

Further Evidence of Within-Marriage Fertility Control in Pre-Transitional England

Cinnirella, Francesco; Klemp, Marc; Weisdorf, Jacob Louis

Published in:
Demography

DOI:
[10.1007/s13524-019-00787-1](https://doi.org/10.1007/s13524-019-00787-1)

Publication date:
2019

Document version
Accepted manuscript

Citation for pulished version (APA):
Cinnirella, F., Klemp, M., & Weisdorf, J. L. (2019). Further Evidence of Within-Marriage Fertility Control in Pre-Transitional England. *Demography*, 56(4), 1557-1572. <https://doi.org/10.1007/s13524-019-00787-1>

Terms of use

This work is brought to you by the University of Southern Denmark through the SDU Research Portal. Unless otherwise specified it has been shared according to the terms for self-archiving. If no other license is stated, these terms apply:

- You may download this work for personal use only.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying this open access version

If you believe that this document breaches copyright please contact us providing details and we will investigate your claim. Please direct all enquiries to puresupport@bib.sdu.dk

Further Evidence of Within-Marriage Fertility Control in Pre-Transitional England

Francesco Cinnirella¹

Marc Klemp²

Jacob Weisdorf³

March 12, 2019

Abstract

The identification of parity effects on the hazard of a next birth in cross-family data requires accounting for heterogeneity in fecundity across couples. Stratifying duration models on the maternal level for this purpose, Cinnirella, Klemp, and Weisdorf (2017) shows that the hazard of a next birth decreases with rising parity in historical England. Clark and Cummins (2019) takes issue with this finding, claiming that the result is a statistical artefact caused by stratifying on the maternal level while controlling for the mother's age at the beginning of her birth intervals. This paper documents that Cinnirella, Klemp, and Weisdorf's finding is robust to dealing with Clark and Cummins' critique.

¹ University of Southern Denmark, D-IAS, CESifo, and CEPR.

² University of Copenhagen, Brown University, and CEPR.

³ University of Southern Denmark, Scuola Superiore Sant'Anna Pisa, and CEPR.

1. Introduction

Traditional scholarship has suggested that marital birth control was uncommon before the demographic transition. Cinnirella, Klemp, and Weisdorf (2017) contests this view, showing that couples in pre-transitional England regulated their birth intervals in response to real-wage changes (parity-*independent* birth control) and to the number of surviving offspring (parity-*dependent* birth control). Their results are based on duration analysis carried out on individual-level statistics from the Cambridge Group's *Family Reconstitution* data covering the period 1540 to 1850.

Clark and Cummins (2019) suggests that Cinnirella, Klemp, and Weisdorf's estimates of parity-*dependent* birth control are biased. Using simulated data, they assert that the estimated coefficients of net parity are biased in models that control for the mother's age at the beginning of her birth intervals and, at the same time, stratify on the level of the mother. They also argue (without empirical substance) that Cinnirella, Klemp, and Weisdorf ignore important factors in their analysis of parity-*independent* birth control, claiming that the results could be "potentially explained as the product just of nutrition, living environments, and social patterns" (*ibid.*, 2019, p. 2). Based on this reasoning, Clark and Cummins declare that "the traditional view of an absence of any evidence of birth control through spacing is vindicated" (*ibid.*, 2019, p. 2).

We are grateful to Clark and Cummins for replicating our earlier estimates and for giving us the opportunity to further examine the existence of birth control in pre-transitional England while accounting for their criticism. The analysis presented below establishes that Clark and Cummins' critique has no bearing on the conclusions presented in Cinnirella, Klemp, and Weisdorf (2017). In particular, the evidence of parity-*dependent* birth control is remarkably robust to the use of a range of identification strategies *not* subject to Clark and Cummins' concerns. While the magnitude of the estimates presented in this article depends on the particular method used to account for the heterogeneity in fecundity across households, the conclusion that within-marriage birth control existed in pre-transitional England is incontestable. We also take the opportunity here to recapitulate the empirical strategies and key control variables used in Cinnirella, Klemp, and Weisdorf (2017) to mitigate any potentially confounding factors in the analysis of parity-*independent* birth control. The present study thus solidifies the conclusion that pre-transitional couples engaged in parity-*dependent* as well as parity-*independent* birth control.

In particular, we estimate the effect of net parity on the hazard of a subsequent birth by using three different model-specifications, which are not affected by Clark and Cummins' concerns. That is, (i) we estimate the relationship between *relative* net parity and birth spacing; (ii) we account for between-family differences in fecundity, stratifying by the *protogenetic interval*, i.e., the time-span between the couple's marriage and their first birth; and (iii) we use duration models that account for family-level shared frailty. While Clark and Cummins' culprits for our alleged statistical artefacts are eliminated, we still find statistically-significant, negative effects of net parity on the hazard of a next birth, further strengthening the evidence for the existence of pre-transitional marital birth control.

We proceed as follows: Section 2 summarizes the methodological issues; Section 3 presents the results using each of the three strategies listed above to identify parity-*dependent* effects; Section 4 clarifies important features with regards to our parity-*independent* results; and Section 5 concludes.

2. Methodological issues

As explained in Cinnirella, Klemp, and Weisdorf (2017, p. 415), more fecund couples are more likely to be over-represented at higher parities than less fecund couples. Unobserved heterogeneity in couples' fecundity (and other co-determinants of larger family size and shorter birth intervals) may thus generate an upward omitted-variable bias in the estimates of the association between parity and the hazard of a next birth. This upward bias may obscure parity-dependent fertility limitation in cross-family estimations when heterogeneity in couples' fecundity is not accounted for.

We use Cox proportional hazard models, stratified on the level of the mother, in order to account for unobserved heterogeneity in couples' fecundity (i.e., maternal-level fixed effects). Moreover, in order to control for the fact that maternal fertility is affected by age, we also control for the mother's age at the beginning of each of her birth intervals. Clark and Cummins, however, assert that the dual control for mother-fixed effects and maternal age renders the estimates of net-parity coefficients downward biased. Based on this, they claim that "[t]he parity effects are [...] an artefact of the bizarre fit of the age controls to the data" (Clark and Cummins, 2019, p. 3). Importantly, Clark and Cummins' assertion is not demonstrated in the context of a formal statistical framework. It is loosely inferred from the outcome of a trivial simulation exercise. This implies that the size of a potential bias, as well as the

conditions necessary for its existence, are not clearly identified. However, their concern is nevertheless potentially valid and, as such, deserves a proper consideration.

In the next section, we propose three alternative strategies, each of which accounts for the criticism raised by Clark and Cummins and thus serves to examine the robustness of the results presented in Cinnirella, Klemp, and Weisdorf (2017). Indeed, we document in the following that net-parity effects are evident also when we do *not* stratify on the level of the mother. The sensitivity checks establish that the finding of birth-control effects before the demographic transition is even more robust to the use of different statistical specifications than already demonstrated in Cinnirella, Klemp, and Weisdorf (2017).

3. Alternative models

This section presents further evidence corroborating the existence of net-parity effects on the hazard of a subsequent birth in historical England. As discussed above, it is crucial, when identifying deliberate birth-limitation through net-parity effects, to account for heterogeneity in couples' fecundity. Clark and Cummins claim that stratification on the level of the mother alongside the inclusion of control variables for maternal age at the beginning of the mother's birth intervals prevent the production of "a reliable causal parity effect estimate" (Clark and Cummins, 2019, p. 9). Below, we propose three alternative models that account for heterogeneity in fecundity between couples, but without stratifying on the level of the mother.

a. Models with relative birth order

The first specification assesses the correlations between birth spacing, parity, and family size. It demonstrates that net-parity coefficients smaller than one do not depend on the simultaneous inclusion of mother-fixed effects and controls for mother's age at childbirth. In particular, we account for heterogeneity in couples' fecundity by focusing on the *relative net parity*. By focusing on measures that capture the *relative* position of births among all family births, we provide a first attempt to disentangle family size from birth order, implicitly accounting for the unobserved heterogeneity in fecundity across families, but *without* including mother-fixed effects.

The relative net-parity measure is defined as $(n-1)/(N-1)$, where n is the net parity (i.e. the birth order of the surviving children) of the child that marks the beginning of the birth interval, and where N is the total number of births in the family (see Ejrnæs and Pörtner, 2004). Using this measure, the first surviving child has a

relative parity equal to zero, and the last surviving child a relative net parity equal to one. We also generate a series of dummy variables indicating the quartiles of the relative net parity measure in order to assess the impact of different values of relative net parity on birth intervals. The estimates of the duration models using the relative net parity are presented in Table 1. In column 1, we show the estimates of using a continuous measure for relative net parity. The coefficient for relative net parity establishes that higher relative net parity is highly-significantly associated with a lower hazard of a subsequent child birth. In column 2, we include indicators for quartiles of the measure for net relative parity. The gradient of the coefficients documents that the association between relative net parity and the spacing of births is monotonic: higher parities are related to a lower risk of having a next birth.

The estimated relationship between relative net parity and birth intervals could be biased by the inclusion of the total household fertility in the relative birth order index since total fertility is potentially endogenous.⁴ Indeed, *ceteris paribus*, shorter birth intervals are expected to increase the final family size, resulting in a reverse-causality problem and thus potentially generating biased estimates. We attempt to shed light on this issue by controlling for net parity and family size separately. The results of this specification are presented in column 3. The coefficients for net parity are very similar to the relative coefficients reported in column 1, while total household fertility is associated with shorter birth intervals. However, an endogeneity bias cannot be ruled out in the context of the present empirical strategy. Our empirical strategy further below is therefore designed to avoid potential issues of endogeneity entirely.

Finally, as discussed in detail in Cinnirella, Klemp, and Weisdorf (2017, p. 430-432), households drop out of the sample when mothers leave the parish in which they are observed. If there is a systematic relationship between birth spacing and sample attrition, then this could potentially bias our estimates. In particular, if longer birth intervals at higher parities are underrepresented because of sample attrition, then we might be underestimating the existence of parity-dependent birth control. In order to avoid this issue, we constrain the sampled households to *completed* marriages. This restricts the sample to marriages in which both spouses survive (in marriage) until the

⁴ The results presented in this subsection are robust to the use of a *predicted* measure of the final family size calculated for each birth interval by dividing the time remaining until the age of onset of age-related sterility with the average length of birth intervals. Using the same constant values for age-related sterility and average birth-interval length for all families (e.g., sample averages), such a predicted measure is not affected by concerns of endogeneity (related to, e.g., an effect of birth spacing on final family size) which may potentially affect the observed measures.

wife turns 50, thus ending her reproductive period. Estimates of relative net-parity effects among *completed* marriages are reported in column 4. Indeed, the coefficients for relative net parity decrease in size compared to the estimates based on the full sample, suggesting that sample attrition might affect net parity coefficient estimates.

Table 1: Parity-dependent fertility control – Relative net parity

	(1)	(2)	(3)	(4) Completed families
Relative net parity	0.420*** (0.012)			
Relative net parity (2 nd quartile)		0.824*** (0.009)		0.722*** (0.016)
Relative net parity (3 rd quartile)		0.732*** (0.010)		0.630*** (0.017)
Relative net parity (4 th quartile)		0.629*** (0.013)		0.513*** (0.021)
Net parity (2 nd quartile)			0.799*** (0.009)	
Net parity (3 rd quartile)			0.681*** (0.009)	
Net parity (4 th quartile)			0.565*** (0.011)	
Total household fertility			1.115*** (0.003)	
Real wage	1.095*** (0.008)	1.094*** (0.008)	1.098*** (0.008)	1.070*** (0.015)
Mother's age at birth	0.910*** (0.006)	0.897*** (0.006)	0.945*** (0.007)	0.875*** (0.013)
Mother's age at birth (squared)	1.002*** (0.000)	1.002*** (0.000)	1.001*** (0.000)	1.002*** (0.000)
Infant death	1.984*** (0.030)	1.903*** (0.029)	1.701*** (0.024)	1.983*** (0.057)
Last birth interval	0.691*** (0.008)	0.662*** (0.007)	0.728*** (0.008)	0.524*** (0.012)
Occupation	Yes	Yes	Yes	Yes
Subjects	71086	71086	71086	19579

Note: Cox proportional hazard model with time-varying real wages. Hazard ratios are reported. Real-wages are standardized with zero mean and unit standard deviation. Mother's age is measured at the beginning of the interval and varies within the family. Models are stratified by quarter century and parish. Standard errors, shown in parentheses, are clustered by household. * p<0.10, ** p<0.05, *** p<0.01.

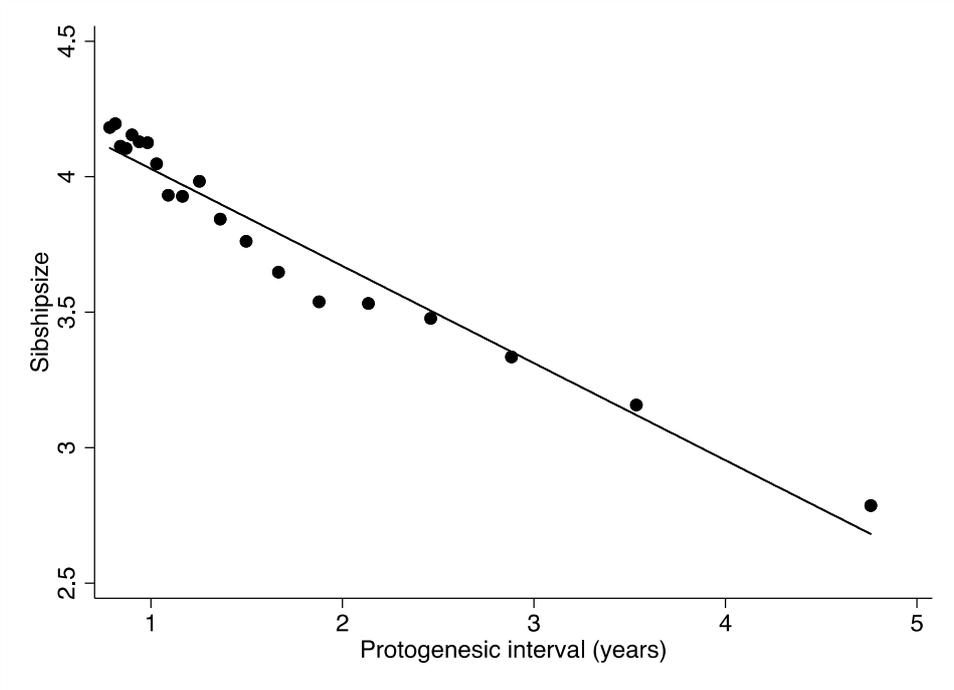
The concept of *relative* net parity helps us to disentangle parity progression from family size. The relative net parity specifications demonstrate that the observed net-parity effects are *neither* due to the use of mother-fixed effects, *nor* to some “impossible” mother's age effect, as Clark and Cummins imagine. However, the use of relative net parity, or the inclusion of family size, may generate an endogeneity bias of our coefficients of interest, as mentioned above. In what follows, we thus propose a set

of alternative specifications, which account, in various ways, for unobserved heterogeneity while simultaneously avoiding issues of endogeneity.

b. Models with the protogenesic interval

Our first alternative model, designed to estimate the effect of parity on birth spacing, makes use of the *protogenesic interval* as a proxy for a couple’s fecundity. The protogenesic interval captures the time elapsed between a couple’s marriage and their first birth. As detailed in Klemp and Weisdorf (2019) and Galor and Klemp (2019), longer protogenesic intervals are indicative of lower levels of parental fecundity among parents who did not conceive before marriage. Moreover, the protogenesic interval is highly-significantly correlated with the length of subsequent birth intervals in historical populations (Klemp and Weisdorf 2019; Galor and Klemp 2019). Therefore, couples with shorter protogenesic intervals and hence higher levels of fecundity are able, *ceteris paribus*, to reach higher parities within a given time frame, as reflected by the fact that the protogenesic interval is a strong predictor of completed family size among historical populations (*ibid.*). The negative association between the protogenesic interval and family size is illustrated in Figure 1, using the Cambridge Group’s *family reconstitution* data.

Figure 1: Relationship between total family fertility and protogenesic interval



Note: Binned scatterplot of family size and protogenesic intervals between 40 weeks and five years.

Moreover, studies of modern data have suggested that the “time to pregnancy” – a concept closely related to the protogenesic interval – is not systematically related to unobserved potentially confounding factors (Klemp and Weisdorf 2019). Likewise, Klemp and Weisdorf (2019) have documented, on the basis of the current historical dataset, that the protogenesic interval is not statistically-significantly associated with otherwise potentially confounding factors, such as parental socio-economic characteristics including occupational skills and wealth. The protogenesic interval does contain a random component, which means that the interval is not a perfect proxy for fecundity. But controlling for the length of the protogenesic interval still serves to mitigate any omitted-variable bias associated with between-family heterogeneity in fecundity.

To this end, we restrict the sample to couples whose protogenesic intervals are at least 40 weeks long (i.e. the first child was conceived after marriage) and less than five years long (since long protogenesic intervals could be indicative of an unobserved first birth). Next, we group couples into deciles of the distribution of the protogenesic intervals and stratify our baseline Cox proportional hazard model on these groups. Importantly, as per Clark and Cummins’ critique, we do *not* stratify on the level of the mother. This specification then allows couples with different protogenesic intervals (and thus different levels of fecundity) to have different baseline hazards. In light of the plausible existence of an upward bias in the estimation of the effect of parity on the hazard of a birth in models that do not account for heterogeneity in couple fecundity, as described in Section 2, our prior is that the estimated effect of net parity on the hazard of a next birth is consistently *lower* than when controlling for couples’ protogenesic intervals.

The estimates of the duration models are detailed in Table 2. Column 1 reports the baseline coefficient estimate on the net-parity variable when the model is not stratified by the protogenesic interval and therefore does not account for heterogeneity in fecundity. The specification does stratify by parish and quarter-century and also includes a range of baseline control variables.⁵ The coefficient estimate for net parity is not significantly different from one. In column 2, we show the results of estimating the same model, but this time with dummies for net parity instead. The coefficients for net parity initially decrease, but eventually increase with higher parities. While the initial decrease is consistent with parity-dependent birth control, the subsequent

⁵ These include indicators for the occupation of the father in order to account for additional sources of relevant heterogeneity between families. Our results do not depend on the inclusion of these variables, however.

increase suggests the existence of omitted-variable bias. It concurs with the notion that the omitted-variable bias associated with *not* accounting for between-family heterogeneity in fecundity can obscure a true underlying negative effect of parity on the subsequent hazard of a birth, and is consistent with the findings in Cinnirella, Klemp, and Weisdorf (2017, Table 4, column 1).

One of the central concerns of Clark and Cummins is related to the estimated relationship between maternal age and birth spacing in the models that are stratified on the level of the mother. The coefficients for the quadratic polynomial in mother's age in columns 1 and 2 of Table 2 suggest a convex relationship between age and the hazard of a birth. In Figure 2, we report the predicted marginal effects of maternal age on the hazard ratio for the specification without stratification by protogenetic intervals (Models 1 and 2 in Figure 2).⁶ The observed relationship is consistent with the notion that fecundity declines with maternal age. The inclusion of the control for the last birth interval generates a slight upward trend in the quadratic fit at the highest maternal ages. This is revealed in Model 3 of Figure 2, where we report the marginal effects of mother's age at childbirth for a specification without a control for the last birth interval.

In column 3 of Table 2, we use a more flexible non-parametric specification, which includes dummy variables for mother's age at birth in five-year intervals. The coefficients for the dummy variables lead to the same conclusion, namely that birth hazards decrease with maternal age. As our new analysis establishes, there is nothing "impossible" about the estimated relationship between age and the hazard of a successive birth.

In columns 4 and 5 of Table 2, we show the estimates for models that are stratified by the protogenetic interval. A single net-parity variable is used in column 4, and the net-parity-indicator variables are used in column 5. The estimates establish that the coefficient for net parity is significantly less than one, once we account for differences in fecundity proxied by the protogenetic interval. This reveals that higher parities are associated with longer birth intervals and thus lower fertility (column 4), a clear sign of parity-dependent fertility control. Likewise, when using dummy variables for net parity, the coefficients tend to decrease with parity and stabilize at values below one (column 5).

Figure 3 reports the marginal effects of mother's age at childbirth for different specifications with stratification by protogenetic intervals. The graphs report a total of

⁶ The marginal effect of mother's age is computed for net parity 3, no mortality of the previous child, no last birth interval, occupation craftsmen, and not born on January 1st, January 11th, not December 25th.

four models: (i) the baseline specification of column 4 with a continuous variable for net parity (Model 1), (ii) the specification of column 5 with net parity dummies (Model 2), (iii) a modified version of the latter that omits a dummy for the last birth interval (Model 3), and (iv) a version that includes neither controls for net parity nor for the last interval (Model 4). Figure 3 shows that the hazard of a next birth declines with maternal age regardless of the specification. As for the case without stratification on the protogenesic intervals, Model 3 shows that it is the inclusion of the last birth-interval dummy which generates a slight upwards slope in the quadratic fit towards the end of the reproductive period. Importantly, Model 4, by excluding both controls for net parity and last birth interval establishes that the inclusion of the parity variable in the other models *does not* qualitatively affect the shape of the quadratic fit.

Since the use of a quadratic fit may be overly restrictive, we also exploit a more flexible specification with dummy variables for maternal age (column 6, Table 2). We find that the risk of a next birth decreases with net parity independently of the negative effect of age on birth risk, which strongly indicates the existence of parity-dependent birth control. Figure 4 reports the coefficients for the dummies of mother's age at childbirth estimated in columns 3 and 6 of Table 2. The maternal age-fecundity relationship estimated in models which stratify by the protogenesic interval does not present any "impossibility" as shown in Figures 3 and 4.

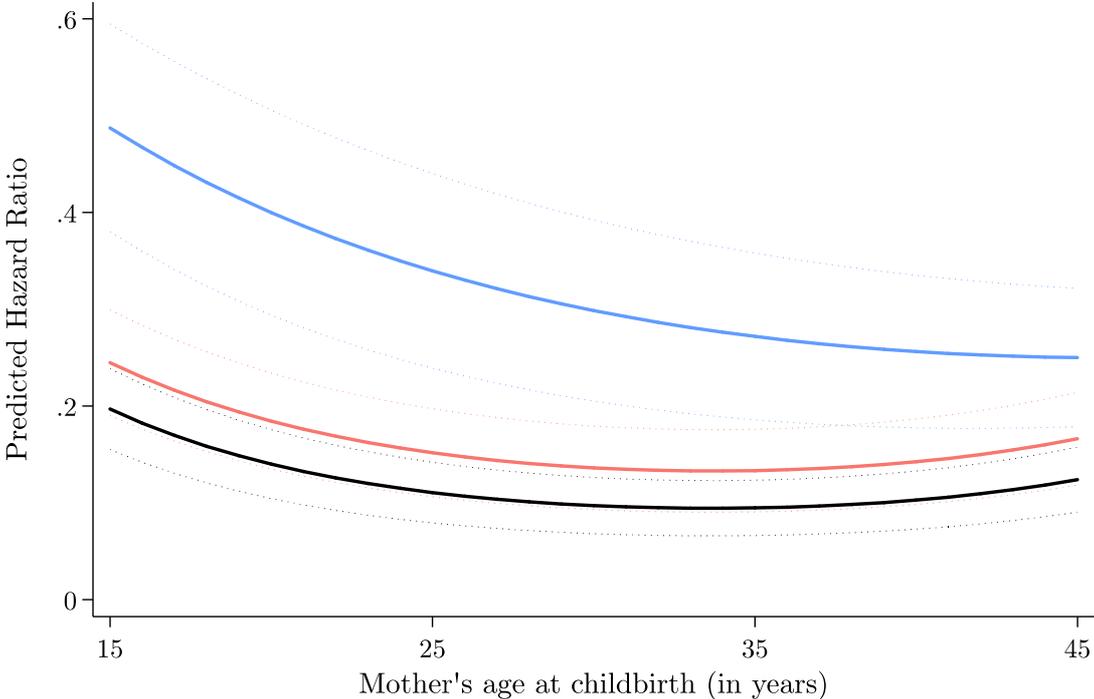
Finally, Figure 5 provides a graphical representation of the coefficients for net parity reported in columns 3 and 6 of Table 2. The graph shows that the fecundity-adjusted coefficients for net parity are systematically smaller, as expected, compared to the unadjusted coefficients, and also that they are systematically below one. This set of results suggests (i) that accounting for fecundity mitigates the omitted-variable bias associated with between-family heterogeneity, and (ii) that net parity has a negative effect on the hazard of a next birth. Again, these results on net parity are not obtained *because of or in the presence of* an "impossible" age-fecundity relationship, as claimed by Clark and Cummins. These (new) results support the notion of a parity-dependent birth-spacing effect and, by implication, are inconsistent with Clark and Cummins' notion of random pre-transitional marital fertility.

Table 2: Parity-dependent birth spacing accounting for the protogenesic interval

	Not stratified by protogenesic interval group			Stratified by protogenesic interval group		
	(1)	(2)	(3)	(4)	(5)	(6)
Net parity	1.001 (0.005)			0.975*** (0.005)		
Net parity 2		0.896*** (0.014)	0.888*** (0.014)		0.856*** (0.013)	0.850*** (0.013)
Net parity 3		0.885*** (0.016)	0.879*** (0.016)		0.818*** (0.015)	0.813*** (0.015)
Net parity 4		0.921*** (0.019)	0.916*** (0.019)		0.841*** (0.018)	0.838*** (0.018)
Net parity 5		0.928*** (0.023)	0.930*** (0.022)		0.824*** (0.022)	0.827*** (0.022)
Net parity 6		0.985 (0.029)	0.995 (0.029)		0.869*** (0.028)	0.878*** (0.028)
Net parity 7		1.058 (0.043)	1.082** (0.043)		0.911** (0.040)	0.931* (0.040)
Net parity 8		1.117* (0.067)	1.151** (0.068)		0.938 (0.059)	0.963 (0.060)
Net parity 9		0.977 (0.092)	1.023 (0.095)		0.839* (0.076)	0.875 (0.079)
Net parity 10		1.138 (0.143)	1.185 (0.144)		0.863 (0.114)	0.895 (0.115)
Mother's age at birth	0.868*** (0.008)	0.889*** (0.009)		0.869*** (0.009)	0.897*** (0.009)	
Mother's age at birth sq.	1.002*** (0.000)	1.002*** (0.000)		1.002*** (0.000)	1.002*** (0.000)	
Mother's age (15-19)			1.129*** (0.049)			1.127*** (0.050)
Mother's age (25-29)			0.911*** (0.016)			0.921*** (0.017)
Mother's age (30-34)			0.866*** (0.017)			0.879*** (0.019)
Mother's age (35-39)			0.907*** (0.021)			0.933*** (0.023)
Mother's age (40-44)			0.912*** (0.031)			0.948 (0.034)
Real wage	1.089*** (0.011)	1.088*** (0.011)	1.089*** (0.011)	1.094*** (0.012)	1.093*** (0.012)	1.093*** (0.012)
Infant death	1.733*** (0.033)	1.731*** (0.033)	1.734*** (0.033)	1.847*** (0.036)	1.848*** (0.036)	1.849*** (0.036)
Last birth interval	0.577*** (0.008)	0.575*** (0.008)	0.584*** (0.008)	0.573*** (0.008)	0.573*** (0.008)	0.581*** (0.008)
Occupation	Yes	Yes	Yes	Yes	Yes	Yes
Subjects	37208	37208	37208	37208	37208	37208

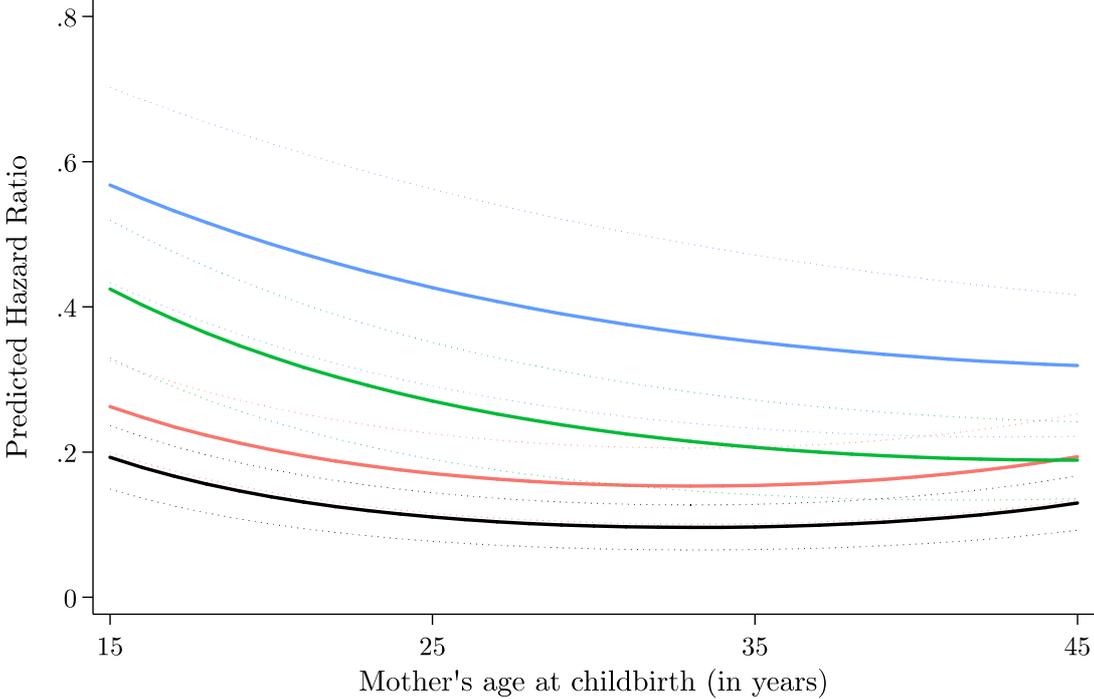
Note: Cox proportional hazard model with time-varying real wages. Hazard ratios are reported. Real-wages are standardized with zero mean and unit standard deviation. Mother's age is measured at the beginning of the interval and varies within the family. Sample constrained to households with protogenesic interval between 40 weeks and five years. Models in column 1 and 2 are stratified by parish and quarter century. Models in columns 3 and 4 are additionally stratified by deciles of the protogenesic interval. Standard errors, shown in parentheses, are clustered by household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 2: Marginal effects of mother's age at childbirth without stratification by protogenesic intervals.



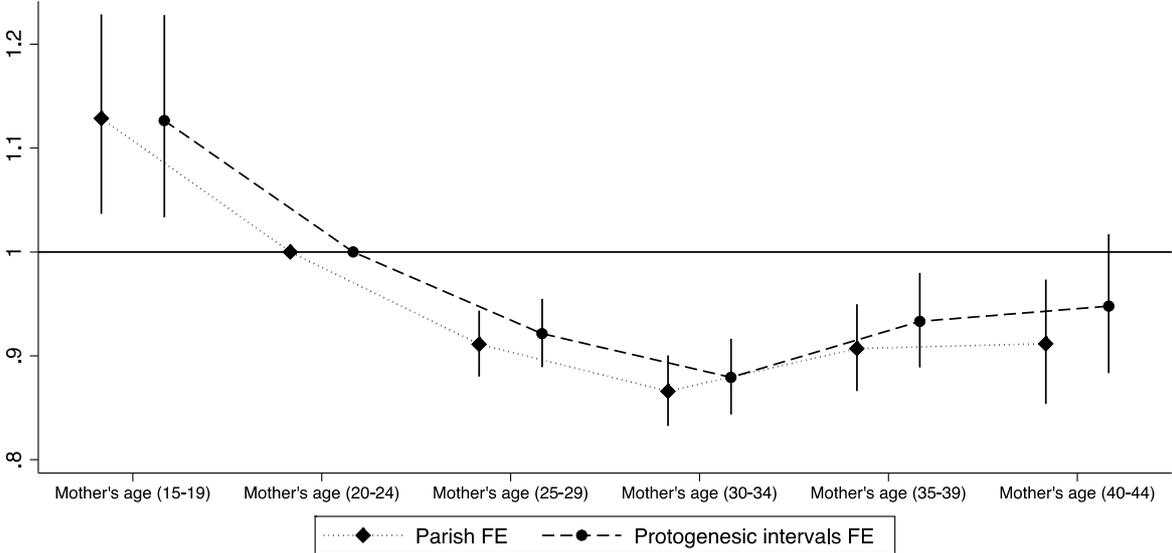
Note: The marginal effects of mother's age (solid lines) and 95% confidence bands (dotted lines) are computed for net parity 3, with no mortality of the previous child, no last birth interval, fathers' occupation craftsmen, and not born on January 1st, January 11th, nor on December 25th. Model 1 (black): with continuous variable for net parity. Model 2 (orange): with net parity dummies. Model 3 (blue): with net parity dummies, without control for last birth interval.

Figure 3: Marginal effect of mother's age at childbirth with stratification by protogenesic intervals.



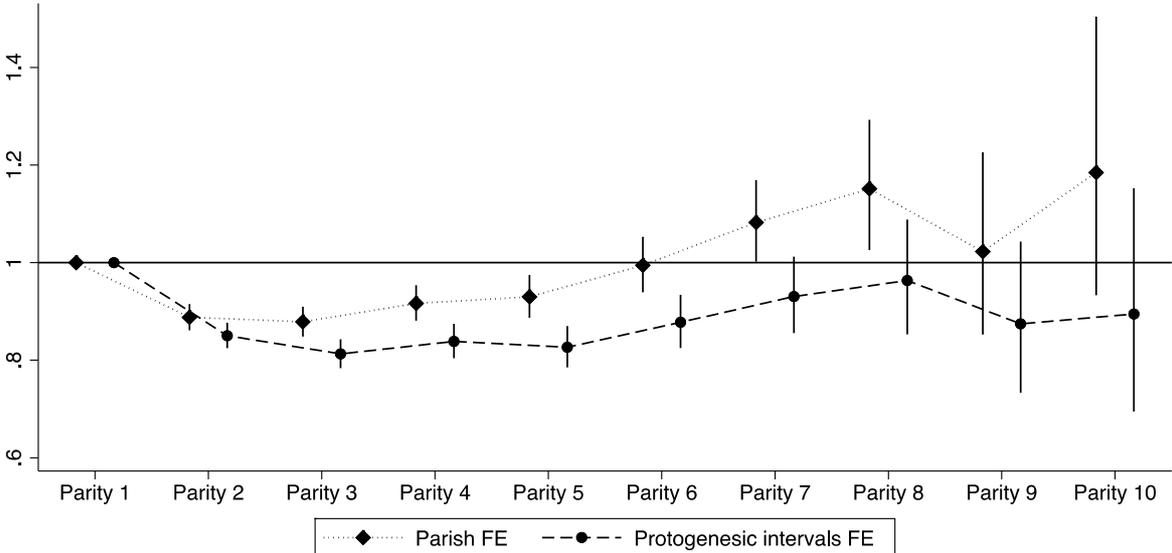
Note: The marginal effects of mother's age (solid lines) and 95% confidence bands (dotted lines) at childbirth are computed for net parity 3, with no mortality of the previous child, no last birth interval, father's occupation craftsmen, and not born on January 1st, January 11th, nor on December 25th. Model 1 (black): with continuous variable for net parity. Model 2 (orange): with net parity dummies. Model 3 (blue): with net parity dummies, without control for last birth interval. Model 4 (green): both without dummies for net parity and control for last birth interval.

Figure 4: Mother's age at childbirth effects from non-parametric specification.



Note: Coefficients for the dummy variables of mother's age at childbirth with 95% confidence intervals from columns 3 and 6, Table 2. Point estimates of the two models are not aligned along the x-axis for visual purposes.

Figure 5: Net parity effects with and without stratification by protogenetic interval.



Note: Coefficients for net parity dummies with 95% confidence interval from estimates in columns 3 and 6, Table 2. Point estimates of the two models are not aligned along the x-axis for visual purposes.

c. Models with shared frailty

Shared frailty models present yet another way to deal with Clark and Cummins' concern. Duration models with shared frailty are effectively random-effects models in which, in their simplest form, an unobserved random factor modifies the hazard function of an individual or a group of individuals in a multiplicative way. Duration models with shared frailty are therefore useful to account for heterogeneity in couples' fecundity *without* stratification on the level of the mother. Models with shared frailty have been previously used in the historical demographic literature (e.g., Dribe and Scalone, 2010; Bengtsson and Dribe, 2006). A complication with the shared frailty model is its high computational cost, as the model estimates random effects for each of the sampled households. Hence, we estimated shared frailty models on a 5% and a 10% randomly selected sub-sample of the dataset.

The results of these specifications for both samples are reported in Table 3. In columns 1 and 3 we include a continuous variable for net parity, whereas in columns 2 and 4 we include dummy variables for net parity. The relative coefficients for net parity from columns 2 and 4 are displayed in Figure 6. Consistent with the previous results, we find that higher net parities have a monotonic negative effect on the hazards of subsequent births, further supporting the hypothesis of deliberate birth control. Also, as expected, maternal age has a positive association with the length of birth intervals. Importantly, we also find a significant frailty effect, suggesting the existence of potentially important unobserved heterogeneity across households.

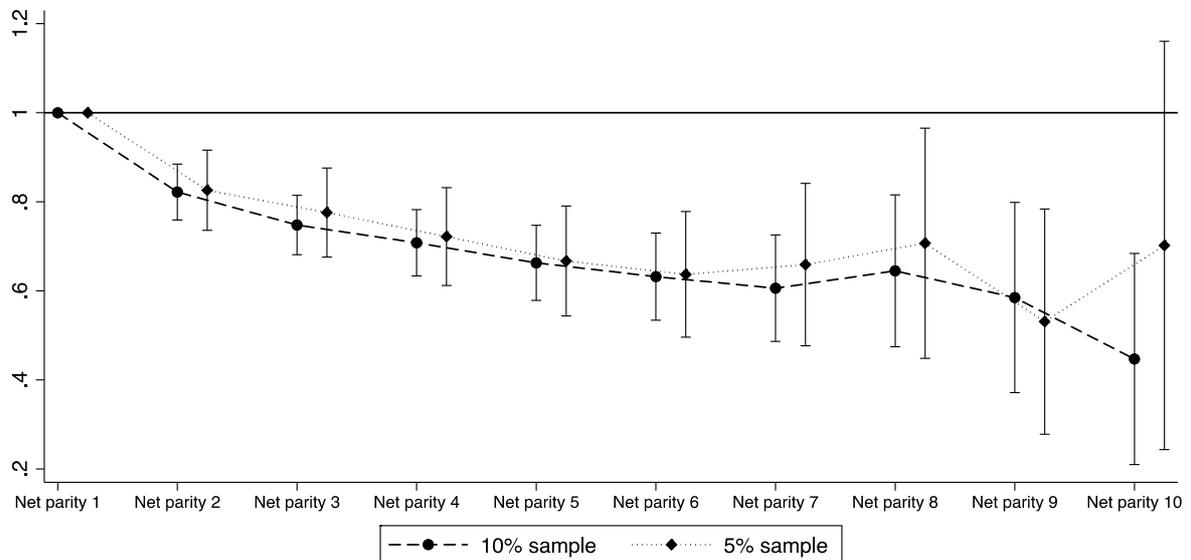
Table 3: Parity-dependent birth spacing accounting for shared frailty

	Models with shared frailty			
	5% sample		10% sample	
	(1)	(2)	(3)	(4)
Net parity	0.931*** (0.015)		0.921*** (0.011)	
Net parity 2		0.826*** (0.046)		0.822*** (0.032)
Net parity 3		0.776*** (0.051)		0.748*** (0.034)
Net parity 4		0.722*** (0.056)		0.708*** (0.038)
Net parity 5		0.667*** (0.063)		0.663*** (0.043)
Net parity 6		0.637*** (0.072)		0.632*** (0.050)
Net parity 7		0.659*** (0.093)		0.606*** (0.061)
Net parity 8		0.707*		0.645***

		(0.132)		(0.087)
Net parity 9		0.531***		0.585***
		(0.129)		(0.109)
Net parity 10		0.702		0.447***
		(0.234)		(0.121)
Real wage	1.062**	1.063**	1.089***	1.088***
	(0.031)	(0.031)	(0.022)	(0.022)
Mother's age at birth	0.886***	0.909***	0.860***	0.883***
	(0.028)	(0.030)	(0.019)	(0.020)
Mother's age at birth (squared)	1.002***	1.002***	1.002***	1.002***
	(0.000)	(0.001)	(0.000)	(0.000)
Infant death	1.873***	1.875***	1.925***	1.932***
	(0.108)	(0.108)	(0.077)	(0.078)
Last birth interval	0.602***	0.604***	0.569***	0.571***
	(0.030)	(0.030)	(0.020)	(0.020)
Frailty	0.287***	0.290***	0.289	0.292***
	(0.035)	(0.035)	(0.024)	(0.025)
Subjects	3405	3405	6897	6897

Note: Cox proportional hazard models with shared frailty. Frailty is assumed to be Gamma distributed. Hazard ratios are reported. Real wages and mother's age vary within the family. Real-wages are standardized with zero mean and unit standard deviation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 6: Net-parity effects of models with shared frailty (5% and 10% random sample)



Note: Coefficients for net parity with 95% confidence intervals from estimates in column 2 and 4, Table 3. Point estimates of the two models are not aligned along the x-axis for visual purposes.

4. The effect of the real wage on birth spacing

In a crude and highly speculative fashion, Clark and Cummins also attempt to raise doubts about the parity-independent birth-control analysis reported in Cinnirella, Klemp, and Weisdorf (2017). There, we present evidence that couples responded to changes in real wages by regulating the length of their birth intervals. In particular, we

show that couples tended to postpone the birth of the successive child during periods of relatively low real wages, and vice versa. Clark and Cummins claim that what we show are just “connections”; that the relationship between occupational groups and the spacing of their births could “well just be a product of the location and lifestyle of that group”; and that the relationship between spacing and real wages “could be just a product of nutritional stress” (Clark and Cummins, 2019, p. 12). Given the arbitrary nature of these claims, along with a complete lack of any empirical demonstration which contradicts our findings, we suspect that their claims are based on a misreading of Cinnirella, Klemp, and Weisdorf (2017). Below, we therefore briefly recapitulate how we accounted for potential confounding factors in our empirical analysis.

When studying the relationship between occupational groups and birth spacing, we stratify by parish and quarter-century. Stratification by parish accounts for time-invariant characteristics that vary across parishes, thus controlling for “location” or “environment” effects mentioned by Clark and Cummins. Stratification by quarter-century accounts for potential “lifestyles” that change with time and that affect all the parishes in the same way.⁷

In addition, we further conduct an analysis on data restricted to the six different occupational groups, while stratifying on the level of the mother and quarter-century (Cinnirella, Klemp and Weisdorf, 2017, Table 8). For all six groups, real wages are negatively correlated with birth intervals. When the most affluent groups, i.e. farmers, merchants and gentry, are combined (columns 5 and 6 of Table 8 in Cinnirella, Klemp and Weisdorf, 2017), the real-wage coefficient is also statistically significant (coefficient estimate: 1.073, p -value: 0.049). These results again suggest that the negative association between real wages and birth intervals are *not* a product of differences in living environments or social patterns.

We also consider the effects among subsamples including stayers, migrants, families with relatively short birth intervals, and *completed* marriages (*ibid.*, Table 7). The real-wage coefficient is significantly larger than one in all subsamples, suggesting that heterogeneity in environments and social patterns among these subgroups did not account for the observed *preventive-check* effect within marriages.

Lastly, in an additional analysis (*ibid.*, Table S6), we accounted for the potentially confounding effects of climatic conditions (proxied by temperatures) as well as nutrition and the living environments (proxied by crude death rates). The effect of real

⁷ The joint stratification by parish and quarter century furthermore accounts for location or environment effects specific to each quarter century.

wages on birth intervals remains highly statistically significant. Overall, the analysis in Cinnirella, Klemp, and Weisdorf (2017) clearly demonstrates that social patterns, living environments, and nutrition are unlikely to be accountable for the observed parity-independent birth-control behavior.

5. Conclusion

The identification of net-parity effects on subsequent births using duration models requires that heterogeneity in couples' fecundity is accounted for. Cinnirella, Klemp, and Weisdorf (2017) used duration models in which heterogeneity is dealt with by stratifying observations on the level of the mother. Based on family-reconstitution data from historical England, our study revealed a negative and statistically-significant effect of net parity on the risk of a subsequent birth. We interpreted this finding as evidence of birth control within marriage prior to the demographic transition. Clark and Cummins (2019) took issue with this finding, claiming it is a statistical artefact caused by stratifying on the maternal level while controlling for the mother's age at the beginning of each birth interval and concluding that marital fertility in pre-transitional England was random.

In this study, we have demonstrated that Clark and Cummins' claim, as well as their conclusion, is wrong. In particular, we show that the negative effect of net parity persists even after accounting for heterogeneity in fecundity through stratification by the couples' protogenetic intervals and the use of models with shared frailty. Although Clark and Cummins could have been correct in assuming that the joint stratification on the maternal level and inclusion of control variables for maternal age might have downward biased net-parity coefficients, the existence of deliberate fertility control in the form of net parity effects remains evident. We show that estimates of the hazard of a next birth increases significantly with net parity independently of the identification strategy used. Moreover, we show that Clark and Cummins are also wrong in asserting that our findings are dependent on "impossible" maternal age estimates. Maternal age is always overall negatively related to the hazard of a next birth according to our new estimations. Furthermore, in their arbitrary criticism of our parity-independent birth-control results, Clark and Cummins ignored our extensive set of control variables, stratification strategies, and sub-sample analyses used to account for potential confounding factors.

Three important lessons spring from the additional analyses presented above: (1) any serious attempt to identify deliberate birth control cannot ignore unobserved

heterogeneity in fecundity across households; (2) the estimated magnitude of net parity effects on the risk of a next birth depends on the specific method used to account for such heterogeneity; and (3) regardless of the method used to account for unobserved heterogeneity, ample evidence of deliberate marital birth control exists for pre-demographic transition economies, including historical England.

References

Amialchuk, A. and E. Dimitrova (2012): "Detecting the evolution of deliberate fertility control before the demographic transition in Germany", *Demographic Research* 27, 507-542.

Anderton, D.L., and L.L. Bean (1985): "Birth spacing and fertility limitation: A behavioral analysis of a nineteenth century frontier population", *Demography* 22, 169-83.

Bengtsson, T. and M. Dribe (2006): "Deliberate control in a natural fertility population: Southern Sweden, 1766-1864", *Demography*, 43, 727-746.

Cinnirella, Francesco, Marc Klemp, and Jacob Weisdorf (2017), Malthus in the Bedroom: Birth Spacing as Birth Control in Pre-Transition England, *Demography*, 2017, 54(2), 413-436.

Clark, Gregory and Neil Cummins (2018), The Child Quality-Quantity Tradeoff, England, 1780-1880: A Fundamental Component of the Economic Theory of Growth is Missing, mimeo.

Clark, Gregory and Neil Cummins (2019) "Randomness in the Bedroom: There is No Evidence for Fertility Control in Pre-Industrial England," *Demography*, forthcoming.

David, P. A., and Mroz, T. A. (1989a): "Evidence of Fertility Regulation among Rural French Villagers, 1749-1789: A Sequential Econometric Model of Birth-Spacing Behaviour (Part 1)", *European Journal of Population*, 5(1), 1-26.

David, P. A., and Mroz, T. A. (1989b): "Evidence of Fertility Regulation among Rural French Villagers, 1749-1789: A Sequential Econometric Model of Birth-Spacing Behaviour (Part 2)", *European Journal of Population*, 5(2), 173-206.

Dribe, M. and F. Scalone (2010), "Detecting Deliberate Fertility Control in Pre-Transitional Populations: Evidence from six German villages, 1766-1863", *European Journal of Population/Revue européenne de Démographie*, 26, 411-434.

Ejrnæs, M. and C. C. Pörtner (2004), "Birth Order and the Intrahousehold Allocation of Time and Education", *Review of Economics and Statistics*, 86(4), 1008-1019.

Galor, O. and M. Klemp (2019), "Human Genealogy Reveals a Selective Advantage to Moderate Fecundity", *Nature Ecology & Evolution*, forthcoming.

Klemp, M., and J. Weisdorf (2019), "Fecundity, Fertility, and the Formation of Human Capital", *Economic Journal*, 129, 925-960.

Kolk, M. (2011): “Deliberate Birth Spacing in Nineteenth Century Northern Sweden”, *European Journal of Population*, 27, 337-359.

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 Fertility Transition: A Critical Review of Methods”, *Population (English edition)*, 59, 117–128.

Van Bavel, J. and J. Kok (2004): “Birth Spacing in the Netherlands: The Effects of Family Composition, Occupation and Religion on Birth Intervals, 1820-1885”, *European Journal of Population*, 20, 119–140.

Van Bavel, J. and J. Kok (2010): “A mixed effects model of birth spacing for pre-transition populations”, *The History of the Family*, 15(2), 125-138.