



HAL
open science

Measuring sex-selective abortion: How many women abort?

Aditi Dimri, Véronique Gille, Philipp Ketz

► **To cite this version:**

Aditi Dimri, Véronique Gille, Philipp Ketz. Measuring sex-selective abortion: How many women abort?. 2021. halshs-03495964

HAL Id: halshs-03495964

<https://shs.hal.science/halshs-03495964>

Preprint submitted on 20 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



WORKING PAPER N° 2021 – 67

Measuring sex-selective abortion: How many women abort?

**Aditi Dimri
Véronique Gille
Philipp Ketz**

JEL Codes: C13, D10, J13

Keywords: Sex ratio, Sex-selective abortion, Measurement, India, Constrained maximum likelihood estimation, Random constraint(s)



Measuring sex-selective abortion: How many women abort?*

Aditi Dimri[†] Véronique Gille[‡] Philipp Ketz[§]

December 15, 2021

Abstract

Current measurement of sex-selective abortion is based on observing an imbalance between the sex ratio at birth and the natural sex ratio, providing us with the number of missing female fetuses. However, this measure does not tell us how widespread this phenomenon is, i.e., how many women abort, which will not be equal to the number of sex-selective abortions if there is repeated sex-selective abortion. In this paper, we show that the number of women that abort *between two consecutive births* and whether they do so repeatedly can be inferred using sex ratios and information on birth spacing. We apply our model to Indian DHS data to estimate how many women abort and to assess whether they do so repeatedly between two births. The results depend on the birth order and siblings composition: For example, we find that women whose first born is a girl abort at most once before the birth of the second child, i.e., (almost) none of them abort a second time if again pregnant with a girl after a first abortion. In contrast, we find evidence of repeated sex-selective abortion before the birth of the third child among women whose first two children are girls. We also introduce a novel *constrained* maximum likelihood estimator that imposes a (set of) *random* constraint(s) and that may be of independent interest.

JEL Classification: C13, D10, J13.

Keywords: sex ratio, sex-selective abortion, measurement, India, constrained maximum likelihood estimation, random constraint(s).

*We thank François Maniquet as well as seminar participants at CORE and Delhi School of Economics for useful comments.

[†]University of Warwick. E-mail address: aditidimri@gmail.com

[‡]IRD, UMR LEDa, PSL, Université Paris-Dauphine. E-mail address: veronique.gille@ird.fr

[§]Paris School of Economics, CNRS. E-mail address: philipp.ketz@psemail.eu

1 Introduction

The imbalance in the (observed) sex ratio at birth provides us with the number of missing girls at birth. Chao et al. (2019) estimate that, between 1970 and 2017, 45 million female births were missing in the world, and that India and China account for more than 97% of these missing female births.¹ Prenatal sex-determination followed by sex-selective abortion have been found to account for a majority of these missing female births in India (Jha et al., 2006). However, little is known about how prevalent sex-selective abortion is. Specifically, while the sex ratio at birth tells us the number of sex-selective abortions, it does not tell us how many women, or households, perform sex-selective abortion if some women sex-selectively abort multiple times.² In this paper, we develop a methodology that uses information on the sex ratio and birth spacing—thus only requiring widely available survey data—to infer the number of women that abort between two consecutive births and to what extent they do so repeatedly. The main contribution of our proposed methodology is that it does *not* rely on standard sources of abortion data such as administrative data from abortion facilities/providers and survey data on self-reported abortion by women or households. The former fails to provide accurate measures on the number of women who abort because sex-selective abortion is illegal and likely to be carried out at non-certified/non-registered clinics. Similarly, the latter is unlikely to provide accurate information due to under-reporting (of any kind of abortion).

The adverse consequences of an unbalanced sex ratio are widely studied and acknowledged;³ governments across the globe have tried to tackle the issue of sex-selective abortion and to reduce imbalances in the sex ratio at birth. But why should we care about the number of women that abort and repeated abortions apart from the total number of abortions? First, we argue that knowledge of the presence of repeated abortions has value on its own. This is because safe abortions remain largely unavailable which, in turn, leads to concerns of morbidity and mortality for women (Grimes et al., 2006; Jain et al., 2004). Given that sex-selective abortions are illegal and more likely to be carried out in unsafe conditions, their repeated use can be extremely detrimental to maternal health, both physically and mentally. Second, the number of missing female births is sometimes used as a proxy for son preference. Thus, breaking down the number of missing female births into the number of women that abort and repeated abortions opens the door for more refined definitions of son preference.

¹In a comprehensive policy report, Guilmoto (2012) records other countries with unbalanced sex ratio as Singapore, Vietnam, Pakistan, Azerbaijan, Armenia, Georgia, Albania and Montenegro.

²In this paper, we do *not* model the decision making process underlying sex-selective abortion and, therefore, refer to our observational units as women or households interchangeably, without taking a stance on who the decision maker is.

³See Anukriti et al. (2021) for a review.

Last, furthering the understanding of the pattern of sex-selective abortion may prove helpful for improving targeting of policies against sex-selective abortion.

In this paper, we model spacing between two births in the presence of sex-selective abortion. We show that the fact that each sex-selective abortion simultaneously increases spacing and changes the sex ratio implies that the distribution of spacing differs by the gender of the next-born.⁴ Based on this observation, we show that 1) differences in mean spacing by gender (of the next-born) can be used to infer whether women perform sex-selective abortion (between two births) and, in some cases, whether they do so repeatedly, 2) differences in mean spacing by gender along with information on the sex ratio can be used to compute bounds on the share of women that abort once and the share of women that abort twice, and 3) the entire distribution of spacing, when combined with distributional assumptions, can be used to estimate the share of women that abort once and twice. Importantly, using differences in mean spacing by gender does not require any assumption on the natural sex ratio, which is a contested topic (Chao et al., 2019). A limitation of our model is that we assume that there is no correlation between how long households wait before trying to conceive another child and their decision to abort if pregnant with a girl.⁵ While we discuss the implications of this assumption, we leave an extension of our model that relaxes it for future research.

The aforementioned estimation procedure that relies on additional distributional assumptions is a “standard” maximum likelihood (ML) estimator. Motivated by the “standard” method of moments estimator that relates the observed sex ratio at birth to the natural sex ratio and that is used to estimate the number of missing female births, we also introduce a novel *constrained* maximum likelihood (CML) estimator that imposes the corresponding moment as a *random* (equality) constraint. We show that this estimator is (locally) equivalent to a particular generalized method of moments (GMM) estimator. This equivalence result provides us with a useful interpretation that not only helps motivating the estimator but also delivers a simple formula for (the estimator of) the asymptotic variance of the (CML) estimator. The proposed estimator along with its interpretation may also prove useful in other contexts.

To illustrate the usefulness of our proposed methodology we use the 2015-2016 round of the DHS survey for India, which contains basic household level reproductive data. Abortion

⁴We use the term “next-born” throughout the paper to refer to the child born *after* the spacing time under consideration. For example, when we discuss spacing between the first and second child, the sex of the next-born refers to the sex of the second child.

⁵As mentioned above, we do not model the decision making process underlying sex-selective abortion. We, therefore, abstract from the fact that total fertility or abortion behaviour at other birth orders may interact with the decision to abort.

has been legal in India since the *Medical Termination of Pregnancy Act, No. 36, Government of India (1971)*, which allows a pregnancy to be terminated up to 20 weeks gestation by registered allopathic medical practitioners at certified abortion facilities. However, sex-selective abortion is not legal and is controlled through the *Pre-Natal Diagnostic Techniques (Regulation and Prevention of Misuse) Act, No. 57, Government of India (1994)*, which prohibits the misuse of antenatal diagnostic tests for the purpose of sex determination, prohibits the advertisement of such tests, requires registration of all facilities that use them, and prohibits those conducting the tests from revealing the sex of the foetus to the expectant parents. Despite this law, sex-selective abortion is widely practised: Chao et al. (2019), for example, estimate that, in India, 20.7 millions female foetuses have been aborted between 1970 and 2017.

We apply our methodology to subsamples of our data, or “groups of women”, that are defined by (different) birth orders and sibling compositions. We identify several groups with and without balanced sex-ratios. We say that a sex ratio is balanced if the observed sex ratio does not (statistically significantly) exceed the natural sex ratio that we take to be 0.513 as suggested by Chao et al. (2019) and Dubuc and Coleman (2007). For groups with balanced sex ratios, such as women whose first-born is a son, we find no difference in (mean) spacing by gender of the next-born. For the groups with unbalanced sex ratios, our different methods largely agree with each other. For the group of women that had a girl as first-born, there is evidence of sex-selective abortion (before the birth of the second child), but no evidence of repeated sex-selective abortion. Our (C)ML estimation results suggest that $\sim 9\%$ of women that get pregnant with a girl abort once but only once, i.e., women do not abort a second time when pregnant again with a girl after a first abortion. For the group of women that first had two girls, we not only find evidence of sex-selective abortion but also evidence of repeated sex-selective abortion (before the birth of the third child). Our (C)ML estimation results suggest that $\sim 13\%$ of women that get pregnant with a girl abort and $\sim 54\%$ of those abort again when pregnant again with a girl after a first abortion.

Our paper contributes to the literature on the measurement of sex-selective abortions. Since the seminal paper of Sen (1992) on missing women, the literature has focused on refining the methodology centred around comparing the observed sex ratio to an assumed natural sex ratio, to improve the estimates on the number of missing women (Anderson and Ray, 2010) and the number of missing female births (see for example Bongaarts and Guilmoto, 2015; Bhalotra and Cochrane, 2010; Chao et al., 2019; Klasen and Wink, 2003). Our paper complements this literature with a new measure, the number of women who abort, which helps understanding how widespread sex-selective abortion is. While we are not the first to note that sex-selective abortion impacts spacing (see Portner, 2020; Anukriti et al., 2021),

to the best of our knowledge there is no other paper that exploits the fact that sex-selective abortion implies differential spacing by gender of the next-born to add another dimension to the measurement of sex-selective abortion. In doing so, our paper also contributes to the literature interested in understanding birth spacing behaviour and its consequences on women’s and children’s health and economic outcomes (see for example Barclay and Kolk, 2017; Bhalotra and Van Soest, 2008; Jayachandran and Kuziemko, 2011; Milazzo, 2018).

The rest of the paper is organized as follows. Section 2 describes our theoretical model and presents the data. Section 3 presents the two propositions derived from our model and applies them to the data. Section 4 introduces and applies our proposed estimation method and section 5 concludes. All proofs are collected in Appendix A.

2 Sex-selective abortion and the distribution of spacing by sex of the next-born

Section 2.1 proposes a statistical model that captures the fact that the time between two births depends on the number of abortions performed between them.⁶ Section 2.2 then shows that, in the presence of sex-selective abortion, the distribution of spacing differs by the sex of the next-born child. Section 2.3 provides an illustration using Indian data.

2.1 Birth spacing in the presence of sex-selective abortions

The time between two consecutive births, or simply (birth) spacing, is denoted by T_b^A (*time to birth*) and depends on the *unobserved*, or latent, number of sex-selective abortions, $A \in \{0, 1, \dots, n_0\}$, performed between the two births, where n_0 denotes the (known) maximal number of “possible” abortions. We sometimes suppress the dependence of T_b^A on A and write $T_b = T_b^A$. Here and in what follows, we let capital letters denote random variables and the corresponding lower case letters their (possible) realizations. We postulate the following model

$$T_b^a = T_w + \sum_{k=1}^a (T_s + T_{c,k}) + T_p. \quad (1)$$

T_w denotes the *waiting time* (until conception), i.e., the time that the household, after the birth of a child, waits before trying to conceive again *plus* the time it takes to conceive. T_s denotes the *time to screening*, i.e., the time it takes from the moment of conception until

⁶We do not model the number of abortions that a women performs over her lifetime. Between two births refers to two consecutive births. For the first-born this refers to the time between marriage and birth of the first child.

the sex of the foetus is known.⁷ This is the time the household has to wait before it can perform a sex-selective abortion (if applicable). We do not index T_s by the abortion cycle, because we take it to be constant. $T_{c,k}$ denotes the *time to conception* after the k^{th} abortion, i.e., the time it takes for the woman to get pregnant after having aborted k times. Lastly, T_p denotes the *time of pregnancy*, which we take to be equal to nine months. For $a = 0$, i.e., when there is no sex-selective abortion, birth spacing equals $T_w + T_p$.

The above model makes the following assumptions. First, we assume that, after an abortion, the household tries to conceive again without any (additional) waiting time.⁸ Second, we assume that abortions not performed for the purpose of sex-selection as well as miscarriages are equally likely for male or female foetuses and are thus taken into account by $T_{c,k}$, as they are independent of the sex of the foetus. In what follows, we assume that $T_{c,k}$ is i.i.d. across k , with $T_{c,k} \sim T_c$. Even if there is heterogeneity, we assume it to be independent of the abortion decision. Lastly, we assume T_w to be independent of the willingness to abort. While households may have a different T_w depending on whether they plan to abort if pregnant with a girl, we do not model it. However, we discuss the implications of this assumption in Section 3.1.

2.2 Difference in spacing by sex of the next-born child

The model introduced in the previous section implies that sex-selective abortion impacts the distribution of spacing, $P(T_b < t)$. Moreover, it does so differently depending on the sex of the next-born child, i.e., $P(T_b < t|Y = 1) \neq P(T_b < t|Y = 0)$ for at least some $t \in \mathcal{T}$, where $Y \in \{0, 1\}$ denotes the *sex of the next-born child*, with $Y = 0$ indicating the birth of a girl and $Y = 1$ indicating the birth of a boy, and where \mathcal{T} denotes the support of T_b .

Before we demonstrate this, we introduce the main parameters of our model, $\alpha_k \in [0, 1]$ for $k \in \{1, \dots, n_0\}$, where α_k denotes the probability that a household performs the k^{th} sex-selective abortion when being pregnant with a girl. For what follows, it is convenient to introduce α_0 which is set equal to one and α_{n_0+1} which is set equal to zero. Furthermore, we let $n \leq n_0$ denote the maximal number of abortions *in the population*, i.e., n is such that $\alpha_k > 0$ for $k \in \{1, \dots, n\}$ and $\alpha_{n+1} = 0$.

⁷For simplicity, we assume that every household knows the sex of the foetus. However, the implications of the model are identical if we assume that those who do not plan to sex-selectively abort do not know the sex of the foetus.

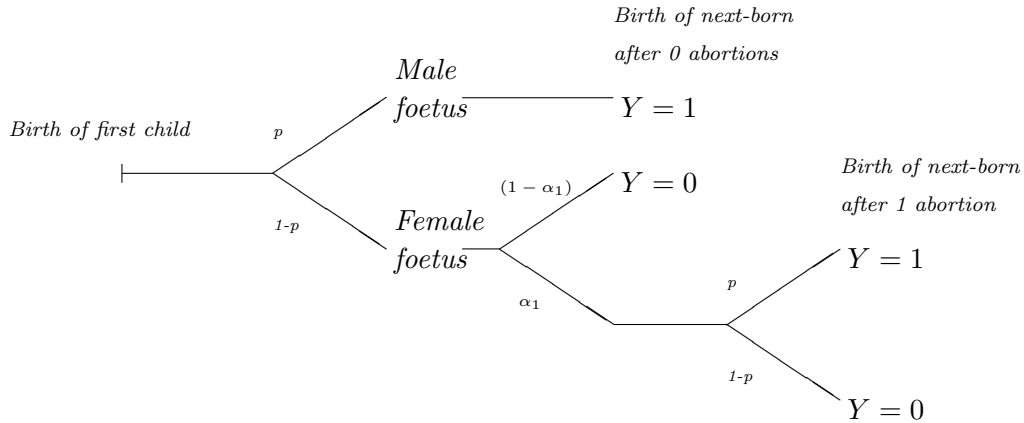
⁸The time to screening T_s is, thus, better thought of as the time it takes to know the sex of the foetus *plus* the time the body needs to recover after an abortion. In the empirical analysis we check the robustness of our results to different values of t_s . The assumption that there is no waiting time in addition to the time it takes for the body to recover from an abortion comes from the fact that the decision to have another child has already been made before the first pregnancy.

2.2.1 Example with (at most) one sex-selective abortion ($n_0 = 1$)

To help fix ideas, let us look at possible abortion patterns between the first and the second birth. If there is no sex-selective abortion between the first child and the second child, the distribution of spacing is the same if households have a boy or a girl as a second child, i.e., $T_b|Y = 1 \sim T_b|Y = 0$, where \sim denotes “is distributed as”. But if there is sex-selective abortion the distribution of spacing when the second child is a boy will be different from the distribution of spacing when the second child is a girl, i.e., $T_b|Y = 1 \not\sim T_b|Y = 0$, where $\not\sim$ denotes “is not distributed as”.⁹

To illustrate that $T_b|Y = 1 \not\sim T_b|Y = 0$ in the presence of sex-selective abortion, consider Figure 1. In this example, we assume that women abort at most once, i.e., $A \in \{0, 1\}$. Furthermore, let p denote the “natural” sex ratio or, equivalently, the “natural” probability of getting pregnant with a boy. We use these two terms interchangeably.¹⁰

Figure 1: Only one abortion



At conception, a woman is pregnant with a male foetus with probability p and pregnant with a female foetus with probability $1 - p$. If she is pregnant with a male foetus, assuming there is no other abortion than sex-selective, the boy is born after 9 months, i.e., $T_b = T_b^0$. If she is pregnant with a female foetus, the household aborts with probability α_1 and with probability $1 - \alpha_1$ the household keeps the girl. If the household keeps the girl, she is born

⁹Formally, two random variables, X_1 and X_2 , have the same distribution if $P(X_1 < x) = P(X_2 < x)$ for all $x \in \mathcal{X}$, where $\mathcal{X} = \mathcal{X}_1 \cup \mathcal{X}_2$ and where \mathcal{X}_i denotes the support of X_i with $i \in \{1, 2\}$. Similarly, X_1 and X_2 have different distributions if $P(X_1 < x) \neq P(X_2 < x)$ for some $x \in \mathcal{X}$.

¹⁰Note that the “natural” sex ratio is usually expressed as the number of boys born divided by the number of girls born.

after the same time as the boy, i.e., $T_b = T_b^0$. If the female foetus is aborted, then a boy or a girl is born later, i.e., $T_b = T_b^1$.

Given that $A \in \{0, 1\}$, we can write $T_b = T_b^1 A + T_b^0(1 - A)$. In words, birth spacing is a mixture of *time to birth* with one abortion, T_b^1 , and *time to birth* with zero abortions, T_b^0 . Since T_b^1 first-order stochastically dominates T_b^0 , it is easy to see that $T_b|Y = 1 \not\prec T_b|Y = 0$ if T_b^1 and T_b^0 have different “mixture weights” depending on the sex of the next-born, i.e., if $P(A = 1|Y = 1) \neq P(A = 1|Y = 0)$.¹¹ From Figure 1, it follows that

$$P(A = 1|Y = 1) = \frac{P(A = 1, Y = 1)}{P(Y = 1)} = \frac{(1 - p)\alpha_1}{(1 - p)\alpha_1 + 1} \quad (2)$$

and

$$P(A = 1|Y = 0) = \frac{P(A = 1, Y = 0)}{P(Y = 0)} = \frac{(1 - p)\alpha_1}{(1 - p)\alpha_1 + 1 - \alpha_1}. \quad (3)$$

Therefore, $P(A = 1|Y = 1) \neq P(A = 1|Y = 0)$ as long as $\alpha_1 > 0$ or, equivalently, $n = 1$ ($= n_0$), i.e., as long as there is sex-selective abortion. We conclude that in the presence of sex-selective abortion spacing differs by sex of the next-born.

2.2.2 Generalization to $n_0 \geq 1$ sex-selective abortions

More generally, our model assumes that

$$P(A = a, Y = 1) = (1 - p)^a \prod_{k=0}^a \alpha_k \quad (4)$$

and

$$P(A = a, Y = 0) = (1 - p)^a (1 - \alpha_{a+1}) \prod_{k=0}^a \alpha_k. \quad (5)$$

By definition, $P(Y = 1) = \sum_{a=0}^{n_0} P(A = a, Y = 1)$ and $P(Y = 0) = \sum_{a=0}^{n_0} P(A = a, Y = 0)$.

2.3 Data illustration

2.3.1 Data source and sample

We use data from the 2015-2016 round of the Demographic and Health Surveys (DHS) for India. The DHS provides retrospective but precise information on the fertility (full birth

¹¹To see this, note that

$$P(T_b < t|Y = y) = P(T_b^1 < t)P(A = 1|Y = y) + P(T_b^0 < t)(1 - P(A = 1|Y = y)).$$

history) of each woman that is between 15 and 49 years old at the time of the survey. As we want to study abortion patterns at different birth orders and sibling compositions, we focus on women who had at least two children. We also restrict the sample to women who had their first child in or after 2000 and before or in 2010.¹² We only keep women who have ever been married once and did not have any birth before marriage in our sample. Births reported as multiple or twin births, before the mother was 14, and with birth spacing less than 9 months or more than 10 years are dropped as well.

2.3.2 Sex ratios in the data

Table 1 shows the sex ratio for several birth orders and siblings compositions in our sample. 1st born refers to the first-born child, B to the second-born child when the first-born child is a boy, G to the second-born child when the first-born child is a girl, and GG to the third-born child when the first two children are girls.

Table 1: Sex ratio

	Sample size	Sex ratio	Std. error	p-value
Balanced sex ratio				
1 st born	153,740	0.506	0.001	1.000
B*	82,646	0.502	0.002	1.000
Unbalanced sex ratio				
G*	80,369	0.535	0.002	0.000
GG*	25,888	0.555	0.003	0.000

1st born refers to the first-born, B (G) to the second-born when the first-born is a boy (girl), and GG to the third-born when the first two children are girls. The p-value is for a one-sided t-test that tests whether the sex ratio is lower than the natural sex ratio, taken equal to 0.513.

We compare the sex ratio at these different birth orders and sibling compositions to the natural sex ratio, which we take to be 0.513 as computed by Chao et al. (2019) and Dubuc and Coleman (2007). In line with the literature, we find that the sex ratio is “balanced” for

¹²We focus on women who had their first child in or after 2000 to minimize the misreporting of birth information, which has been shown to increase with the recall period of the survey respondent (Schoumaker, 2014; Portner, 2020). We limit to first births before or in 2010 to minimize the number of women that did not have a second child by the time of the survey (77.6% of the women who had their first child in 2010 had a second child by the time of the survey compared to 69.5% for women who had their first child in 2011).

first-born children and for second-born children when the first-born child is a boy.¹³ On the contrary, when the first-born child is a girl, the sex ratio for second-born children is 0.535, and when the first two children are girls, the sex ratio is 0.555. Both these sex ratios are significantly higher than the natural sex ratio. These unbalances in the sex ratio indicate the presence of sex-selective abortion for these birth orders and siblings compositions.

2.3.3 Distribution of spacing in the data

We use the groups as defined in Table 1 to illustrate that the distribution of spacing differs by sex of the next-born (child) in the presence of sex-selective abortion. We first look at the two groups for which the sex ratio does not show any imbalance. The top left graph in Figure 2 shows the histogram of time between marriage and first birth by gender of the first-born. The top right graph shows the corresponding density plots. The bottom two graphs show the corresponding plots for spacing between first-born and second-born when the first-born is a boy.

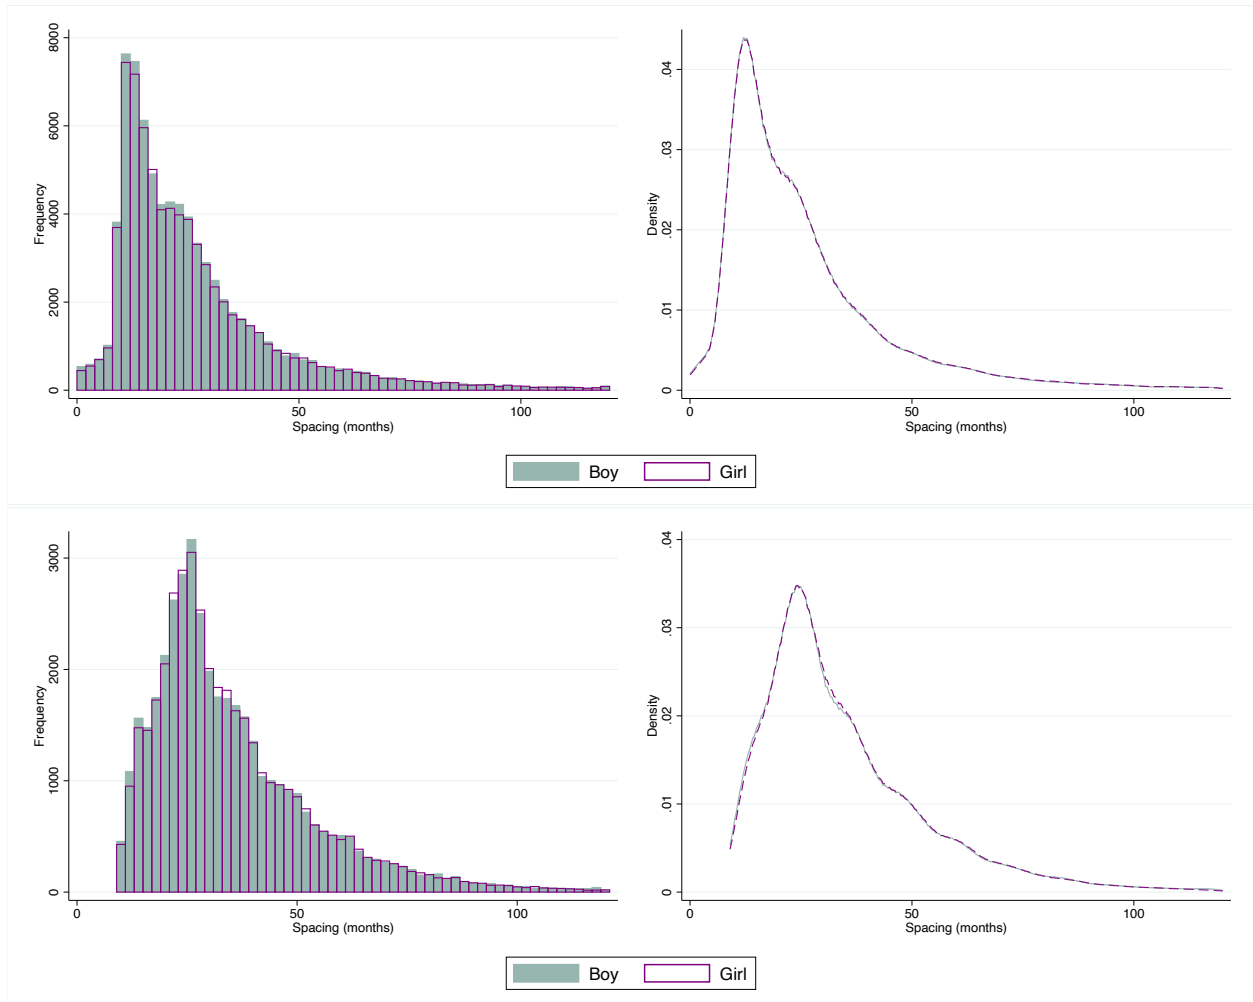
The distributions of spacing for the two groups with balanced sex ratio do not differ by gender of the next-born child. The histograms show that there are more boys born than girls (almost) “each month”, which reflects the fact that the natural sex ratio is higher than 0.5, but the density plots do not show any difference in spacing by gender.

Figure 3 reproduces Figure 2 for the two groups for which the sex ratio is unbalanced, second-born children when the first-born child is a girl (top two graphs) and third-born children when the first two children are girls (bottom two graphs).

Figure 3 shows that the distribution of spacing differs by gender of the next-born child for the groups with unbalanced sex ratios, which becomes particularly apparent in the density plots. As above, the histograms show that “each month” there are more boys born than girls with the differences being more pronounced compared to Figure 2, which is line with high(er) sex ratios.


¹³The sex ratio is actually lower than 0.513 for these two groups. This could be due to the fact that we restrict our sample to households with at least two children: if the probability of having a second child depends on the sex of the first child, then households who had a girl as a first born are likely to be overrepresented.

Figure 2: Distribution of spacing in the absence of sex-selective abortion



Note: Top left: Histogram of time between marriage and first birth by gender of the first-born. Top right: Corresponding density plots. Bottom left: Histogram of time between first-born and second-born when the first-born is a boy by gender of the second-born. Bottom right: Corresponding density plots.

Figure 3: Distribution of spacing in the presence of sex-selective abortion



Note: Top left: Histogram of time between first-born and second-born when the first-born is a girl by gender of the second-born. Top right: Corresponding density plots. Bottom left: Histogram of time between second-born and third-born when the first two children are girls by gender of the third-born. Bottom right: Corresponding density plots.

3 Propositions

The previous section highlights that sex-selective abortion changes the distribution of spacing and that it does so differently depending on the sex of the next-born. These differences can be exploited to get new insights into sex-selective abortion patterns. In Section 3.1, we derive a proposition from our model that can easily be applied to the data to test for (repeated) sex-selective abortion without imposing any further assumptions. In Section 3.2, we assume that there is at most two abortions between two consecutive births to derive bounds on the share of women that abort once and the share of women that abort twice.

3.1 Proposition 1: Repeated abortions or not

Proposition 1 provides easy-to-check (sufficient) conditions that allow us to infer whether there is sex-selected abortion and whether there is repeated sex-selective abortion between two births.

Proposition 1.

- (i) If $E(T_b|Y = 0) \neq E(T_b|Y = 1)$, then $n > 0$ or, equivalently, $\alpha_1 > 0$.
- (ii) If $E(T_b|Y = 0) < E(T_b|Y = 1)$, then $n > 1$ or, equivalently, $\alpha_1, \alpha_2 > 0$

Proposition 1(i) can be used to identify sex-selective abortion and therefore complements the sex ratio. It means that as soon as we observe a different mean spacing “for” girls and boys, we can infer that there is sex-selective abortion.¹⁴ The advantage of using spacing as an indicator of sex-selective abortion is that it is independent of the natural sex ratio, which is not always easily agreed upon.

Proposition 1(ii) can be used to identify repeated sex-selective abortion. When the mean spacing for girls is lower than the mean spacing for boys, there is repeated sex-selective abortion, i.e., a strictly positive fraction of the population aborts a second time when pregnant again with a girl.

As discussed before, the model and hence Proposition 1 assumes that T_w is independent of whether a household aborts if pregnant with a female foetus. In Appendix A, we show that Proposition 1 continues to hold true if T_w is allowed to depend on whether a household aborts, or rather would abort, if pregnant with a female foetus as long as the expected *waiting time*, $E(T_w)$, is (weakly) smaller, or “shorter”, for a household that aborts if pregnant with a female foetus than for a household that does not. The latter condition, albeit not testable,

¹⁴More generally, our model implies that if $P(T_b < t|Y = 0) \neq P(T_b < t|Y = 1)$ for at least some $t \in \mathcal{T}$, then $n > 0$. Therefore, an alternative, potentially more powerful test for the presence of sex-selective abortion would consist in testing whether the two conditional distributions are equal. Here, we focus on the “difference-in-means” for its simplicity and the expositional coherence with the remainder of the paper.

seems likely to hold, as households that plan to abort if pregnant with a girl may anticipate the potentially longer spacing that an abortion implies.

3.1.1 Data illustration continued

We take Proposition 1 to the data to explore what can be learned from differences in mean spacing on sex-selective abortion for our different groups and whether there is evidence of repeated abortion for the birth orders and sibling compositions for which we observe unbalanced sex ratios.

Table 2: Birth Spacing (in months)

	Sample size		$E(T_b Y = y)$		Difference	Std. error	H_0 : diff = 0 p-value
	$y = 0$	$y = 1$	$y = 0$	$y = 1$			
Balanced sex ratio							
1 st born	75,925	77,815	25.81	25.68	0.12	0.10	0.20
B*	41,041	41,278	34.67	34.59	0.07	0.13	0.57
Unbalanced sex ratio							
G*	37,347	43,022	34.28	34.29	-0.01	0.13	0.93
GG*	11,521	14,367	32.96	34.24	-1.28	0.22	0.00

1st born refers to the first-born, B (G) to the second-born when the first-born is a boy (girl), and GG to the third-born when the first two children are girls. Recall that $y = 0$ ($y = 1$) indicates a girl (boy).

Table 2 reports the average spacing in months by gender of the next-born for our four groups. For the two groups without imbalance in the sex ratio at birth, cf. Table 1, there is no (statistically) significant difference in mean spacing by sex of the next-born. Among the two groups with unbalanced sex ratios, the difference in mean spacing is only significantly different from zero for the group where the first two children are girls. Thus, Proposition 1 does not provide further evidence of sex-selective abortion (beyond the unbalanced sex ratio) among households whose first-born is a girl. However, the difference in mean spacing between the second-born and the third-born in the group of women that first had two girls is 1.28 months shorter when the third-born is a girl than when the third-born is a boy. Proposition 1, therefore, provides additional evidence of sex-selective abortion (beyond the unbalanced sex ratio) and, furthermore, suggests the presence of repeated sex-selective abortion in this group.

3.2 Proposition 2: Bounds on α

Next, we assume that $n \in \{1, 2\}$, i.e., we assume that there is at least one abortion ($\alpha_1 > 0$) and at most two abortions.¹⁵ In addition, we assume that the value of the natural sex ratio, p , is known. Using these additional assumptions together with the mean spacing by sex of the next-born, we are able to derive bounds on our model parameters, α_1 and α_2 . These bounds are given in Proposition 2 below and based on two observations. First, if $n \in \{1, 2\}$ (and p is known) then the sex ratio (SR) (in the population) is given by

$$\text{SR} = P(Y = 1) = p + p(1 - p)\alpha_1 + p(1 - p)^2\alpha_1\alpha_2. \quad (6)$$

Solving this equation for α_2 , we obtain the following function of α_1

$$\alpha_2^{\text{SR}}(\alpha_1) = \frac{\text{SR} - p}{p(1 - p)^2\alpha_1} - \frac{1}{1 - p}.$$

Second, as shown in the proof of Proposition 2, the model implies that the difference in mean spacing between girls and boys, $E(T_b|Y = 0) - E(T_b|Y = 1)$, is negative, positive, or zero if and only if α_2 is greater than, less than, or equal to¹⁶

$$\alpha_2^{\text{sign}}(\alpha_1) = \frac{1 - 2(1 - p)\alpha_1 - \sqrt{1 - 4(1 - p)\alpha_1}}{2\alpha_1(1 - p)^2},$$

respectively. Equating $\alpha_2^{\text{SR}}(\alpha_1)$ and $\alpha_2^{\text{sign}}(\alpha_1)$ yields $\alpha_1^* = \frac{(\text{SR}-p)(2p-\text{SR})}{p^2(1-p)}$ and $\alpha_2^* = \frac{\text{SR}-p}{(1-p)(2p-\text{SR})}$. Combining this with the fact that α_1 and α_2 are bounded above by one, we obtain the following Proposition, whose usage is illustrated in the subsequent section.

Proposition 2. *Assume $n \in \{1, 2\}$ and $p \leq 3/4$. If $\alpha_2^* > 1$, then $\alpha_1 \in \left[\frac{\text{SR}-p}{p(1-p)(2-p)}, \min\left\{ \frac{\text{SR}-p}{p(1-p)}, 1 \right\} \right]$ and $\alpha_2 \in [\max\{0, \alpha_2^{\text{SR}}(1)\}, 1]$. If $\alpha_2^* \leq 1$ and*

- (i) $E(T_b|Y = 0) < E(T_b|Y = 1)$, then $\alpha_1 \in \left[\frac{\text{SR}-p}{p(1-p)(2-p)}, \alpha_1^* \right)$ and $\alpha_2 \in [\alpha_2^*, 1]$;
- (ii) $E(T_b|Y = 0) = E(T_b|Y = 1)$, then $\alpha_1 = \alpha_1^*$ and $\alpha_2 = \alpha_2^*$;
- (iii) $E(T_b|Y = 0) > E(T_b|Y = 1)$, then $\alpha_1 \in \left(\alpha_1^*, \frac{\text{SR}-p}{p(1-p)} \right]$ and $\alpha_2 \in [0, \alpha_2^*]$.

3.2.1 Data illustration continued


To take Proposition 2 to the data, we focus on the two groups with unbalanced sex ratios, the group of women that had a girl as a first-born (G) and the group of women that first

¹⁵Assuming a natural sex ratio of 0.513, the probability to be pregnant with a girl three times in a row is 11.6%. Given the limited evidence of repeated sex-selective abortion that we find, the fraction of people that may undergo a third abortion seems negligible.

¹⁶For $\alpha_1 > \frac{1}{4(1-p)}$, $E(T_b|Y = 0) - E(T_b|Y = 1)$ is (strictly) positive for any $\alpha_2 \in [0, 1]$.

had two girls (GG). For the purpose of this exercise, we rely on the results in Table 2 to determine the sign of the difference in mean spacing between girls and boys. We, however, ignore the sampling variation in the observed sex ratio and treat it as the true unknown sex ratio in the population. The construction of “true” confidence sets that take into account all of the sampling variation is left for future research.

Figure 4: Proposition 2



Note: The solid line represents $\alpha_2^{\text{sign}}(\alpha_1)$; the dashed red line and dash-dotted green line represent $\alpha_2^{\text{SR}}(\alpha_1)$ for the G and GG groups, respectively, with the natural sex ratio taken equal to $p=0.513$.

Figure 4 allows us to graphically illustrate the usage of Proposition 2. To the left of the blue line, $\alpha_2^{\text{sign}}(\alpha_1)$, we have $E(T_b|Y=0) < E(T_b|Y=1)$ and to its right we have $E(T_b|Y=0) > E(T_b|Y=1)$, while $E(T_b|Y=0) = E(T_b|Y=1)$ for values on the blue line. The dashed red line and the dash-dotted green line depict $\alpha_2^{\text{SR}}(\alpha_1)$ for the groups G and GG, respectively. Throughout, the natural sex ratio taken equal to 0.513. For the G group, the (observed) sex ratio is 0.535 and the sign of the difference between $E(T_b|Y=0)$ and $E(T_b|Y=1)$ is undetermined since it is not statistically significantly different from 0. Therefore, the set of

values that α_1 and α_2 can take are simply given by the dashed red line with no restriction on α_2 .¹⁷ For the GG group, the sex ratio is 0.555 and $E(T_b|Y = 0) < E(T_b|Y = 1)$ such that the set of values that α_1 and α_2 can take are given by the line segment of dash-dotted green line that is above the solid blue line. Table 3 reports the corresponding “numbers”.

Table 3: Bounds on α_1 and α_2

	α_1	α_2
G	[0.06, 0.089]	[0, 1]
GG	[0.113, 0.154]	[0.183, 1]

G refers to the second-born when the first-born is a girl and GG to the third-born when the first two children are girls.

The bounds on α_1 are relatively tight for both groups, G and GG, and let us conclude that there is more sex-selective abortion, i.e., α_1 is greater, in the group GG than in the group G. The bounds on α_2 are completely uninformative for group G and we cannot exclude the possibility that there is no repeated sex-selective abortion. For group GG, the bounds on α_2 are also very wide but exclude 0 and are, thus, indicative of the presence of repeated sex-selective abortion.

4 Estimation of the model

In the previous section, we derive two propositions that require very little input information as they only exploit mean spacing by sex of the next-born (Proposition 1) and sex ratios (Proposition 2). These propositions also do not require any assumptions on the distribution of T_c and T_w . Alternatively, we can exploit the entire distribution of spacing, by making distributional assumptions on T_c and T_w , to obtain point estimates for our main parameters of interest, the share of women that abort if pregnant with a girl (α_1) and the share of women that abort twice if again pregnant with a girl after a first abortion (α_2).

4.1 Maximum Likelihood

We observe two “outcome” variables: the time to birth, T_b^A , and the sex of the next-born, Y . Let θ denote the parameter vector that subsumes $\alpha = (\alpha_1, \dots, \alpha_{n_0})'$ as well as the parameters

¹⁷Alternatively, one could choose to accept the null hypothesis that $E(T_b|Y = 0)$ and $E(T_b|Y = 1)$ are equal, in which case Proposition 2 would suggest $\alpha_1^* = 0.084$ and $\alpha_2^* = 0.092$ as “estimates” of α_1 and α_2 .

governing the distributions of T_w and T_c , which we denote by $\theta_{-\alpha}$. Making the dependence of probabilities on model parameters explicit,¹⁸ the individual likelihood function can be written as follows

$$\begin{aligned}
L(\theta; t, y) &\equiv P(T_b^A = t, Y = y; \theta) \\
&= \sum_{a=0}^{n_0} P(T_b^a = t, Y = y, A = a; \theta) \\
&= \sum_{a=0}^{n_0} P(T_b^a = t | Y = y, A = a; \theta_{-\alpha}) P(Y = y, A = a; \alpha) \\
&= \sum_{a=0}^{n_0} P(T_b^a = t; \theta_{-\alpha}) P(Y = y, A = a; \alpha) \\
&= \left[\sum_{a=0}^{n_0} P(T_b^a = t; \theta_{-\alpha}) P(Y = 1, A = a; \alpha) \right]^y \\
&\quad \times \left[\sum_{a=0}^{n_0} P(T_b^a = t; \theta_{-\alpha}) P(Y = 0, A = a; \alpha) \right]^{1-y}.
\end{aligned}$$

Here, the fourth line follows from the fact that $P(T_b^a = t | Y = y, A = a; \theta_{-\alpha}) = P(T_b^a = t | Y = y; \theta_{-\alpha}) = P(T_b = t^a; \theta_{-\alpha})$. In words, given the number of abortions, time to birth no longer depends on the sex of the next-born. $P(Y = 1, A = a; \alpha)$ and $P(Y = 0, A = a; \alpha)$ are given in (4) and (5), respectively. Next, we specify $P(T_b^a = t; \theta_{-\alpha})$.

Recall that, by assumption, the randomness in T_b^a stems from T_w and T_c alone. We assume that T_w and T_c are independent and gamma-distributed, with shape parameters γ_w and γ_c and scale parameters β_w and β_c . That is, the pdf of, for example, T_w is given by

$$f_{\Gamma}(t; \gamma_w, \beta_w) = \begin{cases} \frac{t^{\gamma_w-1} e^{-\frac{t}{\beta_w}}}{\beta_w^{\gamma_w} \Gamma(\gamma_w)} & \text{if } t > 0 \\ 0 & \text{otherwise,} \end{cases}$$

where $\Gamma(\cdot)$ denotes the gamma function. This distributional assumption allows us to analytically compute the pdf of T_b^a , $P(T_b^a = t; \theta_{-\alpha})$.¹⁹ In particular, we have

$$P(T_b^a = t; \theta_{-\alpha}) = \begin{cases} f_{\Gamma}(t - 9; \gamma_w, \beta_w) & \text{if } a = 0 \\ f_{\Sigma\Gamma}(t - 9 - t_s a; \gamma_w, \beta_w, \gamma_c a, \beta_c) & \text{if } a > 0, \end{cases}$$

¹⁸For expositional purposes, we assume that T_w and T_c are discrete random variables.

¹⁹The formula for $f_{\Sigma\Gamma}(t; \gamma_1, \beta_1, \gamma_2, \beta_2)$, given in Appendix B, takes the form of an infinite series. In order to numerically evaluate $f_{\Sigma\Gamma}(t; \gamma_1, \beta_1, \gamma_2, \beta_2)$ we rely on a finite series approximation. In particular, we truncate the series at 30 terms. The resulting approximation error is immaterial.

where $f_{\Sigma}(t; \gamma_1, \beta_1, \gamma_2, \beta_2)$ denotes the pdf of the sum of two independent gamma-distributed random variables with shape parameter γ_i and scale parameter β_i for $i \in \{1, 2\}$; see Appendix B for the exact formula, which is due to Moschopoulos (1985). Recall that the *time to pregnancy*, T_p , is taken equal to 9 and that the *time to screening*, T_s , is taken equal to a fixed value, t_s . Note also that here we are using the fact that the sum of two independent gamma-distributed random variables with the same scale parameter, say β , but (possibly) different shape parameters, say γ_1 and γ_2 , is gamma-distributed with shape parameter $\gamma_1 + \gamma_2$ and scale parameter β .

Given the above distributional assumptions, we could in principle estimate θ by maximizing the log-likelihood function, $\sum_{i=1}^n \log L(\theta; t_i, y_i)$, over the parameter space $[0, 1]^2 \times [0, C]^4$ for some large $C < \infty$, where $i = 1, \dots, n$ indexes the households in the estimation sample. Here, in order to “aid” estimation we instead calibrate γ_c and β_c using “external” data.²⁰ In what follows, we further impose that $n_0 = 2$ and, with a slight abuse of notation, let $\theta = (\alpha_1, \alpha, \gamma_w, \beta_w)'$ with the understanding that $\log L(\theta; t_i, y_i)$ denotes $\log L((\theta', \tilde{\gamma}_c, \tilde{\beta}_c)'; t_i, y_i)$, where $\tilde{\gamma}_c$ and $\tilde{\beta}_c$ denote the calibrated values of γ_c and β_c , respectively. Formally, our (unconstrained) “maximum likelihood” estimator is defined as

$$\hat{\theta}_{\text{ML}} = \arg \max_{\theta \in \Theta} \sum_{i=1}^n \log L(\theta; t_i, y_i),$$

where $\Theta = [0, 1]^2 \times [0, C]^2$.

We also consider an additional estimator that maximizes the log-likelihood function subject to an equality constraint. The latter equates the right hand side of equation (6) with the *observed* sex ratio at birth. That is, we consider the following constrained “maximum likelihood” estimator

$$\begin{aligned} \hat{\theta}_{\text{CML}} &= \arg \max_{\theta \in \Theta} \sum_{i=1}^n \log L(\theta; t_i, y_i) \\ \text{s.t. } &p + p(1-p)\alpha_1 + p(1-p)^2\alpha_1\alpha_2 = \bar{y}, \end{aligned} \quad (7)$$

where $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$. Recall that if $\alpha_2 = 0$ then the value of α_1 that solves (7) corresponds

²⁰The calibrated values of γ_c and β_c are obtained by maximum likelihood, fitting a gamma distribution to the time between marriage and first birth (minus 9) for the women in our sample. The underlying assumptions are 1) that once people get married they do not wait before trying to conceive their first child, such that the time between marriage and first birth (minus 9) is a good proxy for the time it takes to conceive 2) there is no sex-selective abortion for the first child. This second assumption is corroborated by the observed sex ratio for this group (see section 2.3.2) as well as the observed difference in mean spacing by gender (see section 3.1.1) and widely accepted in the literature (see for example Dahl and Moretti, 2008; Milazzo, 2018; Heath and Tan, 2018). We discard times greater than 5 years.

to (a one-to-one transformation of) the “standard” method of moments estimator for the number of missing female births.²¹ Therefore, the estimator $\hat{\theta}_{\text{CML}}$ can be interpreted as a method of moments estimator that only relies on a single moment and that uses the maximum likelihood principle to “achieve” identification. In fact, $\hat{\theta}_{\text{CML}}$ can be shown to be equivalent to a particular generalized method of moments (GMM) estimator where the sample moment (equation (7)) gets “infinite” weight and the remaining moments, given by the score, are weighted by the inverse of the observed Fisher information matrix; see Appendix C for details. This equivalence also facilitates the estimation of the asymptotic variance matrix of $\hat{\theta}_{\text{CML}}$; see equation (8) in Appendix C.

The preceding interpretation of $\hat{\theta}_{\text{ML}}$ also motivates its use. A researcher may, for example, strongly believe that certain moments hold in the population, cf. equation (6), but may be less sure of the distributional assumptions she is making and, therefore, prefers to impose the moments in estimation (rather than only in the limit), cf. equation (7), to be, in some sense, less dependent on the distributional assumptions. This (heuristically motivated) gain in robustness comes, of course, at the cost of a loss in efficiency if the model is well specified—in general, $\hat{\theta}_{\text{CML}}$ will not attain the Cramér-Rao lower bound.

A comment on statistical inference based on $\hat{\theta}_{\text{ML}}$ and $\hat{\theta}_{\text{CML}}$ is in order. When reporting confidence intervals and standard errors below, we maintain the assumption that the true value of α_1 , say α_1^* , satisfies $\alpha_1^* \geq c$ for some small $c > 0$. This assumption ensures that the parameter α_2 is (strongly) identified; note that $\sum_{i=1}^n \log L(\theta; t_i, y_i)$ does not depend on α_2 for any θ with $\alpha_1 = 0$ such that α_2 is not identified when $\alpha_1 = 0$ (see e.g., Andrews and Cheng, 2012). Similarly, we assume that $\gamma_w^* \geq c$ and $\beta_w^* \geq c$ which corresponds to standard practice (albeit rarely made explicit), where as above * indicates the true value. Importantly, however, we do not impose any restrictions on the true value of α_2 , α_2^* , beyond the restriction that $\alpha_2^* \in [0, 1]$ when basing inference on $\hat{\theta}_{\text{ML}}$. This implies that α_2^* is allowed to be (near or) at the boundary of the parameter space. It is well-known that “standard” estimators that are restricted to the (true) parameter space are not asymptotically normally distributed when the true parameter is (near or) at the boundary (see e.g., Andrews, 1999; Ketz, 2018). Nevertheless, the standard errors for $\hat{\theta}_{\text{ML}}$ reported below are informative and “valid”, in the sense that they are, in fact, standard errors, for a suitably defined “quasi-unconstrained” estimator (Ketz, 2018). Furthermore, the standard (Wald-type) confidence intervals, reported below, that are based on $\hat{\theta}_{\text{ML}}$ uniformly control asymptotic size; this follows from Corollary 1 in Ketz (2019). When basing inference on $\hat{\theta}_{\text{CML}}$, we implicitly assume that $(\alpha_1^*, \alpha_2^*) \in [c, 1 - c]^2$, i.e., we implicitly assume away any “boundary issues”.

²¹The number of missing female births can be obtained by multiplying α_1 by the number of women pregnant with a girl, i.e., $n \times (1 - p)$.

Extending the notion of a “quasi-unconstrained” estimator to settings where, for example, an equality constraint is maintained is left for future research.

4.2 Results

We estimate our model separately for the two groups for which we found an unbalanced sex ratio, allowing for a maximum of two abortions, i.e., $n_0 = 2$. Table 4 reports the estimates of α_1 and α_2 for the unconstrained and the constrained estimator, introduced in Section 4.1. Here, the natural sex ratio is taken equal to 0.513 and t_s is taken equal to 4; robustness checks for different values of t_s are presented in Appendix D.

Table 4: Estimates

	Sample size	Estimator	α_1	α_2
G	80,369	Unconstrained	0.092 (0.013)	0.002 (0.211)
		Constrained	0.089 –	0.002 –
GG	25,888	Unconstrained	0.130 (0.017)	0.541 (0.248)
		Constrained	0.134 (0.020)	0.530 (0.300)

G refers to the second-born when the first-born is a girl and GG to the third-born when the first two children are girls. Standard errors are reported in parentheses.

The estimates are very similar across estimators within groups, which suggests that the underlying distributional assumptions are reasonable. Among the women that had a girl as a first-born, $\sim 9\%$ are estimated to abort when pregnant with a girl. The estimate of α_2 is essentially zero, suggesting that, in this group, none of the women perform a second (sex-selective) abortion when pregnant again with a girl. In the group of women that first had two girls, $\sim 13\%$ of women are estimated to abort when pregnant with a girl. Furthermore, the estimate of α_2 is (approximately) equal to 54% and statistically significantly different from zero (using a one-sided t-test). This suggests the presence of repeated sex-selective abortion in this group.

5 Conclusion

Measuring how many women abort and if they do so repeatedly is not feasible using “direct” survey or administrative data. We propose an innovative methodology that exploits information contained in birth spacing to measure these quantities. In an application using Indian data, we find evidence of repeated sex-selective abortion for some “groups” of women but not for others. Given that our methodology only requires information on the birth history of women (i.e., sex and birth spacing), which is largely available, our methodology should prove useful in other countries and contexts.

References

- Anderson, Siwan and Debraj Ray**, “Missing Women: Age and Disease,” *The Review of Economic Studies*, 2010, *77* (4), 1262–1300.
- Andrews, Donald W. K.**, “Estimation When a Parameter Is on a Boundary,” *Econometrica*, 1999, *67* (6), 1341–1383.
- **and Xu Cheng**, “Estimation and Inference With Weak, Semi-Strong, and Strong Identification,” *Econometrica*, 2012, *80* (5), 2153–2211.
- Anukriti, S, Sonia R. Bhalotra, and Eddy H. F. Tam**, “On the Quantity and Quality of Girls: Fertility, Parental Investments and Mortality,” *Economic Journal*, 2021. Forthcoming.
- Barclay, Kieron J. and Martin Kolk**, “The Long-Term Cognitive and Socioeconomic Consequences of Birth Intervals: A Within-Family Sibling Comparison Using Swedish Register Data,” *Demography*, 2017, *54* (2), 459–484.
- Bhalotra, Sonia R. and Arthur Van Soest**, “Birth-spacing, fertility and neonatal mortality in India: Dynamics, frailty, and fecundity,” *Journal of Econometrics*, 2008, *143* (2), 274–290.
- **and Tom Cochrane**, “Where have all the young girls gone? Identification of sex selection in India,” *IZA Discussion Paper*, 2010.
- Bongaarts, John and Christophe Z. Guilmoto**, “How Many More Missing Women? Excess Female Mortality and Prenatal Sex Selection, 1970–2050,” *Population and Development Review*, 2015, *41* (2), 241–269.
- Chao, Fengqing, Patrick Gerland, Alex R. Cook, and Leontine Alkema**, “Systematic assessment of the sex ratio at birth for all countries and estimation of national imbalances and regional reference levels,” *Proceedings of the National Academy of Sciences*, 2019, *116* (19), 9303–9311.
- Dahl, Gordon B. and Enrico Moretti**, “The Demand for Sons,” *The Review of Economic Studies*, 2008, *75* (4), 1085–1120.
- Dubuc, Sylvie and David Coleman**, “An increase in the sex ratio of births to India-born mothers in England and Wales: evidence for sex-selective abortion,” *Population and Development Review*, 2007, *33* (2), 383–400.

- Grimes, David A., Janie Benson, Susheela Singh, Mariana Romero, Bela Ganatra, Friday E. Okonofua, and Iqbal H. Shah**, “Unsafe abortion: the preventable pandemic,” *The Lancet*, 2006, *368* (9550), 1908–1919.
- Guilmoto, Christophe Z.**, “Sex Imbalances at Birth: Current trends, consequences and policy implications,” *UNFPA Asia and Pacific Regional Office*, 2012.
- Heath, Rachel and Xu Tan**, “Worth fighting for: Daughters improve their mothers’ autonomy in South Asia,” *Journal of Development Economics*, 2018, *135*, 255–271.
- Jain, Vanita, Subhas C. Saha, Rashmi Bagga, and Sarala Gopalan**, “Unsafe abortion: A neglected tragedy. Review from a tertiary care hospital in India,” *Journal of Obstetrics and Gynaecology Research*, 2004, *30* (3), 197–201.
- Jayachandran, Seema and Ilyana Kuziemko**, “Why do mothers breastfeed girls less than boys? Evidence and implications for child health in India,” *The Quarterly Journal of Economics*, 2011, *126* (3), 1485–1538.
- Jha, Prabhat, Rajesh Kumar, Priya Vasa, Neeraj Dhingra, Deva Thiruchelvam, and Rahim Moineddin**, “Low male-to-female sex ratio of children born in India: national survey of 1.1 million households,” *The Lancet*, 2006, *367* (9506), 211–218.
- Ketz, Philipp**, “Subvector inference when the true parameter vector may be near or at the boundary,” *Journal of Econometrics*, 2018, *207* (2), 285–306.
- , “On asymptotic size distortions in the random coefficients logit model,” *Journal of Econometrics*, 2019, *212* (2), 413–432.
- Klasen, Stephan and Claudia Wink**, ““Missing women”: Revisiting the debate,” *Feminist Economics*, 2003, *9* (2–3), 263–299.
- Milazzo, Annamaria**, “Why are adult women missing? Son preference and maternal survival in India,” *Journal of Development Economics*, 2018, *134*, 467–484.
- Moschopoulos, Peter G.**, “The distribution of the sum of independent gamma random variables,” *Annals of the Institute of Statistical Mathematics*, 1985, *37* (3), 541–544.
- Portner, Claus C.**, “Birth Spacing and Fertility in the Presence of Son Preference and Sex-Selective Abortions: India’s Experience Over Four Decades,” *Demography*, 2020, *conditionnally accepted* (7189).

Schoumaker, Bruno, *Quality and consistency of DHS fertility estimates, 1990 to 2012*, ICF International Rockville, Maryland, 2014.

Sen, Amartya, "Missing women.," *BMJ: British Medical Journal*, 1992, *304* (6827), 587.

A Proofs and additional results

Proof of Proposition 1. To prove (i), note that $T_b = T_b^0$ if $n = 0$. Therefore, if $n = 0$, T_b is independent of Y such that $E(T_b|Y = 0) = E(T_b) = E(T_b|Y = 1)$ and the results follows.

To prove (ii), note that $E(T_b|Y = 0) > E(T_b|Y = 1)$ if $n = 1$. This together with (i) yields the desired result. To see that $E(T_b|Y = 0) > E(T_b|Y = 1)$ if $n = 1$, note that

$$E(T_b|Y = y) = E(T_b^1)P(A = 1|Y = y) + E(T_b^0)(1 - P(A = 1|Y = y))$$

for $y \in \{0, 1\}$. The result then follows by noting that $E(T_b^1) > E(T_b^0)$, since $T_s > 0$ and $T_c \geq 0$, and $P(A = 1|Y = 0) > P(A = 1|Y = 1)$, cf. equations (2) and (3). \square

Next, we show that Proposition 1(ii) continues to hold, under the condition stated in the main text, when we allow T_w to depend on whether a household aborts (or would abort) if pregnant with a female foetus (for the first time between two births). Let $W \in \{0, 1\}$ denote whether a household aborts if pregnant with a female foetus ($W = 1$) or not ($W = 0$). With this extension of the model, we have $P(A = 1, W = 1|Y = y) = P(A = 1|Y = y)$ (as $A = 1 \Rightarrow W = 1$), $P(A = 1, W = 0|Y = y) = 0$ (as $W = 0 \Rightarrow A = 0$), $P(A = 0, W = 0|Y = 0) = P(A = 0|Y = 0) = 1 - P(A = 1|Y = 0)$ (as $\{A = 0\} \cap \{Y = 0\} \Rightarrow W = 0$), $P(A = 0, W = 1|Y = 1) = \frac{\alpha_1}{(1-p)\alpha_1+1}$, and $P(A = 0, W = 0|Y = 1) = \frac{1-\alpha_1}{(1-p)\alpha_1+1}$ for $y \in \{0, 1\}$. In order to allow T_w to depend on W , let $T_w = T_w^W = WT_w^1 + (1 - W)T_w^0$, where T_w^1 (T_w^0) denotes the *waiting time* for a household that aborts (does not abort) if pregnant with a female foetus. Let

$$T_b \equiv T_b^{A,W} = T_w^W + T_c + A(T_c + T_s) + T_p$$

for $A, W \in \{0, 1\}$. Then, the desired result follows as

$$\begin{aligned} E(T_b|Y = 1) &= E(T_b^{1,1})P(A = 1|Y = 1) + E(T_b^{0,1})P(A = 0, W = 1|Y = 1) \\ &\quad + E(T_b^{0,0})P(A = 0, W = 0|Y = 1) \\ &\leq E(T_b^{1,1})P(A = 1|Y = 1) + E(T_b^{0,0})(1 - P(A = 1|Y = 1)) \\ &< E(T_b^{1,1})P(A = 1|Y = 0) + E(T_b^{0,0})(1 - P(A = 1|Y = 0)) = E(T_b|Y = 0), \end{aligned}$$

where the first (weak) inequality holds by assumption, namely $E(T_w^1) \leq E(T_w^0)$ and the second inequality by the same argument used in proving part (ii) of Proposition 1, with $T_b^{1,1}$ ($T_b^{0,0}$) replacing T_b^1 (T_b^0).

Proof of Proposition 2. First, we note that the set of possible or admissible values of α_1 and α_2 is restricted by equation (6) as well as the bounds $0 < \alpha_1 \leq 1$ and $0 \leq \alpha_2 \leq 1$.

The set of admissible values of α_1 (α_2) is, thus, a subset of $A_1 = \{0 < \alpha_1 \leq 1 : \alpha_2 = \alpha_2^{\text{SR}}(\alpha_1) \text{ for some } 0 \leq \alpha_2 \leq 1\}$ ($A_2 = \{0 \leq \alpha_2 \leq 1 : \alpha_2 = \alpha_2^{\text{SR}}(\alpha_1) \text{ for some } 0 < \alpha_1 \leq 1\}$).

To further restrict the set of admissible values of α_1 (α_2), we use the information contained in the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$. Using equations (1) and (5), together with $T_{c,k} \stackrel{i.i.d.}{\sim} T_c$, we get

$$E(T_b|Y = 0) = E(T_w + T_p) + (1 - p)\alpha_1 \frac{(1 - \alpha_2) + 2(1 - p)\alpha_2}{1 - p\alpha_1 - p(1 - p)\alpha_1\alpha_2} E(T_c + T_s).$$

Similarly, using equations (1) and (4), together with $T_{c,k} \stackrel{i.i.d.}{\sim} T_c$, we get

$$E(T_b|Y = 1) = E(T_w + T_p) + (1 - p)\alpha_1 \frac{1 + 2(1 - p)\alpha_2}{1 + (1 - p)\alpha_1 + (1 - p)^2\alpha_1\alpha_2} E(T_c + T_s).$$

Subtracting $E(T_b|Y = 1)$ from $E(T_b|Y = 0)$, we obtain

$$\begin{aligned} E(T_b|Y = 0) - E(T_b|Y = 1) &= \frac{(1 - p)\alpha_1 E(T_c + T_s)}{(1 + (1 - p)\alpha_1 + (1 - p)^2\alpha_1\alpha_2)(1 - p\alpha_1 - p(1 - p)\alpha_1\alpha_2)} \\ &\quad \times ((1 - p)^2\alpha_1\alpha_2^2 + (2(1 - p)\alpha_1 - 1)\alpha_2 + \alpha_1). \end{aligned}$$

As the first factor is (strictly) positive, the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$ only depends on the second factor, which is a polynomial of degree 2 with respect to α_2 . For $\alpha_1 > \frac{1}{4(1-p)}$ no real root exists and the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$ is (strictly) positive for any value of α_2 . For $\alpha_1 \leq \frac{1}{4(1-p)}$ the two real roots are

$$\frac{1 - 2(1 - p)\alpha_1 \pm \sqrt{1 - 4(1 - p)\alpha_1}}{2\alpha_1(1 - p)^2}.$$

However, for $\alpha_1 \leq \frac{1}{4(1-p)}$, we have that

$$\frac{1 - 2(1 - p)\alpha_1 + \sqrt{1 - 4(1 - p)\alpha_1}}{2\alpha_1(1 - p)^2} \geq \frac{1 - 2(1 - p)\alpha_1}{2\alpha_1(1 - p)^2} \geq \frac{1}{1 - p} > 1$$

for any $p \in (0, 1)$ such that only the first root “matters”. Denoting the first root by $\alpha_2^{\text{sign}}(\alpha_1)$ the conclusion in the main text concerning the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$ follows.

Next, note that $\alpha_2^{\text{SR}}(\alpha_1)$ is strictly decreasing (for $\alpha_1 > 0$) and that $\alpha_2^{\text{sign}}(\alpha_1)$ is strictly increasing on its support $(0, \frac{1}{4(1-p)}]$. Furthermore, note that $p \leq 3/4$ implies that $(0, \frac{1}{4(1-p)}] \subseteq (0, 1]$, which ensures that α_1^* , defined in what follows, is less than or equal to 1. The two functions, $\alpha_2^{\text{SR}}(\alpha_1)$ and $\alpha_2^{\text{sign}}(\alpha_1)$, intersect at $\alpha_1^* = \frac{-2(p-SR)(p-\frac{1}{2}SR)}{p^2(1-p)}$ with $\alpha_2^* \equiv \alpha_2^{\text{SR}}(\alpha_1^*) =$

$\alpha_2^{\text{sign}}(\alpha_1^*) = \frac{\text{SR}}{(1-p)(2p-\text{SR})}$ if $\alpha_2^{\text{SR}}(\frac{1}{4(1-p)}) \leq \alpha_2^{\text{sign}}(\frac{1}{4(1-p)}) = \frac{1}{1-p}$, which is equivalent to $\frac{\text{SR}}{p} \leq \frac{3}{2}$.

Noting that $\alpha_2^* > 1$ for $\frac{\text{SR}}{p} > \frac{3}{2}$ (even though α_2^* does not correspond to the point of intersection in this case), we have the following: If $\alpha_2^* > 1$ then $E(T_b|Y = 0) - E(T_b|Y = 1) > 0$ for any $\alpha_1 \in A_1$ ($\alpha_2 \in A_2$) such that no further restriction (using the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$) is possible. Using $\alpha_2^{\text{SR}}(\frac{\text{SR}-p}{p(1-p)(2-p)}) = 1$ and $\alpha_2^{\text{SR}}(\frac{\text{SR}-p}{p(1-p)}) = 0$, the first part of the statement of Proposition 2 follows.

If $\alpha_2^* \leq 1$, then the sign of $E(T_b|Y = 0) - E(T_b|Y = 1)$ allows us to further restrict the set of admissible values of α_1 (α_2): If $E(T_b|Y = 0) - E(T_b|Y = 1) < 0$, then $\alpha_2 \in A_2 \cap \{\alpha_2 : \alpha_2 > \alpha_2^*\} = (\alpha_2^*, 1]$ and $\alpha_1 \in A_1 \cap \{\alpha_1 : \alpha_1 < \alpha_1^*\} = \left[\frac{\text{SR}-p}{p(1-p)(2-p)}, \alpha_1^*\right)$, where we have used $\alpha_2^{\text{SR}}(\frac{\text{SR}-p}{p(1-p)(2-p)}) = 1$ and the fact that $\alpha_2^{\text{SR}}(\alpha_1)$ is strictly decreasing. This proves part (i). If $E(T_b|Y = 0) - E(T_b|Y = 1) = 0$, then trivially $\alpha_1 = \alpha_1^*$ and $\alpha_2 = \alpha_2^*$, which proves part (ii). If $E(T_b|Y = 0) - E(T_b|Y = 1) > 0$, then $\alpha_2 \in A_2 \cap \{\alpha_2 : \alpha_2 < \alpha_2^*\} = [0, \alpha_2^*)$ and $\alpha_1 \in A_1 \cap \{\alpha_1 : \alpha_1 > \alpha_1^*\} = \left(\alpha_1^*, \frac{\text{SR}-p}{p(1-p)}\right]$, where we have used $\alpha_2^{\text{SR}}(\frac{\text{SR}-p}{p(1-p)}) = 0$, $\frac{\text{SR}-p}{p(1-p)} < 1$ if $\alpha_2^* \leq 1$, and the fact that $\alpha_2^{\text{SR}}(\alpha_1)$ is strictly decreasing. This proves (iii) and, thereby, completes the proof of Proposition 2. \square

B Moschopoulos' (1985) formula

Assume without loss of generality that $\beta_1 < \beta_2$. Then,

$$f_{\Sigma\Gamma}(t; \gamma_1, \beta_1, \gamma_2, \beta_2) = \begin{cases} C \sum_{k=0}^{\infty} \delta_k \frac{t^{\gamma_1+\gamma_2+k-1} e^{-\frac{t}{\beta_1}}}{\beta_1^{\gamma_1+\gamma_2+k} \Gamma(\gamma_1+\gamma_2+k)} & \text{if } t > 0 \\ 0 & \text{otherwise,} \end{cases}$$

where $C = \left(\frac{\beta_1}{\beta_2}\right)^{\gamma_2}$ and where δ_k is defined recursively by

$$\delta_{k+1} = \frac{\gamma_2}{k+1} \sum_{i=1}^{k+1} \left(1 - \frac{\beta_1}{\beta_2}\right)^i \delta_{k+1-i}$$

for $k = 0, 1, 2, \dots$ with $\delta_0 = 1$. See Moschopoulos (1985) for more details.

C Maximum likelihood imposing moments

We consider a generic constrained maximum likelihood estimator, $\hat{\theta}_{\text{CML}}$, that satisfies $\hat{\theta}_{\text{CML}} \in \Theta$ and maximizes

$$l(\theta; w) \quad \text{s.t.} \quad c(\theta; w) = 0,$$

where $l(\theta; w) = \sum_{i=1}^n l(\theta; w_i) = \sum_{i=1}^n \log f(w_i; \theta)$ and where $c(\theta; w) = c_1(\theta) - c_2(w)$. Here, $c_1(\theta)$ denotes a known d_c -dimensional function of θ and $c_2(w)$ is $(d_c \times 1)$ vector of “sample moments”, e.g., $c_2(w) = \bar{w}_1$. Throughout, we assume that $d_\theta > d_c$, where d_θ is the dimension of θ . Let

$$s(\theta; w) = \frac{\partial l(\theta; w)}{\partial \theta}$$

and

$$C(\theta) = \frac{\partial c(\theta; w)}{\partial \theta'} = \frac{\partial c_1(\theta)}{\partial \theta'}.$$

Then, $\hat{\theta}_{\text{CML}}$ satisfies

$$s(\hat{\theta}_{\text{CML}}; w) = C'(\hat{\theta}_{\text{CML}})\hat{\lambda},$$

where $\hat{\lambda}$ denotes the $(d_c \times 1)$ vector of (“estimated”) Lagrange multipliers.

Now, consider the following GMM estimator

$$\hat{\theta}_{\text{GMM}} = \arg \min_{\theta \in \Theta} g(\theta)'WG(\theta),$$

where

$$g(\theta) = \begin{bmatrix} s(\theta; w) \\ c(\theta; w) \end{bmatrix}$$

and

$$W = \begin{bmatrix} \left(\frac{\partial^2 l(\hat{\theta}_{\text{CML}}; w)}{\partial \theta' \partial \theta} \right)^{-1} & 0 \\ 0 & \text{diag}(|\hat{\lambda}|)M \end{bmatrix}.$$

Here, $\frac{\partial^2 l(\hat{\theta}_{\text{CML}}; w)}{\partial \theta' \partial \theta}$ is shorthand for $\frac{\partial^2 l(\theta; w)}{\partial \theta' \partial \theta} \Big|_{\theta = \hat{\theta}_{\text{CML}}}$ and M denotes a large constant, whose role will be made clear below. Let

$$G(\theta) = \begin{bmatrix} \frac{\partial s(\theta; w)}{\partial \theta'} \\ C(\theta) \end{bmatrix} = \begin{bmatrix} \frac{\partial^2 l(\theta; w)}{\partial \theta' \partial \theta} \\ C(\theta) \end{bmatrix}$$

and let c^* be a vector of 1s and -1s such that $|\hat{\lambda}| \circ c^* = -\hat{\lambda}$, where \circ denotes element-by-element multiplication operator and where we suppress the dependence of c^* on $\hat{\lambda}$ for notational ease. Then, the first order conditions (FOCs) of the GMM estimation problem, $G'(\hat{\theta}_{\text{GMM}})Wg(\hat{\theta}_{\text{GMM}}) = 0$, can be written as

$$\frac{\partial^2 l(\hat{\theta}_{\text{GMM}}; w)}{\partial \theta' \partial \theta} \left(\frac{\partial^2 l(\hat{\theta}_{\text{CML}}; w)}{\partial \theta' \partial \theta} \right)^{-1} s(\hat{\theta}_{\text{GMM}}; w) + C'(\hat{\theta}_{\text{GMM}}) \left(\hat{\lambda} + |\hat{\lambda}| \circ (c(\hat{\theta}_{\text{GMM}}; w)M - c^*) \right) = 0.$$

Note that the FOCs require that $c(\hat{\theta}_{\text{GMM}}; w) \rightarrow 0$ as $M \rightarrow \infty$ such that θ_{CML} approximately

satisfies the FOCs for M large enough (in the sense that $\|\hat{\theta}_{\text{GMM}} - \hat{\theta}_{\text{CML}}\| \rightarrow 0$ as $M \rightarrow \infty$). In practice, it suffices to choose M large enough such that $\hat{\theta}_{\text{GMM}}$ and $\hat{\theta}_{\text{CML}}$ are equal up to some small numerical error.²²

The (near-) numerical equivalence of $\hat{\theta}_{\text{GMM}}$ and $\hat{\theta}_{\text{CML}}$ (for M large enough) is useful for two reasons. First, it provides us with an alternative interpretation for the constrained maximum likelihood estimator. That is, $\hat{\theta}_{\text{CML}}$ can be interpreted as an “augmented” method of moments estimator, where a set of d_c moments, $c(\theta; y)$, is “augmented” by maximum likelihood to achieve identification. The equivalence with $\hat{\theta}_{\text{GMM}}$ reveals that the moments $c(\theta; w)$ get “infinite” weights and that the remaining moments $s(\theta; w)$ are weighted by the inverse of the observed Fisher information matrix, $\left(\frac{\partial^2 l(\hat{\theta}_{\text{CML}}; w)}{\partial \theta' \partial \theta}\right)^{-1}$, i.e., they are weighted relative to their “information”.

Second, this equivalence provides us with an easy formula for the estimator of the asymptotic variance of $\hat{\theta}_{\text{CML}}$, which is given by the standard sandwich formula for GMM estimators, namely

$$\left(G'(\hat{\theta}_{\text{CML}})WG(\hat{\theta}_{\text{CML}})\right)^{-1} G'(\hat{\theta}_{\text{CML}})W\hat{\Omega}WG(\hat{\theta}_{\text{CML}}) \left(G'(\hat{\theta}_{\text{CML}})WG(\hat{\theta}_{\text{CML}})\right)^{-1}, \quad (8)$$

where $\hat{\Omega}$ is a consistent estimator of $E[g(\theta^*)g'(\theta^*)]$ divided by the sample size, where θ^* denotes the true value of θ .

²²In Matlab, for example, it suffices to choose M large enough for the “initial point [to be] a local minimum” if $\hat{\theta}_{\text{CML}}$ is used as starting value of the GMM estimation problem.

D Robustness checks

Table 5: Estimates for different values of t_s

	(1)	(2)	(3)	(4)		
	Estimator	t_s	α_1	α_2		
G	Unconstrained	4	0.092	0.002		
		5	0.092	0.002		
		6	0.090	0.002		
		7	0.086	0.007		
		Constrained	4	0.089	0.002	
			5	0.089	0.002	
			6	0.089	0.002	
	7		0.089	0.007		
	GG	Unconstrained	4	0.130	0.582	
			5	0.130	0.541	
			6	0.125	0.598	
			7	0.115	0.735	
			Constrained	4	0.131	0.578
				5	0.134	0.530
6				0.131	0.574	
7		0.126		0.684		

G refers to the second-born when the first-born is a girl and GG to the third-born when the first two children are girls.