Response to Flournoy (2021): Reflections on the benefits and challenges of communicating the results of a pre-registered study

Monica E. Ellwood-Lowe¹, Ruthe Foushee^{1,2}, Mahesh Srinivasan¹

¹ University of California, Berkeley

² University of Chicago

In 2020, we posted a preprint online presenting the results of two pre-registered studies, now published in revised form (Ellwood-Lowe et al., 2021; original preprint archived at <https://osf.io/ktsdp/>). While undergoing the journal review process, circulation of this preprint encouraged open feedback from peers, including a thoughtful comment on our studies and pre-registration best practices by Dr. John Flournoy (Flournoy, 2021). Here we respond to the points raised in Flournoy's comments, and discuss the benefits and challenges of pre-registering "high-risk" studies, particularly as early career researchers. We begin by reflecting on our scientific process between the first pre-registration and the ultimate dissemination of results—pulling back a curtain to reveal a piece of the research process that is often occluded. Next, we address Flournoy's comments directly and explain how they were helpful in shaping the final version of our paper. Finally, we make suggestions for pre-registered studies in the future.

Post-pre-registration, pre-paper process

Science is slow; science with young children especially so. We pre-registered Study 1 of our paper in the fall of 2017, the first semester of the first year of the first author's doctoral program. A relatively new endeavor for all involved, it was expectedly imperfect. Despite having already spent three months piloting the study and perfecting the design,

there were unanticipated hurdles when we began data collection in earnest. We were interested in whether caregivers would speak less with their children after primed to reflect on recent experiences of scarcity, compared to a control condition. This study hinged on caregivers engaging earnestly with a manipulation survey, and then staying put with their three-year-old children in a windowless testing room for ten minutes (under the guise of waiting for a researcher to correct an error), while a video camera subtly filmed them. To put it simply, a caregiver, a three-year-old, and an old lab camera all find ways of not doing exactly what you expect them to. Caregivers found an opportune moment to use the bathroom; children a great time to run the halls; cameras an event to disconnect from power or have an empty battery or full SD card.

Despite difficulties with unexpectedly high levels of exclusions, early results looked promising. An analysis of the first 25 data points, performed for a conference submission, confirmed our pre-registered hypothesis that caregivers would talk less after reflecting on experiences of resource scarcity (word types: $t(19.21)=-3.45, p=.003$; word tokens: $t(16.45)=-3.43, p=.003$). By the end of data collection, however, these highly significant findings had become marginally insignificant. Also, although most parents in piloting had reflected on financial scarcity, many parents in our actual study chose to reflect on other kinds

of scarcity that were arguably less serious and chronic (e.g., not having enough milk); a post-hoc analysis showed that only the parents who had reflected on financial scarcity showed the predicted pattern. After two years of data collection, we agreed the study was too operationally costly to repeat. We were unsure of whether our study had effectively primed scarcity (we had seen failed replications of scarcity priming in the literature), as opposed to identifying parents in our sample who were already grappling with scarcity in their lives. How else could we test our hypothesis?

An idea came to us after we began to meet regularly with economists who shared our interests. Perhaps we could follow their lead and use a natural experiment approach to test whether regular variation in families' financial resources over the course of a pay cycle affected their speech to their children. We found that we could make use of already-collected all-day LENA audio recording data that included multiple recordings from families over time (<https://homebank.talkbank.org/>), and so we pre-registered a second study. This one was more straightforward, it seemed, as the data were already collected for us. The results were less straightforward—some significant, others not, though all in the direction we expected.

By the end of the entire data collection and analysis process for both studies, we had obtained a collection of results that provided *some* support for our general hypotheses. However, some of our primary pre-registered analyses were not significant, and some of our most important results derived from secondary pre-registered analyses or post-hoc analyses. Rather than continue to run new studies or replications, we decided to disseminate our findings, with the hope that our ideas and studies would spur future research.

The paper

We wrote our paper in the hopes that it would be readable for a popular audience, and

we targeted interdisciplinary journals. While we reported all of our analyses, much of our narrative focused on the results that were traditionally statistically significant, including those that were exploratory. After releasing this version of the preprint, Flournoy's careful read and comments gave us the opportunity to reflect on writing best practices after pre-registration. Accordingly, we implemented some changes to the exposition of our results as the paper proceeded through the review process. Below, we respond directly to Flournoy's comments and explain how our final paper changed as a result.

As noted before, in Study 1, we tested the hypothesis that caregivers who were asked to reflect on their recent experiences of scarcity would talk less to their children than caregivers in a control group. Flournoy's primary critique of this study was that our discussion focused on exploratory results, which would not survive correction for multiple comparison. Indeed, while our preprint noted which results were exploratory, it was not structured in a way that would be particularly clear to a reader which analyses to give most weight to. As a result, we rearranged our results and discussion based upon the structure of our pre-registration, as Flournoy suggested, making clear which analyses were primary pre-registered tests of our hypothesis, secondary pre-registered tests of our hypothesis, and exploratory analyses. In our discussion, we are more even about distinguishing significant from non-significant results, though we do believe we may have been underpowered to detect an effect, rendering our non-significant results not inherently meaningful.

In Study 2, we tested the hypothesis that caregivers would talk less to their children at the end of the month, when they were likely to be experiencing more financial scarcity, than during the rest of the month. Flournoy had two critiques: first, the under-specification of which analyses of those we'd planned would bear directly on our hypotheses; second, the

interpretation of results in relation to the preregistration.

Flournoy rightly pointed out that for our primary prediction, we pre-registered four separate primary analyses—one in each of three corpora, and one across all three corpora—each of which could have yielded a significant result. We found significant effects in one of these four tests, at $p = .006$, with non-significant effects in the same direction for the rest. Correcting for multiple comparisons, $p = .05/4 = .0125$, this result would still be considered significant. Eight more pre-registered secondary analyses explored the same question, but with additional controls to help isolate the key effect, and again found a significant effect in the same corpus, at $p = .003$, and an effect across all corpora, at $p = .006$. Correcting for eight ($p = .006$) tests, the effects in the one corpus and across all data points would remain significant; and correcting for twelve ($p = .004$) tests, the effect in the one corpus would remain significant. We note that an additional primary analysis in which income was tested as an interactive effect answered a different question—whether monthly fluctuations in child-directed speech varied as a function of family income—and was performed only once, combining the two corpora that had sufficient income data.

In response to Flournoy's critique, the final version of our paper includes a re-written discussion that is in line with this conclusion, beginning: "in one corpus, families engaged in significantly fewer conversational turns in the last week of the month, when they were likely to be experiencing the most financial scarcity..." We also added qualifications regarding the interpretation of secondary and exploratory results. Finally, as with our first study, we rearranged our results section to more clearly delineate primary from secondary and exploratory tests. In the future, we agree that a pre-registration should more clearly delineate which test will primarily test which hypothesis.

Moving forward

As Flournoy's critique makes clear, hinging findings on p -values is deceptively difficult, though the process can be made slightly more legible by clear pre-registrations and a strong understanding of the logic behind the statistic. At the same time, pre-registration and the kind of tests it spurs has drawn critique, among them that it removes the joy from the scientific process (Collins et al., 2021). Certainly, many of us early in our careers have felt bogged down by developing a scientifically perfect study, and feeling trapped when it doesn't go as predicted. (Though many of us have also felt at least as bogged down by a scientific literature that includes studies that may not be reproducible or replicable.) But we suggest that this feeling could be diminished by a widespread acceptance of straightforward scientific reporting, particularly in scientific journals.

Of the changes we made to arrive at the final draft of our paper, one that we see as particularly valuable is following Flournoy's advice to structure results in terms of primary, secondary, and exploratory analyses. Doing so allowed us to be especially clear about what exactly was predicted (rather than leaving it to the narrative)—and therefore what tests were run first, and the evidentiary value of these tests—while also elaborating explorations we pursued after the fact. Structuring papers in this way may allow for both stringent hypothesis testing and acknowledgement of the mystery and joy of scientific exploration. We believe this has allowed us to navigate the tension of staying close to the results in the discussion, while also telling the kind of story that may spur future research forward.

Overall, we have presented a series of findings that provide *some* support our hypotheses. They are not a slam dunk. Like any study testing a novel hypothesis, they raise more questions than they answer. That our manuscript was published even so reflects

highly on the editor and journal. We believe that publishing these *kind of* findings—presented as such—may be crucial for both spurring future research that resolves the uncertainty they present and to removing some of the pressure faced by early career researchers. In the future, we hope that even beyond findings like ours which provide some support for a hypothesis, journals will embrace non-significant results as long as they test an important hypothesis, are even in their interpretation, and offer directions for future research.

Collins, H. K., Whillans, A. V., & John, L. K. (2021). Joy and Rigor in Behavioral

Science. *Organizational Behavior and Human Decision Processes*, 164, 179–191.

Ellwood-Lowe, M. E., Foushee, R., & Srinivasan, M. (2021). What causes the word gap? Financial concerns may systematically suppress child-directed speech. *Developmental Science*, 1–16. <https://doi.org/https://doi.org/10.1111/desc.13151>

Flournoy, J. (2021). A comment on Ellwood-Lowe et al. (2020/2021): Presentation and discussion of results should take the lead from the preregistration. *PsyArXiv*. <https://doi.org/10.31234/osf.io/23qsw>