of women completing lifetime reproduction in one of these two ways. It is very possible that, during her lifetime, a woman may move between the two approaches. For example, a woman who ultimately stops at a desired target parity may have followed any number of strategies to achieve that parity, including postponing and spacing births based on current parity. Or a woman who starts with a specific parity in mind may change to another preferred parity later in life or move away from a parity preference

altogether. Finally, women in either of these broad categories may use a wide range of strategies to accomplish their goals during the course of their reproductive lives (Kingsley 1963). The crucial distinction between parity-specific stopping and parity-independent decisions is that parity-specific stopping is defined by ending lifetime reproduction at a specific number of children based on a desire for that number of children. The model described here aims to estimate the relative mixture of distributions generated by these two kinds of decision-making.

In the following sections we first outline our model, with some examples, and then explain the method in detail. In the 'Results' section we: (1) fit the model to 313 empirical distributions from the DHS and the US Census; (2) check how well our population-level model performs against micro-level data, by assessing whether our estimate of parity-specific stopping correlates with two indicators from birth records across 301 DHS data sets; (3) examine how variation in our two key estimates—the proportion of parity-specific stoppers and mean completed fertility among those following a parity-independent approach—are associated with fertility decline and other variables associated with fertility decline, such as education and wealth; and (4) compare the relative importance of these two strategies for fertility decline across different world regions.

Outline description of the model

The model we use here assumes, to a first approximation, that individuals can follow one of two approaches to lifetime reproduction, which we refer to as parity-independent decision-making and parity-specific stopping. The first approach relies on behavioural management and physiological constraints on the timing of births, rather than on achievement of a specific ideal number of children. The pace and timing of births arise from a number of strategies, including delaying first birth, extending reproduction based on the age of the youngest child, postponing for other reasons, and even stopping (provided this does not depend on achieving a specific parity). Thus, there are multiple variants of this approach, including the use of both spacing and postponement (Timæus and Moultrie 2008, 2013; Casterline and Odden 2016b). Its unifying feature is that a woman following this approach disregards current parity in decisions to end reproduction, a mode of decision-making that has been observed in a number of ethnographic and historical studies (Bledsoe and Banja 2002; Johnson-Hanks 2002; Van Bavel 2004a; Timæus and Moultrie 2008). Parity-independent fertility can include substantial variation *within* individuals, such as increasing birth intervals with age. It can also include substantial variation *between* individuals, such as variation in average interbirth intervals and varying windows of opportunity for reproduction (as defined by the onsets of puberty and menopause). Importantly, many of these diverse sources of variation in pace, timing, and postponement will lead to completed fertility distributions approximating a mixture of Poisson distributions (Winkelmann 1995). As an example, the model does not require a constant rate of reproduction across an individual woman's life, since a population with a time-varying rate of reproduction (e.g., increasing birth intervals with age) can still produce a Poisson mixture distribution of completed fertility (Winkelmann 1995). Moreover, a population with variation in both windows of opportunity for reproduction and average interbirth intervals will also produce a mixture of Poisson distributions with varying lifetime fertility. Thus, many kinds of parity-independent strategies can result in distributions that resemble a mixture of Poisson distributions.

We model this mixture of Poisson processes as a Gamma–Poisson (or negative binomial) model (Brass 1958; Wood 1994). Figure 1(a) illustrates a sample from Bangladesh in 1996 (women aged 45–49) for which the predicted distribution from such a process is a close approximation to the observed distribution of completed fertility. An important point about such distributions is that they can be summarized completely by two parameters: expected lifetime fertility and the variance in lifetime fertility between women in the population. Expected lifetime fertility is the number of children a woman is expected to bear over her lifetime. However, under a count process (e.g., a Poisson process) the realized completed fertility for a specific woman will not perfectly reflect this expected value because of stochastic variation introduced by the count process (Hruschka and Burger 2016). Thus, the expected lifetime fertility estimated for a Poisson process reflects a summary measure of a stochastic process and not some 'master schedule' of planned interval lengths (Bongaarts and Potter 2013). It is also important to emphasize that the estimation of expected lifetime fertility for a population does not assume that women's fertility is constant over the life course. For example, if age-specific fertility rates decline over the life course, estimated lifetime fertility will represent an average of these. Notably, as long as

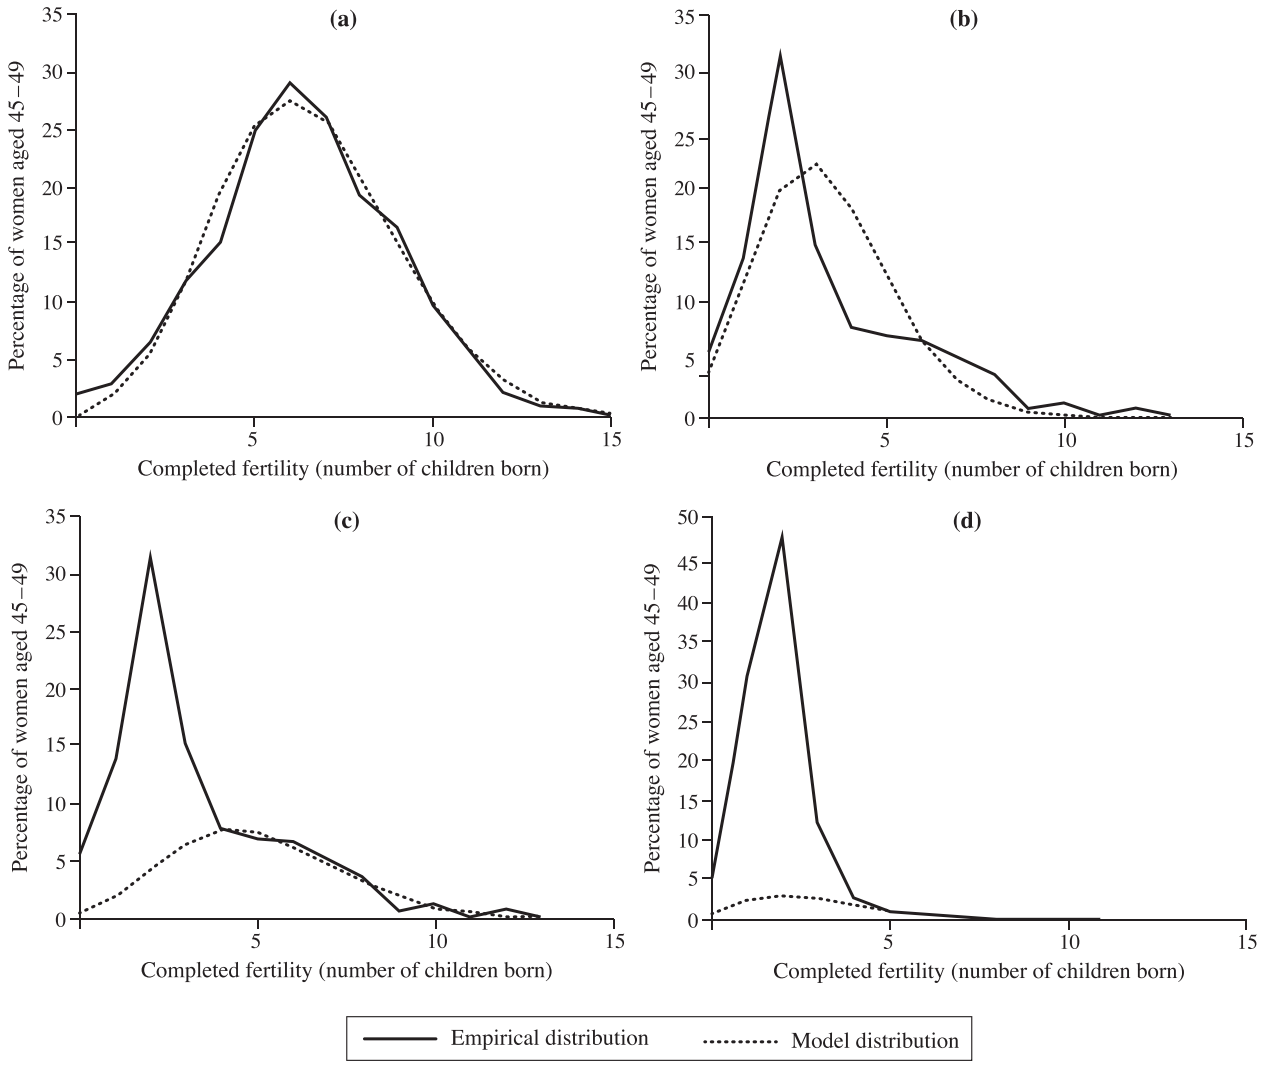

Figure 1 Empirical completed fertility distributions compared with Gamma–Poisson portion of best fitting model distributions (women aged 45–49) for (a) Bangladesh 1996, (b) and (c) Kazakhstan 1995, and (d) Ukraine 2007

Note: (a) Bangladesh 1996 DHS vs. pure Gamma–Poisson ($n = 649$); (b) Kazakhstan 1995 DHS vs. pure Gamma–Poisson ($n = 370$); (c) Kazakhstan 1995 DHS vs. Gamma–Poisson permitting parity-specific stopping at 0–5 children ($n = 370$); (d) Ukraine 2007 DHS vs. Gamma–Poisson permitting parity-specific stopping at 0–5 children ($n = 1{,}082$).

Source: Authors' calculations from DHS.

women are not stopping at specific parities, continuously declining age-specific fertility rates will still produce a Poisson mixture distribution of the kind used in this model (Winkelmann 1995).

The second approach, parity-specific stopping, targets a specific set of ideal parities and then stops or severely limits reproduction when such parities are reached (Henry 1961; Coale 1973). A woman or couple might have several possible ideal parities (e.g., two or three children) that may vary over the life course, with varying weights on how important each is (Coombs 1974). However, women or couples following the parity-specific approach will tend to stop reproduction within that small set of ideal parities. To achieve the desired parity, a woman may follow any number of strategies, including changing the timing of births. However, this need not follow a 'master schedule' with pre-planned intervals to achieve a specific parity (Bongaarts and Potter 2013). The crucial distinction between parity-specific stopping and parity-independent fertility, is that in the former there is a clear stopping rule according to number of children. We assume here that when people set specific ideal parities, they are generally in the low range (i.e., zero to five children), with some probability assigned to each of these outcomes. As we show later, estimated parity-specific stopping is concentrated at parities one to three (and to a

lesser extent four) across all samples. Nonetheless, five is still included as a potential target in the models to provide a check on the need for parities above four as potential targets. Ultimately, model estimates indicate that five is almost never a target for parity-specific stopping. We also treat nulliparity as a special case of parity-specific stopping, as it can reflect either lifetime infertility or a choice to remain nulliparous.

An example of a situation where a substantial portion of women are likely to be following parity-specific stopping is the Kazakhstan 1995 DHS sample (Figures 1(b) and (c)). A pure Gamma–Poisson model does not closely approximate the empirical distribution, because the distribution generated by the Gamma–Poisson process is not sufficiently peaked at low parities (Figure 1(b)). When we permit some proportion of the women to engage in parity-specific stopping at parities zero to five, we can then also estimate the Gamma–Poisson distribution of fertility among the remaining set of women following a parity-independent approach (dotted line in Figure 1(c)). Importantly, in this sample, some proportion of the women arrive at low parities via parity-specific stopping (the proportion over the dotted line in Figure 1(c) for parities zero to five), while others arrive at low parities via a parity-independent approach (the proportion under the dotted line in Figure 1(c)). At parities above five, the Gamma–Poisson model applied to the subset of women following a parity-independent approach provides an excellent fit to the empirical distribution. Figure 1(d) illustrates a situation where nearly all women stop at specific parities (Ukraine, 2007). Nearly all cases are between one and three, and there are so few cases above five that a parity-independent approach is very unlikely to have generated much of this distribution of completed fertilities. See the supplementary material (Figures S4 and S5) for comparable figures for all samples in this study.

According to the model, the observed lifetime fertility for any one person may arise from either parity-specific stopping or stochastic events during a Poisson-like process of reproduction. This means we need to model the distribution of completed fertilities as a statistical mixture of the two approaches. Using an extension of classical zero-inflated regression models, we describe a method to estimate from a distribution of completed fertility: (1) the proportion of women following a parity-independent vs. a parity-specific stopping approach in a given sample; and (2) mean lifetime fertility among women following a parity-independent approach.

Materials and methods

Samples

We use data from two sources: DHS 1985–2016, to examine diversity across contemporary international populations (ICF International 1985–2016), and historical US Censuses 1940–90, to examine long-term change across cohorts (Ruggles et al. 2015). First, we analyse 301 DHS across 84 low- and middle-income countries that have collected systematic, comparable data on women's reproductive histories (Arnold 1990). We focus on women for two practical reasons: (1) the DHS Program has primarily aimed its data collection efforts at women and children, with relatively few having collected comparable data for men; and (2) even for DHS data sets with information about men's marital fertility, varying opportunities for extra-marital fertility make it impossible to estimate men's fertility in the same way as women's. We use DHS data sets collected between 1985 and 2016 that include information about fertility and household wealth. To assess the sensitivity of our findings to alternative samples of women, we also analyse and report model estimates for ever-married women and living children.

Second, we analyse data from unweighted subsamples of six consecutive US Censuses (1940, 1950, 1960, and 1970 1 per cent flat sample; 1980 5 per cent; and 1990 unweighted 1 per cent) archived by the Integrated Public Use Microdata Series at the Minnesota Population Data Center (Ruggles et al. 2015). To maintain comparability across censuses that sometimes only collected fertility for ever-married women, we exclude never-married women. We also exclude women with imputed values for number of children ever born or marital status. To control for changing ethnicity distributions in the US over the twentieth century, we focus on white, non-Hispanic women (Stulp et al. 2016). To avoid slow estimation with very large sample sizes, for samples larger than 5,000, we use a random subsample of size 5,000 for estimation of each sample. The list of samples accompanied by key indicators and model estimates are provided in the supplementary material (Table S2).

In analyses of all samples, we estimate completed fertility for women aged 45–49 using total live-born children, as this measure is uniformly available across all censuses and surveys (Eijkemans et al. 2014). In the most comprehensive study to date of six natural fertility populations in Europe and North America, 90 per cent of women ended reproduction by age 45 and nearly 100 per cent by age

50. Thus, focusing on 45–49-year-olds will capture most live births. For the US Census, to fill in the time series, we also independently analyse cohorts of older samples of women (aged 50–54) for a total of twelve independent samples (six censuses with two age categories each).

Model specification

In line with past work on fertility, the modelling approach aims to infer general patterns of decision-making in a population from aggregate distributions of fertility outcomes (Coale and Trussell 1974; Knodel 1977; David et al. 1988; Van Bavel 2004b; Timæus and Moultrie 2008). Compared with these approaches, the current approach has the advantage of requiring only the distribution of completed fertility in a population. Moreover, it provides a straightforward way of estimating how key parameters (e.g., proportion following a parity-independent approach and mean lifetime fertility among those) depend on different factors (e.g., education, wealth, rural residence).

The model is related to multiple hurdle models of fertility (Miranda 2010). However, it does not assume, as multiple hurdle models do, that all women who arrive at a specific final parity get there via a single strategy. Rather, a woman with a final parity of two may have reached it either through parity-specific stopping or through a parity-independent Poisson process. Thus, the model is more closely aligned with a zero-inflated model, but permits inflation at other parities as well (Poston and McKibben 2003). We treat nulliparous stopping as a special case of parity-specific stopping, as it can potentially reflect either a decision to stop at zero or lifelong infertility. Thus, according to the model, there is some probability, q_0, that a woman will be nulliparous (for either of these two reasons). There is an additional probability, q_{PS}, that she will complete childbearing at parities one to five due to parity-specific stopping. Of these parity-specific stoppers, the probability of stopping at parity y is q_y. Although we will show that a target of five children is generally never estimated, we include it to assess the degree to which parity-specific stopping may extend to such high parities. Where parity-specific stopping only occurs at lower parities, the estimates for q_5, and possibly q_4, should statistically be zero.

Finally, there is the complementary probability, $1 - q_0 - q_{PS}$, that a woman is following a parity-independent strategy. Among those following a parity-independent approach, the probability of parity y is a gamma mixture of Poisson processes with mean lifetime fertility, λ, and variance in lifetime fertility, λ^2/θ.

Here, we label the probability of achieving parity y through the parity-independent process as $PDF(y|\lambda, \theta)$. The full probability can be written as:

$$p(y) = \begin{cases} (1 - q_0 - q_{PS}) \times PDF(0|\lambda, \theta) + q_0 & \text{if } y = 0 \\ (1 - q_0 - q_{PS}) \times PDF(y|\lambda, \theta) + q_{PS}q_y & \text{if } y = 1,\ 2,\ 3,\ 4,\ 5 \\ (1 - q_0 - q_{PS}) \times PDF(y|\lambda, \theta) & \text{if } y > 5 \text{ and } < 20 \\ (1 - q_0 - q_{PS}) \times PDF(y \geq 20|\lambda, \theta) & \text{if } y \geq 20 \end{cases}. \quad (1)$$

Given that some data sets from the US Census have a single category of twelve or more births, in the US samples we model a maximum of 12+ births instead of 20+ as a single category. We assume weakly informative prior distributions for the two Gamma–Poisson parameters as $\lambda \sim \text{gamma}(3,\ 0.5)$, which has an expected value of six children and a wide variance around that, and $\theta \sim \text{gamma}(0.001,\ 0.001)$, a commonly used weakly informative prior for the shape parameter for negative binomial distributions.

The probabilities of excess nulliparity, q_0, above that expected from a Gamma–Poisson process and of excess parities due to parity-specific stopping, q_{PS}, are parameterized with binomial logistic models:

$$q_0 = \frac{\exp(\beta_0)}{1 + \exp(\beta_0)} \quad \text{and} \quad q_{PL} = \frac{\exp(\beta_{PS})}{1 + \exp(\beta_{PS})}. \quad (2)$$

The probabilities of stopping at specific parities, q_1, q_2, q_3, q_4, and q_5, are modelled as a multinomial logistic with parity one as the reference category:

$$q_i = \frac{\exp(\beta_i)}{\exp(\beta_1) + \exp(\beta_2) + \exp(\beta_3) + \exp(\beta_4) + \exp(\beta_5)} \quad i = 1, 2, 3, 4, 5 \quad \beta_1 = 0. \quad (3)$$

We assume uniform priors for $\beta_0,\ \beta_{PL}, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5 \sim \textit{logistic}(0, 1)$ that are symmetric for $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5$ by constraining $\beta_1 + \beta_2 + \beta_3 + \beta_4 + \beta_5 = 0$.

Model estimation

We calculate parameter point estimates and credibility intervals from the estimated posterior probability distribution for the model. To estimate the posterior probability distribution, we use a Metropolis–Hastings Markov Chain Monte Carlo (MCMC) process implemented in R (code available from authors). Depending on the convergence of estimates, we use a 'short' or 'long' estimation process. The first phase of the process tunes the proposal distribution with 40 tuning loops of 1,000 iterations each for 'short' runs (500 tuning loops for 'long' runs). After each loop, the proposal distribution's scale is adjusted based on the acceptance rate from the prior 1,000 iterations, and the proposal distribution's covariance structure is estimated from the previous 1,000 iterations. With the final proposal distribution, we then allow 20,000 steps for burn-in of 'short' runs (100,000 for 'long' runs) and, finally, estimate the posterior distribution with 200,000 steps (1,000,000 for 'long' runs).

To ensure the Markov chain is well mixed, we assess stationarity of the Markov chain for each of the eight parameters by using a Geweke test comparing the last 50 per cent of iterations with the first 10 per cent (Geweke 1991). Within each chain, we use a Bonferroni-corrected alpha to reject a null hypothesis of no difference (alpha = 0.5/8 for eight parameters). We also assess whether the effective sample size for the run was less than 100 (which would indicate the estimates to be relatively unreliable) using the 'coda' package in R (Plummer et al. 2006). We apply 'short' runs to the DHS data sets but, given high rates of non-convergence in the US Census samples (>15 per cent) with the 'short' runs, we apply the 'long' runs to all US Census samples. In those rare cases where chain diagnostics indicate a problem on the first run (three for DHS, none for US), we re-estimate with a 'long' run. In no cases did we need to rerun the analyses more than twice.

To assess the reliability of the estimation process, we also run duplicate MCMC chains for all samples and compare duplicates. Key parameter estimates (e.g., q_{PS},λ, θ) are highly correlated across the two runs ($r > 0.95$). In this paper we calculate point estimates for each parameter as the mean of the estimated posterior distribution for that parameter, and the 95 per cent credibility interval as the 2.75 and 97.5 per cent percentiles of the posterior distribution.

Fitting the model to survey data

We fit the model to the completed fertility distribution in each of the 301 DHS surveys and twelve US Census subsamples to estimate: (1) the excess proportion of nulliparous women (q_0); (2) the proportion of women engaged in parity-specific stopping (at parities one to five) (q_{PS}); and (3) the mean (λ) and variance of lifetime fertility (λ^2/θ) among women not following parity-specific stopping. Among those women estimated to have stopped at specific target parities, we also report estimates of the proportions stopping at one, two, three, four, and five children. For each of the 313 samples, we test whether the sample's empirical distribution was different from the fitted model distribution using a chi-square test. The reference distribution for the chi-square statistic is generated with 50,000 replicate samples drawn from the model distribution (using the Holm–Bonferroni method with alpha = 0.05, 313 tests) (R Core Team 2014). For those samples whose predictions significantly deviate from observations, we report the average deviation of predictions from observations. To assess whether standard pure Gamma–Poisson models would also create comparable fits, we estimate the fit of an uninflated Gamma–Poisson model to the same 313 samples using the same chi-square test.

Comparing estimates with micro-level data and assessing sensitivity to alternative specifications

We then assess how well the model estimates of parity-specific stopping correlate with two indicators from birth records estimated across 301 DHS. The first indicator is the concavity of parity progression ratios over increasing parity, which has been used as a measure of the degree of parity-specific limitation (Brass et al. 1997). The second indicator—variance in parity-specific median intervals—examines how far a population's birth intervals deviate from the Poisson model's assumption of a constant median birth interval across parities. Specifically, we use the variance in median interbirth interval length across the first five birth intervals (1–2, 2–3, 3–4, 4–5, and 5–6). As an example, a population with completely constant median birth intervals between births one and five would have a variance of zero, and any deviation from this uniform median birth interval would lead to increasing var-

iance. Comparisons with birth records exclude 32 surveys that did not collect birth records and two countries with only partial birth records (El Salvador, 1985 and Tanzania, 2012). We also assess the sensitivity of estimates when applying the model to alternative samples (e.g., ever-married women aged 45–49) and measures (living vs. live-born children) in the DHS.

Assessing associations of parity-specific stopping and parity-independent fertility with fertility declines

To assess the relationships of different strategies to variables traditionally associated with fertility declines, we estimate how key model outputs—the proportion who stop at specific target parities and completed fertility among those following a parity-independent approach—correlate with sample characteristics traditionally associated with fertility declines. These include the entire sample's completed fertility, the proportions of women with at least primary or at least secondary education, and mean household wealth (estimated according to Hruschka et al. 2015).

To assess the degree to which each of the two strategies contributes to fertility decline in the DHS samples, we examine pairs of surveys from countries that exhibited a decline in completed fertility of more than 0.5 between the first and final surveys. We then estimate the expected decline in fertility in each of two idealized scenarios:

(1) *Decline resulting purely from increasing parity-specific stopping.* Here we assume that the proportion of women engaged in parity-specific stopping ($q_{05} = q_0 + q_{PS}$) and their expected parity follows the observed fertility decline between surveys, but everything else stays the same.
(2) *Decline resulting purely from changes in completed fertility among women following parity-independent strategies.* Here we assume the mean completed fertility among the parity-independent segment follows the observed decline between surveys, but everything else stays the same.

To estimate these expected values, we decompose completed fertility as $q_{05}M_{05} + q_{PI}M_{PI}$, where q_{05} and q_{PI} are the proportions following parity-specific stopping at 0–5 children and a parity-independent approach, respectively, and M_{05} and M_{PI} are the corresponding average completed fertilities of these two populations.

We estimate the decline in scenario (1) as $q'_{05}M'_{05} + q'_{PI}M_{PI} - (q_{05}M_{05} + q_{PI}M_{PI})$, where q'_{05}, M'_{05}, and q'_{PI} are estimates from the second time point, and all other variables are estimates from the first time point. Here, $q'_{05}M'_{05} + q'_{PI}M_{PI}$ is the hypothetical completed fertility at the second survey if the only variables to change between surveys are the proportion of women following parity-specific stopping (vs. a parity-independent approach) and the distribution of ideal family sizes among women following parity-specific stopping.

Similarly, we estimate the decline in scenario (2) as $q_{05}M_{05} + q_{PI}M'_{PI} - (q_{05}M_{05} + q_{PI}M_{PI})$, where M'_{PI} is the average completed fertility of those following a parity-independent approach from the second time point, and all other variables are estimates from the first time point. Thus, $q_{05}M_{05} + q_{PI}M'_{PI}$ represents the hypothetical completed fertility at the second survey, if the only variable to change between surveys is the completed fertility of women following a parity-independent approach.

It is important to note that these two estimates of change—one based purely on changes in the proportion of women engaged in parity-specific stopping and one purely on the change in mean fertility among those engaged in parity-independent strategies—will not necessarily sum to the total observed change, as there is interaction of varying degrees between the changes in means and changes in proportions.

Finally, we assess whether fertility declines not relying on parity-specific stopping are particularly common in sub-Saharan Africa (Moultrie et al. 2012). Specifically, we estimate the proportion of observed declines attributable: (1) solely to parity-specific stopping; and (2) solely to declines in mean fertility among women following a parity-independent approach. We then compare the mean proportions for declines observed across five world regions (sub-Saharan Africa ($n = 18$ countries), South Asia ($n = 3$), Middle East, North Africa, and Central Asia ($n = 7$), Latin America and the Caribbean ($n = 7$), Southeast Asia ($n = 4$)) and for two eras in the US historical samples where we observe declines.

Results

Goodness of fit to empirical distributions

The model shows good fit to observed distributions, with only four of the 301 model distributions from

the DHS samples and two of the twelve model distributions from the US Census samples being significantly different from their respective empirical distributions (bootstrapped chi-square test, 100,000 replicates, Holm–Bonferroni correction for multiple tests with 313 tests). Importantly, the pure Gamma–Poisson model commonly used to model fertility distributions (which assumes no parity-specific stopping) is statistically different from 54.5 per cent of the 301 empirical samples from the DHS and 96.4 per cent of the twelve US Census samples. These results suggest that our proposed model describes the patterns of fertility accurately in nearly all samples and provides a much better fit in many cases than the pure Gamma–Poisson model.

There is little systematic deviation of predicted from observed proportions at each final parity, with the magnitude of the mean difference between predicted and observed parity proportions being less than 0.01. In those rare situations where the predictions and observations are significantly different ($n = 4$, DHS; $n = 2$, US), the model generally overestimated the proportions at parities five and six, and underestimated the proportions at parities eight, nine, and ten for DHS data sets, while for US Census data sets it underestimated nulli parity. However, the magnitude of deviations was never greater than 0.03.

It is not clear what causes the greater occurrence of small but significant deviations of predictions from observed distributions in the US samples. One hypothesis is that limiting parities to twelve in the US data provides less information for fitting the Gamma–Poisson component of the model. This might increase the uncertainty of estimates and thereby reduce fit in any specific estimation. A related hypothesis is that in situations where a larger portion of the population is estimated to be engaging in parity-specific stopping, there may be less information to provide accurate estimates of the Gamma–Poisson component. Future work should examine how small deviations arise from these and other aspects of data availability and model estimation.

What proportion of each sample follows parity-specific stopping (parities one to five)?

The proportion of women estimated to have stopped at specific parities varies from 0.7 per cent (Rwanda, 2000 and Burkina Faso, 2014) to 84 per cent (Ukraine, 2007 and Armenia, 2016). It increases as completed fertility decreases in DHS data sets ($r = -0.65$, $p < 0.001$), but not for US Census samples ($r = -0.26$, $p > 0.10$) (see the supplementary material (Table S2) for all estimates). Figure 2(a) shows the proportion of women in each sample who are estimated to follow parity-specific stopping and its relationship to the sample's mean completed fertility.

Notably, there is substantial variation in completed fertility, even among populations with very low

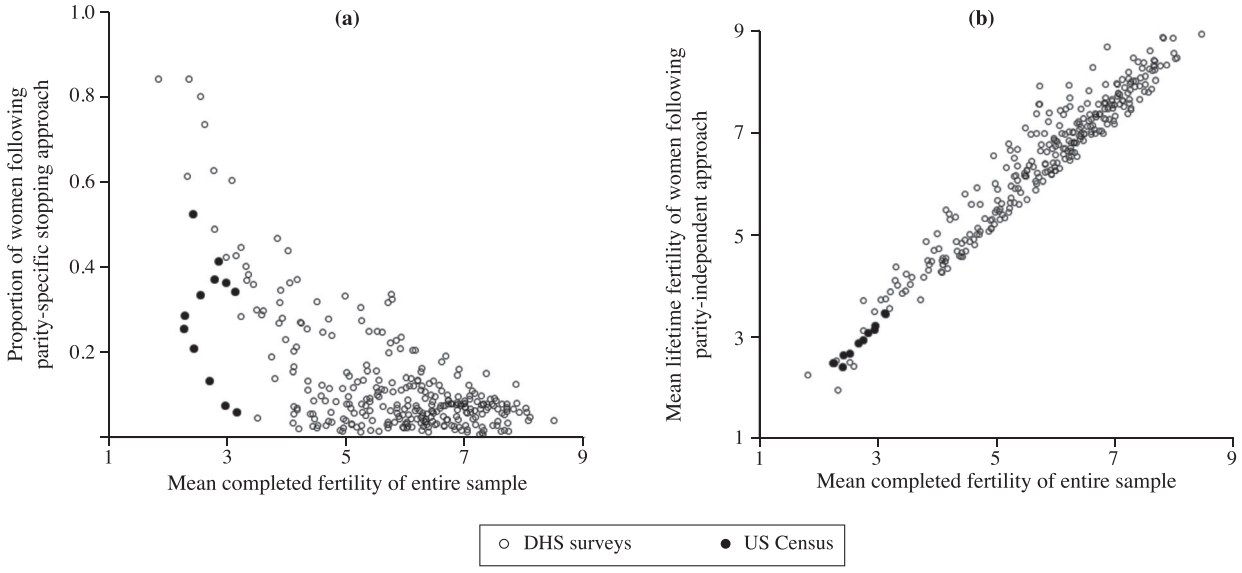

Figure 2 The relationship between mean completed fertility of entire samples and: (a) proportion of women following parity-specific stopping at 1–5 children; and (b) mean completed fertility among the complementary proportion following a parity-independent approach

Note: $n = 12$ US Census samples; $n = 301$ DHS samples.

Source: Authors' calculations from DHS and US Census data.

proportions of parity-specific stopping (Figure 2(a)). This suggests there is ample room for fertility decline without a shift to parity-specific stopping. For example, comparably low estimates for parity-specific stopping (<5 per cent of population) are seen in samples from Cambodia, South Africa, Bangladesh, Jordan, Rwanda, and Niger, yet mean completed fertility ranges from <4.5 in some samples (Cambodia, Myanmar, Vietnam, South Africa, Bangladesh, and Uzbekistan) to over 8.0 in others (Jordan, Rwanda, and Niger). Low mean fertility with little parity-specific stopping is especially characteristic of early US cohorts, with white, non-Hispanic women from the 1940 Census having completed fertility of 2.96 despite only 7.5 per cent being estimated to follow parity-specific stopping.

Conversely, there are samples with very different estimates for the proportion of women following parity-specific stopping, despite comparable overall completed fertility. For example, Turkey in 2004 and South Africa in 1998 exhibit similar completed fertility (4.0 and 4.2, respectively), but strikingly different estimated proportions following parity-specific stopping (44 and 2 per cent, respectively).

Among women following parity-specific stopping, there is clear bias towards parities zero and two in both the DHS and US Census samples (Figure 3). However, there are also non-trivial proportions stopping at parities one, three, and to a lesser extent four, although in the US the proportion stopping at one is also low. There is almost no excess of parity-specific stopping at five children. This suggests that the model is sufficiently flexible—permitting stopping at zero to five—to detect those parities at which parity-specific stopping is most heavily concentrated and to show which parities are unlikely to be a target for stopping (five and to a lesser extent four (and one in the US)).

What is the mean completed fertility of women following a parity-independent approach?

Mean lifetime fertility for women following a parity-independent approach ranges from 2.0 (Armenia, 2016) to 8.9 (Jordan, 1990). It is also negatively correlated with the proportion of parity-specific stoppers in DHS samples ($r = -0.54$, $p < 0.001$) but not US Census samples ($r = -0.32$, $p > 0.10$). Figure 2(b) shows the mean completed fertility of women following a parity-independent approach by the mean completed fertility of the entire sample. Given the low proportion of parity-specific stopping in many of the samples, the overall tight correlation is not surprising.

Associations of parity-specific stopping with micro-level data

Estimates of parity-specific stopping are strongly correlated with an established population-level measure of parity-specific limitation: the steepness of reductions in parity progression ratios (PPRs) at low parities (reduction in PPR2–3 compared with PPR1–2, $r = 0.85$; from PPR3–4 compared with PPR2–3, $r = 0.54$; all $p < 0.001$, $n = 267$). Estimates of parity-specific stopping are also strongly correlated with the magnitude of deviations from a Poisson assumption of equal median interbirth intervals ($r = 0.60$, $p < 0.001$, $n = 267$).

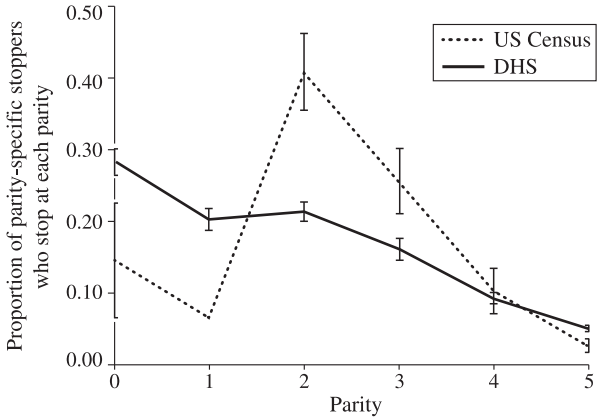

Figure 3 Proportions of parity-specific stoppers targeting parities zero, one, two, three, four, and five, US Census and DHS samples

Note: Error bars show 95 per cent confidence intervals.

Source: As for Figure 2.

How do these results change with alternative samples and measures?

Several checks with alternative samples (e.g., ever-married women aged 45–49) and measures (living vs. live-born children) in the DHS surveys suggest that the model produces robust results. Specifically, key estimates for ever-married DHS samples are strongly correlated with estimates from the full sample ($r = 0.99$ for parity-specific stopping, and both mean and variance of lifetime fertility for parity-independent strategies). As expected from the higher fertility of ever-married women, average lifetime fertility is higher among this group than among all women (0.2-child difference for populations at mean parity two, 0.1-child difference at mean parity four, and 0.05-child difference at mean parity six). Moreover, the variance in lifetime fertility among ever-married women is on average 3.8 per cent lower, indicating more homogenous lifetime fertility than in the full sample.

Key estimates for living children are also strongly correlated with the original estimates for all live-born children (parity-specific stopping $r = 0.93$; for parity-independent strategies $r = 0.92$ for mean and $r = 0.75$ for variance in lifetime fertility). The relationship between parity-specific estimates based on the two measures is near equality (slope = 1.01, intercept = −0.02). As expected from the lower counts of living children, lifetime fertility among women following parity-independent approaches is lower among high-parity populations when considering only living children (0.4-child difference for populations with mean parity four and 0.9-child difference at mean parity six). Notably, the variance in lifetime fertility is also substantially lower when the model is applied to living children, with a 59 per cent reduction compared with live-born children. This indicates that a large portion of heterogeneity in completed fertility among women following a parity-independent approach is due to differences in the number of non-surviving children (Ben-Porath 1976; Hossain et al. 2007).

How are parity-specific stopping and parity-independent strategies associated with education and household wealth?

The estimated proportion of individuals following parity-specific stopping is associated with other factors normally coinciding with declining fertility: sample's mean household wealth ($r = 0.49$, $p < 0.001$) and proportion of women with at least primary education ($r = 0.44$, $p < 0.001$) or secondary education ($r = 0.61$, $p < 0.001$). There is also a strong negative correlation between mean parity-independent lifetime fertility and mean household wealth ($r = -0.58$, $p < 0.001$) as well as the proportion of women with at least primary education ($r = -0.63$, $p < 0.001$) or at least secondary education ($r = -0.72$, $p < 0.001$).

How much of the fertility decline is attributable to declining parity-independent lifetime fertility vs. increasing proportions of parity-specific stopping

Figure 4 illustrates the long-term change across US Census birth cohorts in estimates for: (a) mean completed fertility; (b) the proportions of women engaging in parity-specific stopping; and (c) lifetime fertility among those following a parity-independent approach. Mean completed fertility exhibits an initial sharp decline and then a reversal associated with the baby boom (Figure 4(a)), as does lifetime fertility among women following a parity-independent approach (Figure 4(c)). The initial fertility decline between the 1880 and 1910 cohorts, from 3.2 to 2.3 children, is accompanied by an increase in parity-specific stopping from approximately 5 to 25 per cent of women (Figure 4(b)). However, our decomposition of the decline into that attributable to either of the two strategies, indicates that if the only thing to change during the decline had been parity-specific stopping, we would have expected a decline of only 0.1 children instead of the observed 0.9-child reduction. By contrast, if the only thing to change during that time had been declining fertility among women following a parity-independent approach, we would have expected a decline of 0.8 children, much closer to the observed decline. A similar result is seen when we examine the second decline, from 3.1 to 2.4 children, between cohorts born in the 1930s and 1940s. In this case, a 0.7-child reduction is attributable to parity-independent strategies, while only a 0.2-child reduction is attributable to parity-specific stopping.

We observe similar results when considering the 39 DHS countries that exhibited fertility declines of more than 0.5 children between the first and last surveys (Figure 5). In most countries, the expected contribution to the decline from parity-independent strategies alone is much greater than that from parity-specific stopping alone. Moreover, the pattern is quite consistent across all five world regions, as well as the US, suggesting that the

Figure 4 Historical changes in US fertility: (a) mean completed fertility; (b) estimated proportion following parity-specific stopping; and (c) mean lifetime fertility among those following a parity-independent approach
Note: Twelve cohorts of ever-married, white non-Hispanic women aged 45–49 and 50–54 in the US. Charts show point estimates and 95 per cent posterior bands.
Source: Authors' calculations from 1940–90 US Census data.

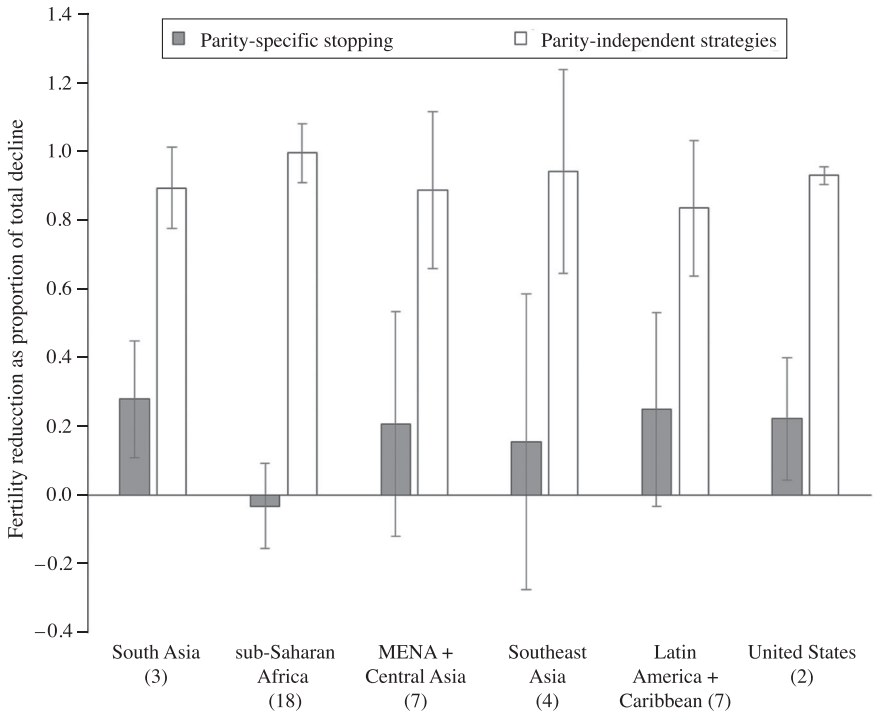

Figure 5 Estimated contributions of parity-specific stopping and parity-independent strategies to fertility declines in five world regions and the US

Note: Includes only countries with at least a 0.5-child reduction over DHS surveys and two declines in the US (1888–1908 and 1933–43 cohorts). Bars show means and lines show 95 per cent confidence intervals for country estimates. Sample size for each region in parentheses. MENA = Middle East and North Africa.

Source: As for Figure 2.

majority of fertility declines (at least in the twentieth century) can be attributed to parity-independent strategies (through spacing, postponement, starting reproduction later, or stopping for other reasons) rather than a switch to parity-specific stopping (Moultrie et al. 2012).

Discussion

Our findings suggest that an increase in parity-specific stopping accompanies many fertility declines, but is not necessarily responsible for much of the overall reduction. Rather, in both the US and the low- and middle-income country samples, the bulk of fertility decline appears to be due to a decline in births among women following parity-independent strategies. Moreover, this parity-independent approach seems to be quite widespread and is not unique to a given world region. This suggests that observations of strategic birth spacing and postponement in sub-Saharan Africa, as well as delayed first birth and parity-independent stopping, are more widespread globally (Bledsoe and Banja 2002; Johnson-Hanks 2007; Moultrie et al. 2012; Van Lith et al. 2013) and may even account for historical declines in higher-income countries, such as the US.

These findings raise the possibility that parity-specific stopping is a by-product of the same societal changes leading to declining fertility, but does not necessarily play a direct causal role in fertility declines. One explanation consistent with this is that increasing opportunities through education encourage a large portion of women to reduce fertility by spacing and postponing their births. Meanwhile, increasing education also changes thinking among a subset of women, so that they begin to engage in parity-specific stopping. Although increases in parity-specific stopping coincide with overall fertility decline, under this scenario parity-specific stopping would, in fact, be a lagging indicator of fertility decline.

The model also identifies potentially significant within-sample heterogeneity in the proportion following a parity-independent approach. This suggests that characterizing entire societies or populations as having one strategy or another (e.g., 'controlled fertility' or 'natural fertility' populations, or societies that focus on child numbers or manage reproduction) masks potentially important variation in strategies

within populations (Johnson-Hanks 2008; Stulp and Barrett 2016; Stulp et al. 2016). Focusing on changes in within-sample heterogeneity might also shed light on observed shifts from over- to under-dispersion in completed fertility relative to a Poisson distribution as samples move from high to low fertility (van Daalen and Caswell 2015; Hruschka and Burger 2016). Specifically, we show that in high- and medium-fertility populations most women follow a parity-independent approach, which creates distributions that are more likely to be over-dispersed relative to a Poisson distribution. By contrast, in samples with lower fertility, a larger portion follow parity-specific stopping at low parities, which creates distributions that are under-dispersed relative to a Poisson distribution.

Estimates of heterogeneity in parity-specific stopping also raise important new questions. Our findings suggest that a much larger proportion of women follow parity-specific stopping in low-fertility than high-fertility settings. Moreover, the samples with high levels of parity-specific stopping exhibit variation, with modes ranging from one (Ukraine, 2007) and two (most countries) to three (Armenia, 2005). This in turn raises important questions about what causes individuals to focus on specific sets of targets once they begin engaging in parity-specific stopping. Are these due to the influence of culturally arbitrary social norms or rather to ecological constraints that lead to different optima (Melkersson and Rooth 2000; Santos Silva and Covas 2000; Winterhalder and Leslie 2002; Sobotka and Beaujouan 2014)?

Our model is intended as a first approximation to population variation in strategies, but also provides a platform for further refinements. Parity-independent strategies comprise a range of behaviours including delaying age at first birth (Allal et al. 2004), spacing of births conditional on age of previous children, differing likelihoods of replacement for prior infant deaths, and postponement for other reasons (Timæus and Moultrie 2008). This raises important questions about the distal causes of variation in these different proximate mechanisms, and whether these arise from greater availability of contraceptive technology, or education leading to later age at first birth or marriage, to name a few possibilities. The modelling framework we outline here provides an avenue for answering these questions, with the possibility of adding fixed effect predictors (e.g., maternal education) of the key model parameters. Our analyses also focus only on completed fertility at the end of women's reproductive careers. However, it is likely that women's strategies vary over time as they enter new life stages and situations. Future work examining the progression of age-specific fertility distributions over the life course may permit a more finely grained analysis of how strategies change over the life course. Finally, future work with the model should also examine how sensitive estimates and fit are to different kinds of data. For example, data limitations from the US required us to limit parities to a maximum of twelve. This may have provided less information for fitting the Gamma–Poisson component of the model, thus increasing the uncertainty of estimates and reducing fit in any specific estimation. It may also account for the fact that in two out of twelve US samples (but only three DHS samples) there is a significant difference between the model prediction and the observed distribution.

As an approximation based on coarse-grained population data, the model has several limitations. Most notably, it can be difficult to infer the underlying criteria people use to make decisions from the patterns of behaviour observed in populations. However, we suggest that this is one approach that can be used alongside other methods in efforts to triangulate the fertility decision-making process. This is bolstered by the associations of our estimates of parity-specific stopping with other established measures of parity-based limitation and deviations from a Poisson process derived from micro-level birth records. At a minimum, the model provides a novel set of statistics on fertility distributions that researchers can begin to include in the existing toolkit for characterizing population variation and change in fertility. Acknowledging these limitations, a key advantage of the method presented here is that it is relatively simple and has minimal data requirements—needing only synchronic data on completed fertility—compared with existing methods for estimating stopping and spacing. Thus, it opens up opportunities for comparing the contributions of parity-specific stopping and parity-independent strategies in a wide range of existing contemporary and historical data sets where accurate information on the distribution of completed fertility is available.

Notes and acknowledgements

1 Daniel J. Hruschka, Joseph Hackman, and Alexandria Drake are at the School of Human Evolution and Social Change, Arizona State University. Rebecca Sear is at the Department of Population Health, London School of Hygiene and Tropical Medicine. Please direct

all correspondence to Daniel J. Hruschka, School of Human Evolution and Social Change, Arizona State University, Tempe, AZ 85287, USA; or by E-mail: Daniel.hruschka@asu.edu

2 We are grateful to Heidi Colleran, Gert Stulp, and Mary Towner for constructive feedback on earlier versions of the paper. The work was supported by the National Science Foundation grant BCS-1150813, jointly funded by Programs in Cultural Anthropology, Social Psychology Program and Decision, Risk, and Management Sciences; and BCS-1658766, jointly funded by Programs in Cultural Anthropology and Methodology, Measurement and Statistics. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

Allal, N., R. Sear, A. M. Prentice, and R. Mace. 2004. An evolutionary model of stature, age at first birth and reproductive success in Gambian women, *Proceedings of the Royal Society of London B: Biological Sciences* 271(1538): 465–470.

Anderton, D. L. and L. L. Bean. 1985. Birth spacing and fertility limitation: a behavioral analysis of a nineteenth century frontier population, *Demography* 22 (2): 169–183.

Arnold, F. 1990. Assessment of quality of birth history data in Demographic and Health Surveys, in *An Assessment of DHS-I Data Quality*. Columbia, Maryland: Institute for Resource Development / Macro Systems, Inc, pp. 81–111.

Ben-Porath, Y. 1976. Fertility response to child mortality: micro data from Israel, *Journal of Political Economy* 84(4, Part 2): S163–S178.

Bledsoe, C. H. and F. Banja. 2002. *Contingent Lives: Fertility, Time, and Aging in West Africa*. Chicago, USA: University of Chicago Press.

Bledsoe, C., F. Banja, and A. G. Hill. 1998. Reproductive mishaps and western contraception: an African challenge to fertility theory, *Population and Development Review* 24(1): 15–57.

Bongaarts, J. and R. E. Potter. 2013. *Fertility, Biology, and Behavior: An Analysis of the Proximate Determinants*. New York, USA: Academic Press.

Brass, W. 1958. The distribution of births in human populations in rural Taiwan, *Population Studies* 12(1): 51–72.

Brass, W., F. Juarez, and A. Scott. 1997. An analysis of parity-dependent fertility falls in tropical Africa, in G. W. Jones, R. M. Douglas, J. C. Caldwell, and R. M. D'Souza (eds), *The Continuing Demographic Transition*. Oxford, UK: Oxford University Press, pp. 80–93.

Casterline, J. B. and C. Odden. 2016a. Exceptional or global? Evidence for postponement in the developing world. Presented at Annual Meeting of the Population Association of America, Washington, DC, March.

Casterline, J. B. and C. Odden. 2016b. Trends in inter-birth intervals in developing countries 1965–2014, *Population and Development Review* 42(2): 173–194.

Cleland, J., S. Bernstein, A. Ezeh, A. Faundes, A. Glasier, and J. Innis. 2006. Family planning: the unfinished agenda, *The Lancet* 368(9549): 1810–1827.

Coale, A. J. 1973. The demographic transition reconsidered. Presented at Congres de l'IUSSP.

Coale, A. J. and T. J. Trussell. 1974. Model fertility schedules: variations in the age structure of childbearing in human populations, *Population Index* 40(2): 185–258.

Coombs, L. C. 1974. The measurement of family size preferences and subsequent fertility, *Demography* 11(4): 587–611.

David, P. A., T. A. Mroz, W. C. Sanderson, K. W. Wachter, and D. R. Weir. 1988. Cohort parity analysis: statistical estimates of the extent of fertility control, *Demography* 25(2): 163–188.

Eijkemans, M. J. C., F. Van Poppel, D. F. Habbema, K. R. Smith, H. Leridon, and E. R. Te Velde. 2014. Too old to have children? Lessons from natural fertility populations, *Human Reproduction* 29(6): 1304–1312.

Feng, W., J. Lee, and C. Campbell. 1995. Marital fertility control among the Qing nobility: implications for two types of preventive check, *Population Studies* 49(3): 383–400.

Geweke, J. 1991. *Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments*. MN, USA: Federal Reserve Bank of Minneapolis, Research Department Minneapolis. Staff report, Vol. 196.

Hayford, S. R. and V. Agadjanian. 2011. Uncertain future, non-numeric preferences, and the fertility transition: a case study of rural Mozambique, *African Population Studies* 25(2): 419–439.

Henry, L. 1961. Natural fertility–observation, theory, results, *Population* 16(4): 625–636.

Hionidou, V. 1998. The adoption of fertility control on Mykonos, 1879–1959: stopping, spacing or both? *Population Studies* 52(1): 67–83.

Hossain, M. B., J. F. Phillips, and T. K. LeGrand. 2007. The impact of childhood mortality on fertility in six rural thanas of Bangladesh, *Demography* 44(4): 771–784.

Hruschka, D. J. and O. Burger. 2016. How does variance in fertility change over the demographic transition? *Philosophical Transactions of the Royal Society B: Biological Sciences* 371(1692): 1–8.

Hruschka, D. J., D. Gerkey, and C. Hadley. 2015. Estimating the absolute wealth of households, *Bulletin of the World Health Organization* 93(7): 483–490.

ICF. 1985–2016. Demographic and Health Surveys (various) [Datasets]. Funded by USAID. Rockville, Maryland: ICF [Distributor].

Johnson-Hanks, J. 2002. On the limits of life stages in ethnography: toward a theory of vital conjunctures, *American Anthropologist* 104(3): 865–880.

Johnson-Hanks, J. 2007. Natural intentions: fertility decline in the African demographic and health surveys, *American Journal of Sociology* 112(4): 1008–1043.

Johnson-Hanks, J. 2008. Demographic transitions and modernity, *Annual Review of Anthropology* 37: 301–315.

Kingsley, D. 1963. The theory of change and response in modern demographic history, *Population Index* 29(4): 345–366.

Knodel, J. 1977. Family limitation and the fertility transition: evidence from the age patterns of fertility in Europe and Asia, *Population Studies* 31(2): 219–249.

Knodel, J. 1979. From natural fertility to family limitation: The onset of fertility transition in a sample of German villages, *Demography* 16(4): 493–521.

Knodel, J. 1983. Natural fertility: age patterns levels and trends, in R. A. Bulatao and R. D. Lee (eds), *Determinants of Fertility in Developing Countries*. New York: Academic Press. pp. 61–102.

Knodel, J. 1987. Starting, stopping, and spacing during the early stages of fertility transition: The experience of German village populations in the 18th and 19th centuries, *Demography* 24(2): 143–162.

Knodel, J. and E. Van de Walle. 1979. Lessons from the past: policy implications of historical fertility studies, *Population and Development Review* 5(2): 217–245.

Mason, K. O. 1997. Explaining fertility transitions, *Demography* 34(4): 443–454.

Mattison, S., C. Moya, A. Reynolds, and M. C. Towner. 2018. Evolutionary demography of age at last birth: integrating approaches from human behavioural ecology and cultural evolution, *Philosophical Transactions of the Royal Society B: Biological Sciences* 373(1743): 1–12.

Melkersson, M. and D.-O. Rooth. 2000. Modeling female fertility using inflated count data models, *Journal of Population Economics* 13(2): 189–203.

Miranda, A. 2010. A double-hurdle count model for completed fertility data from the developing world, in *NCRM Working Paper*. London: Department of Quantitative Social Science, Institute of Education, pp. 1–45.

Moultrie, T. A., T. S. Sayi, and I. M. Timæus. 2012. Birth intervals, postponement, and fertility decline in Africa: A new type of transition? *Population Studies* 66(3): 241–258.

Plummer, M., N. Best, K. Cowles, and K. Vines. 2006. CODA: convergence diagnosis and output analysis for MCMC, *R News* 6(1): 7–11.

Poston Jr, D. L. and S. L. McKibben. 2003. Using zero-inflated count regression models to estimate the fertility of US women, *Journal of Modern Applied Statistical Methods* 2(2): 371–379.

R Core Team. 2014. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. 2013. ISBN 3-900051-07-0.

Reher, D. S. and A. Sanz-Gimeno. 2007. Rethinking historical reproductive change: insights from longitudinal data for a Spanish town, *Population and Development Review* 33(4): 703–727.

Ruggles, S., K. Genadek, R. Goeken, J. Grover, and M. Sobek. 2015. Integrated Public use Microdata Series: Version 6.0 [Machine-readable database]. edited by U.o. Minnesota. Minneapolis.

Santos Silva, J. M. C. and F. Covas. 2000. A modified hurdle model for completed fertility, *Journal of Population Economics* 13(2): 173–188.

Sobotka, T. and É. Beaujouan. 2014. Two is best? The persistence of a two-child family ideal in Europe, *Population and Development Review* 40(3): 391–419.

Stulp, G. and L. Barrett. 2016. Wealth, fertility and adaptive behaviour in industrial populations, *Philosophical Transactions of the Royal Society B: Biological Sciences* 371(1692): 1–11.

Stulp, G., R. Sear, S. B. Schaffnit, M. C. Mills, and L. Barrett. 2016. The reproductive ecology of industrial societies: the association between wealth and fertility, *Human Nature* 27(4): 445–470.

Szreter, S. and E. Garrett. 2000. Reproduction, compositional demography, and economic growth: family planning in England long before the fertility decline, *Population and Development Review* 26(1): 45–80.

Timæus, I. M. and T. A. Moultrie. 2008. On postponement and birth intervals, *Population and Development Review* 34(3): 483–510.

Timæus, I. M. and T. A. Moultrie. 2013. Distinguishing the impact of postponement, spacing and stopping on birth intervals: evidence from a model with heterogeneous fecundity, *Journal of Biosocial Science* 45(03): 311–330.

Towner, M. C., I. Nenko, and S. E. Walton. 2016. Why do women stop reproducing before menopause? A life-history approach to age at last birth, *Philosophical Transactions of the Royal Society B: Biological Sciences* 371(1692): 1–12.

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

Van Bavel, J. 2004b. Detecting stopping and spacing behaviour in historical demography. A critical review of methods, *Population* 59(1): 117–128.

Van Bavel, J. and J. Kok. 2010. A mixed effects model of birth spacing for pre-transition populations: evidence

of deliberate fertility control from nineteenth century Netherlands, *The History of the Family* 15(2): 125–138.

van Daalen, S. and H. Caswell. 2015. Lifetime reproduction and the second demographic transition: stochasticity and individual variation, *Demographic Research* 33: 561–588.

Van de Walle, E. 1992. Fertility transition, conscious choice, and numeracy, *Demography* 29(4): 487–502.

Van Lith, L. M., M. Yahner, and L. Bakamjian. 2013. Women's growing desire to limit births in sub-Saharan Africa: meeting the challenge, *Global Health: Science and Practice* 1(1): 97–107.

Westoff, C. F. and A. Bankole. 2000. Trends in the demand for family limitation in developing countries, *International Family Planning Perspectives* 26(2): 56–97.

Winkelmann, R. 1995. Duration dependence and dispersion in count-data models, *Journal of Business & Economic Statistics* 13(4): 467–474.

Winterhalder, B. and P. Leslie. 2002. Risk-sensitive fertility: the variance compensation hypothesis, *Evolution and Human Behavior* 23(1): 59–82.

Wood, J. W. 1994. *Dynamics of Human Reproduction: Biology, Biometry, Demography*. New York, USA: Transaction Publishers.