Elastic Analysis Procedures--An Incurable (but Preventable) Problem in the Fertility Effect

Literature: Comment on Gildersleeve, Haselton, & Fales (2013)

Christine R. Harris and Harold Pashler

University of California, San Diego

Laura Mickes

Royal Holloway, University of London

Author Note

Christine R. Harris, Department of Psychology, University of California, San Diego;

Harold Pashler, Department of Psychology, University of California, San Diego; Laura Mickes,

Department of Psychology, Royal Holloway, University of London.

Correspondence concerning this article should be addressed to Christine R. Harris,

Department of Psychology, University of California, San Diego, 9500 Gilman Dr. #0109, La

Jolla, CA 92093-0109.

Email: charris@ucsd.edu

Abstract

Gildersleeve, Haselton, and Fales (2014) present a meta-analysis of the effects of fertility on mate preferences in women. Research in this area has categorized fertility using a great variety of methods, chiefly based on self-reported cycle length and time since last menses. We argue that this literature is particularly prone to hidden experimenter degrees of freedom. Studies vary greatly in the duration and timing of windows used to define fertile vs. nonfertile phases, criteria for excluding subjects, the choice of what moderator variables to include, as well as other variables. These issues raise the concern that many or perhaps all results may have been created by exploitation of unacknowledged degrees of freedom ("p-hacking"). Gildersleeve et al. sought to dismiss such concerns, but we contend that their arguments rest upon statistical and logical errors. The possibility that positive results in this literature may have been created, or at least greatly amplified, by p-hacking receives additional support from the fact that recent attempts at exact replication of fertility results have mostly failed. Our concerns are also supported by findings of another recent review of the literature (Wood, Kressel, Joshi, & Louie, in press.) We conclude on a positive note, arguing that if fertility-effect researchers take advantage of the rapidly emerging opportunities for study pre-registration, the validity of this literature can be rapidly clarified.

Keywords: menstrual cycle, ovulatory cycle, p-hacking, fertility effects, mate preferences, researcher degrees of freedom

Meta-analysis can often be very useful in allowing the scientific community to rationally aggregate the information contained in a complex literature, especially one that contains conflicting findings. However, it has been generally acknowledged that the credibility of a meta-analysis can be undermined in some situations. One widely recognized threat arises when unpublished studies are omitted. Another major threat, rarely discussed until the past few years, is at least as injurious to the interpretability of literature syntheses and also much harder to mitigate. This threat involves the bias that arises when the data in the studies being surveyed were originally analyzed in a tendentious fashion by experimenters wittingly or unwittingly exploiting unacknowledged degrees of freedom in order to obtain positive results (Simmons, Nelson, & Simonsohn, 2011). Such biased analyses have come to be referred to as "p-hacking".

We recently pointed out that the evolutionary psychology literature on fertility effects may be unusually prone to p-hacking partially because of the high degree of variability that exists from study to study in the method used to classify different women as fertile or infertile (Harris, Chabot, & Mickes, 2013; Harris, 2013). The idea that positive results in this literature may have been created, or at least greatly amplified, by p-hacking receives support from two observations: (a) recent attempts at exact replication of fertility results have mostly failed, and (b) an examination of just the unpublished studies in the literature finds no effects[1] (Wood, Kressel, Joshi, & Louie, in press).

In their review, Gildersleeve, Haselton, & Fales (2014) mention the problem of analytical elasticity, but they attempted to dismiss it. The purpose of this commentary is to show why this dismissal is unfortunately unpersuasive. We argue that Gildersleeve have made logical errors in

---

[1] One potential way that p-hacking might explain such null findings would be if investigators who failed to p-hack usually found no results, and therefore, did not publish their results, whereas those who used p-hacking tended to succeed in getting positive findings, which they then published.

their analysis of how p-hacking would be expected to manifest itself in the literature. Our comment concludes on a positive note, however, pointing out that even if the current literature is inconclusive because of analytical flexibility and the potential for p-hacking, a solution is rapidly emerging that will allow future studies to definitively avoid these problems and test existing claims.

The arguments offered by Gildersleeve et al. (2014) to dismiss the role of p-hacking involve several misleading sets of claims, which we discuss and rebut below. First, they downplay the actual amount of variation in fertility definitions found in the literature. We point out that the flexibility is quite dramatic—sometimes even enough to reverse the direction of effects. Second, they argue that this potential flexibility is not actually leading to biased effects because investigators decide in advance how they will analyze their data. We acknowledge that the frequency of "fishing expeditions" in this literature is obviously not directly observable, but we argue that the fact that fertile periods often vary *from one paper to another within the same lab* strongly suggests that p-hacking is more than a theoretical possibility (see Figure 1, Harris et al., 2013). Third, they maintain that the use of continuous fertility methods (as in a relatively small subset of the studies that they review) gets around the problem of flexibility; we explain why it does not solve the problem. Finally, Gildersleeve et al. contend that specific quantitative findings from their meta-analysis argue against the idea that results are being inadvertently manufactured through p-hacking. We show that these arguments rest on logical and statistical confusions. We turn now to the first issue.

How much flexibility is there in the literature?

In attempting to argue that p-hacking is not a problem in this literature, Gildersleeve and colleagues (2014) play down the amount of flexibility in analyses that a typical menstrual cycle

study has. They write, "Most aspects of study design are determined in advance of data collection, thereby eliminating concerns about researcher degrees of freedom" (p. 52) and while they note that definition of high and low fertility windows is not always determined in advance they repeatedly claim that they are "usually defined in advance of data collection" (p. 147). In a contemporaneous piece, Gildersleeve et al. (2013) boldly state, "It is implausible that these findings are a mere artifact of 'researcher degrees of freedom.'" (p. 520).

In this literature, there are an unusual number of analytical choices that investigators must make. Ironically, Gildersleeve and colleagues' (2014) own figures illustrate some of these choices (see Figure 3).[2] For example, consider the assignment of women into the high and low fertile groups. The experimenter chooses 1) the number of days to be counted as fertile, 2) the specific days considered fertile, 3) the number of days to be counted as infertile, 4) the specific days considered infertile, 5) which days will be thrown out of analyses all together (for example, exclude all women who have cycle days longer than 28 days), and 6) whether to assess fertility with a forward method, backward method, or some mixture of the two. Together, these decisions can dramatically alter fertility classifications.

For example, based on the numbers in Gildersleeve and colleagues' Figure 3, the number of days considered as fertile ranges from 3 to 15 days and the number of days considered

---

[2]Perusal of this figure may actually leave the reader with an *overestimation* of the consistency in the literature given that a single study often has multiple separate entries. It also should be noted that their figure has errors, of which we note just two: Little, Jones, Burt & Perrett (2007) are listed as having excluded women who were on day 15 or higher in their cycle when in fact these women were included in the nonfertile group. Rupp et al. (2009) are listed as having used a continuous fertility method when they actually placed participants in low fertility (days 1-5 and 17-35) and high fertility (days 6-16) groups.

infertile ranges from 3 to 22 days.[3] Thus, if experimenters simply chose from the range of

options that already exist in the field, they would have at least 13 options for the number of days

to include as high fertility and 20 options for low fertility days. The total number of options

available explodes when one considers all of the choices created by different combinations of

these variables, along with the many options for the specific placement of these windows.

The flexibility in analyses, however, does not end with determining the fertile and

infertile windows. Researchers often commonly make additional decisions about which subjects

to exclude (e.g., women over a particular age, single women) and which moderators to examine

(e.g., relationship status; primary mate's characteristics). Another avenue for flexibility comes

from using different transformations of dependent variables.

If changes in fertility definitions are large, one might still wonder whether they are large

enough to alter results. Although this question could benefit from detailed study applying

simulated classifications to real datasets, there is little doubt that changing fertility windows can

transform the results of a study. For example, Harris (2011) attempted to replicate Penton-Voak

and colleagues' findings that women in the fertile phase prefer more masculine faced men than

women in nonfertile phases (Penton-Voak et al., 1999; Penton-Voak & Perrett, 2000), using

stimuli obtained from Penton-Voak. Harris' primary analyses were performed exactly as Penton-

Voak and Perrett (2000): Women who had cycle days more than 28 days were excluded, women

who were on cycle days 6-14 were placed in the high fertility group and all others were placed in

the low fertility group. Harris found a significant effect (p < .03), but in the *opposite* direction to

that reported by Penton-Voak and colleagues: women in the fertile phase preferred *less*

---

[3] This number is actually greater since some studies include cycle lengths as long as 40 days, but the Gildersleeve et al. figure only displays data for 28 days. Extreme variability also has been documented in other reviews (Harris et al., 2013; Wood et al., in press).

masculine faces relative to women in the nonfertile phase. In subsequent analyses, Harris showed

how this effect could be changed by relatively small shifts in the fertile window. When days 8-16

were considered the fertile days (a shift of only two days from the previous analysis) and all

remaining days were counted as not fertile, the effect disappeared (p >.56). Interestingly,

Gildersleeve and colleagues (2014) chose to include the latter analysis, rather than the former, in

their meta-analysis[4]. In other work, Wood and colleagues reanalyzed data from Frost (1994) and

showed that reducing the fertile phase window by one day could make a previously significant

result no longer significant (p = .283). Thus, even slight variations in analytic strategies

sometimes have drastic effects on whether an analysis produces significant effects.

Is potential analytical flexibility being exploited to seek positive results?

We contend that when analytical flexibility is present, it is only sensible to assume that

some experimenters (whether consciously or unconsciously) will be exploiting this flexibility in

order to find positive publishable effects, and more specifically to arrive at findings they find

theoretically agreeable. In a series of simulations, Simmons et al. (2011) examined the

consequences of a number of research practices such as excluding subsets of participants,

exploring different transformations of dependent measures, not reporting all analyses or

---

[4] Gildersleeve et al. justify their selection by claiming that they chose to include the Harris
analysis that was based on the high-fertility window with the highest estimated average
conception probability according to the values reported by Wilcox and colleagues (2001).
Previously, Gildersleeve et al. (2013) have claimed "the backward counting method is generally
regarded as a more accurate method of estimating cycle position and fertility" (pp. 519).
Therefore, it is odd that they did not choose to use the analyses in which Harris used a backward
fertility estimate, especially since this is the analysis for which the high-fertility window had the
highest estimated average conception probability (contrary to Gildersleeve and colleagues'
claims in footnote 9).

conditions, etc. They showed that such tactics, especially when adopted in combination, can

increase false alarm rates far above the nominal 5% that is assumed to exist in published

research. For example, Simmons et al. found that performing four such practices in combination

resulted in a 60% likelihood of finding an effect that was significant at p = .05 in the absence of

any real effect.

The clearest evidence for exploitation of analytical flexibility arises when the same

researchers adopt and then discard various transformations of the same dependent variables from

one study to the next, without any justification being provided for these shifts. For example, both

Pillsworth and Haselton (2006) and Haselton and Gangestad (2006) examined a mate's sexual

and investment attractiveness. In Pillsworth and Haselton (2006) the types of attractiveness were

analyzed separately, but in Haselton and Gangestad (2006) a difference score of sexual vs.

investment attractiveness was calculated and then used in analyses.

Of course, it should go without saying that many investigators undoubtedly do not engage

in p-hacking, even if common practices would have allowed them to. To produce the slight

preponderance of positive findings that Gildersleeve et al. (2014) contend the literature shows, it

would probably only take a fairly modest number of biased analyses. Moreover, such behaviors

may be undertaken in good faith, related to ignorance of their consequences. The outpouring of

interest that the Simmons et al. (2011) paper has drawn within the scientific community probably

shows that many investigators have been unaware of how small choices can cumulate to easily

produce statistically significant findings built out of sampling error.

Does using continuous calculations for fertility solve the p-hacking issue?

Gildersleeve et al. (2014) argue that restricting their analysis to just those studies that

employ continuous fertility calculations (rather than dividing women into high and low fertility

groups) can resolve concerns about p-hacking, because such studies do not require the investigator to specify windows. They suggest that some effects are still present in this subgroup of studies, and therefore experimenter degrees of freedom are not an issue.

However, while continuous fertility calculation methods generally reduce the number of choices to be made relative to the categorical classification methods, they are not necessarily as straightforward as Gildersleeve et al. (2014) suggest, and many such studies still report making complex and idiosyncratic analysis decisions. Most investigations cite Wilcox et al. (2001) as the source of their continuous fertility numbers. Wilcox et al. provided a table for estimated probability of pregnancy following a single act of unprotected intercourse, which provides a risk estimate for each day in a woman's cycle up to day 40 (separately for women whose cycle lengths are consistent vs. irregular). The most straightforward way to use this table is to assign a woman a probability value based on the day in her cycle (as done by Morrison, Clark, Gralewski, Campbell, & Penton-Voak, 2010). However, other researchers take into account a woman's predicted cycle length and perform transformations to force the woman on a 28/29-day cycle (regardless of what her natural cycle length is) and then try to apply Wilcox et al. table to these transformed numbers (e.g., Thornhill, Chapman, & Gangestad, 2013)[5]. Other studies combine the 28/29-day conversions with additional transformations such as averaging forward and backward fertility calculations (e.g., Gangestad, Garver-Apgar, Simpson, & Cousins, 2007). Although such analytic strategies may be individually defensible, the fact that such a range of choices exists undercuts Gildersleeve and colleagues' argument.

---

[5] This procedure could alter the data quite substantially. Take, for example, a woman who has a 30 day-cycle and is tested on day 14. A straightforward application of Wilcox et al. would give her a pregnancy risk estimate of .085. However, transformations using her cycle length and assumptions that she is fertile 15 days prior to the start of her next cycle would produce a risk estimate of .059.

Quantitative evidence about p-hacking

Gildersleeve et al. (2014) also describe some specific analyses they undertook to shed light on the possible role of p-hacking (pp. 52-55). One such analysis takes all the studies that placed women in discrete low or high fertility categories, and attempts (using a continuous estimator of fertility) to estimate the fertility difference expected from this choice of window. That amounts to a grade of how successfully the study sorted high from low-fertile women. Having graded studies in this way, they look at whether there is a detectable correlation across studies between this grade and the effect size that the study reports for the outcome variable. They report that there is no sign that studies that made a "good choice" of window produce any bigger effects.

The results of this analysis are problematic for the view that this literature reveals real fertility effects. If it does, then studies that better measure fertility should tend to show stronger effects. On the other hand, if fertility effects in the literature were all produced by Type-1 errors based upon sampling error, then the windows that are graded high and those graded low by Gildersleeve and colleagues' procedure should not differ in the size of reported effects. This is what their analysis showed. (The failure of better "quality" studies to show effects also has been reported by Wood et al., in press; e.g., studies that used hormonal assessments did not find more robust effects.)

Having failed to find confirmation for real effects from their analysis, Gildersleeve and colleagues (2014) then display a figure with effect sizes for a subset of studies along with information specifying which days each study counted as fertile or infertile and which were excluded altogether. According to Gildersleeve et al., the reader should be able to see by eye that there are no interesting correlations between larger effect sizes and any of three things,

"a) more variable high- and low-fertility window definitions, b) more poorly placed high- and

low-fertility windows (high-fertility windows that included true low-fertility days of the cycle

and/or low-fertility windows that included true high-fertility days of the cycle), and c) less

frequent use of a continuous fertility variable…" (p. 53). Thus, Gildersleeve et al. conclude:

"we used multiple procedures to assess and adjust for various forms of potential bias. The results

of these procedures do not suggest that these sources of bias account for the robust cycle shifts

observed in this meta-analysis." (p. 53).

        We find it odd that Gildersleeve et al. would suggest that readers can rely on "eyeball"

judgments of complex multidimensional data—especially when the data are presented in tabular

fashion. Clearly, this provides no reliable information about whether any relationships exist.

But even if they *had* demonstrated the lack of any such relationships, the implications for

possible p-hacking would not be self-evident. The patterns produced by p-hacking would

depend upon how the p-hacking was carried out. For example, if there is no true effect of

fertility on the outcome variables being measured, and investigators tried multiple different

candidate windows selected at random until a significant result was obtained (if at all), one

should not expect to see any relation between choice of a fertile window and the strength of the

positive effects. In a universe composed exclusively of type-1 errors, any measurement of the

independent variable is as good (or bad) as any other.

        The most reasonable way to see whether p-hacking could result in a correlation between

window duration and effect size would be if the following conditions happen to hold: (a)

researchers who p-hack begin their analyses with small windows, and then try progressively

larger windows until they get an effect, and (b) researchers vary in how much they persevere in

p-hacking. Interestingly, another recent meta-analysis of this literature (Wood et al., in press)

found that larger fertility windows were *more* likely to show effects. This effect is exactly what one would predict if investigators tend to start with relatively narrow fertility windows, and then some investigators (i.e., those engaging in more energetic "p-hacking") expand the window in search of effects. We are unable to think of any completely benign explanation for this pattern, and it would seem to trump Gildersleeve and colleagues' failure to observe any relationships from looking at a figure.

Looking Forward to Better Research: How to Prevent such Problems in the Future

The general point of the current comment is to say that excess analytical flexibility makes the literature less conclusive than it could and should be. Supporting this interpretation is that those few studies in which we can be sure that analytical flexibility was not present--namely, recent papers that attempted to perform *direct replications*—have generally reported negative results. For example, as described above, Harris (2011) carried out a fairly direct replication of Penton-Voak and colleagues, using the same choices as the original investigators, and found no evidence whatsoever for the predicted cycle shift in facial masculinity preferences (with the effect running in the opposite direction). Two additional investigations that also used these same methods, fertility classifications, etc. failed to find any effect of cycle phase on masculinity preferences, even when the relationship context was specified as short-term (Mickes & Harris, 2014). Another direct replication of shifts in religiosity, political attitudes, and voting preferences (Harris & Mickes, in press) failed to support most of the effects of fertility reported by Durante, Rae, & Griskevicius (2013).

Looking beyond these direct replications, what methodological improvements can help avoid the kinds of problems discussed here? Fortunately, the problem of flexibility and potential for "p-hacking" which bedevils the fertility effect literature can be prevented in a simple and

decisive fashion. What is needed is for investigators to conduct future studies using pre-registration (Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012) of their definition of fertile periods and all other analytical methods, including plans for excluding subjects and potential moderator variables to be included in analyses. Pre-registration is now possible and convenient through the Open Science Framework (OSF.io). (The use of biological measures of fertility may also increase reliability, but the use of these tests does not eliminate many of the forms of elasticity described above, so when used, should be combined with pre-registration.) Pre-registration became the norm some years ago in clinical trials, and there is rapidly increasing awareness of its potential to advance basic research as multiple journals and organizations embrace it. New studies should routinely utilize pre-registration, and key studies in the literature need to be replicated with pre-specified procedures as well. If this plan is adopted, we will soon have a good idea of whether the findings in this literature are solid but relatively small (as Gildersleeve et al., 2014, contend) or whether many effects have been invented out of the whole cloth (as the Wood et al., in press, meta-analysis would seem to suggest).

References

Durante, K.M., Rae, A., & Griskevicius, V. (2013). The fluctuating female vote: Politics, religion, and the ovulatory cycle. *Psychological Science.*doi:10.1177/0956797612466416

Gangestad, S. W., Garver-Apgar, C. E., Simpson, J. A., & Cousins, A. J. (2007). Changes in women's mate preferences across the ovulatory cycle. *Journal of Personality and Social Psychology, 92,* 151-163.

Gildersleeve, K., DeBruine, L., Haselton, M.G., Frederic, D.A., Penton-Voak, I.S., Jones, B.C., & Perrett, D.I. (*2013*). Shifts in women's mate preferences across the cycle: A critique of Harris (2011) and Harris (2012). *Sex Roles*, *69,* 516-524.

Harris, C. R. (2011). Menstrual cycle and facial preferences reconsidered. *Sex Roles*, *64,* 669-681. doi:10.1007/s11199-010-9772-8

Harris, C.R. (2013). Shifts in masculinity preferences across the menstrual cycle: Still not there. *Sex Roles, 69,* 507-515.

Harris, C.R. & Mickes, L. (in press). Women can keep the vote: No evidence that hormonal changes during the menstrual cycle impact political and religious beliefs. *Psychological Science.*

Harris, C.R., Chabot, A. & Mickes. L. (2013). Shifts in methodology and theory in menstrual cycle research on attraction. *Sex Roles, 69,* 525-535.

Haselton, M. G. & Gangestad, S. W. (2006). Conditional expression of women's desires and men's mate guarding across the ovulatory cycle. *Hormones and Behavior, 49,* 509-518. doi:10.1016/j.yhbeh.2005.10.006

Little, A. C., Jones, B. C., Burt, D. M., & Perrett, D. I. (2007). Preferences for symmetry in faces

    change across the menstrual cycle. *Biological Psychology, 76,* 209-216.

    doi:10.1016/j.biopsycho.2007.08.003

Mickes, L. & Harris, C.R. (2013). No Ovulatory Cycle Effects on Masculinity Preferences.

    Unpublished data.

Morrison, E. R., Clark, A. P., Gralewski, L., Campbell, N., & Penton-Voak, I. S. (2010).

    Women's probability of conception is associated with their preference for flirtatious but

    not masculine facial movement. *Archives of Sexual Behavior, 39,* 1297-1304.

Penton-Voak, I. S., & Perrett, D. I. (2000). Female preference for male faces changes cyclically:

    Further evidence. *Evolution and Human Behavior, 21,* 39- 48. doi:10.1016/S1090-

    5138(99)00033-1

Penton-Voak, I. S., Perrett, D. I., Castles, D. L., Kobayashi, T., Burt, D. M., Murray, L. K., &

    Minamisawa, R. (1999). Menstrual cycle alters face preference. *Nature, 399,* 741-742.

    doi:10.1038/21557

Pillsworth, E. G., & Haselton, M. G. (2006). Male sexual attractiveness predicts differential

    ovulatory shifts in female extra-pair attraction and male mate retention. *Evolution and

    Human Behavior, 27,* 247-258. doi:10.1016/j.evolhumbehav.2005.10.002

Rupp, H. A., Librach, G. R., Feipel, N. C., Ketterson, E. D., Sengelaub, D. R., & Heiman, J. R.

    (2009). Partner status influences women's interest in the opposite sex. *Human Nature, 20,*

    93-104.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: undisclosed

    flexibility in data collection and analysis allows presenting anything as significant.

    Psychological Science, 22, 1359–1366. doi:10.1177/0956797611417632

Thornhill, R., Chapman, J. F., & Gangestad, S. W. (2013). Women's preferences for men's

    scents associated with testosterone and cortisol levels: Patterns across the ovulatory

    cycle. *Evolution and Human Behavior, 34,* 216-221.

Wagenmakers, E. J., Wetzels, R., Borsboom, D., van der Maas, H. L., & Kievit, R. A. (2012).

    An agenda for purely confirmatory research. *Perspectives on Psychological Science*,

    7(6), 632-638.

Wood, W., Kressel, L., Joshi, P. D., & Louie, B. (in press). Meta-analysis of menstrual cycle

    effects on women's mate preferences. *Emotion Review.*