Wagenmakers, 2007), and the tendency to overestimate the support in favor of the alternative hypothesis (e.g., Edwards, Lindman, & Savage, 1963; Sellke, Bayarri, & Berger, 2001; Wetzels et al., 2011). Instead, our main analysis tool is the Bayes factor (e.g., Hoijtink, Klugkist, & Boelen, 2008; Jeffreys, 1961; Kass & Raftery, 1995). The Bayes factor BF_{01} quantifies the evidence that the data provide for the null hypothesis (H_0) vis-a-vis an alternative hypothesis (H_1). For instance, when $BF_{01} = 10$, the observed data are 10 times as likely to have occurred under H_0 than under H_1. When $BF_{01} = 1/5 = .20$, the observed data are 5 times as likely to have occurred under H_1 than under H_0. An additional bonus of using the Bayes factor is that it eliminates the problem of optional stopping. As noted in the classic article by Edwards et al. (1963), "the rules governing when data collection stops are irrelevant to data interpretation. It is entirely appropriate to collect data until a point has been proven or disproven, or until the data collector runs out of time, money, or patience" (p. 193; see also Kerridge, 1963).

Hence, we outlined the details of our Bayes factor calculation in the online document:

> Data analysis proceeds by a series of Bayesian tests. For the Bayesian t-tests, the null hypothesis H_0 is always specified as the absence of a difference. Alternative hypothesis 1, H_1, assumes that effect size is distributed as Cauchy (0,1); this is the default prior proposed by Rouder et al. (2009). Alternative hypothesis 2, H_2, assumes that effect size is distributed as a half-normal distribution with positive mass only and the 90th percentile at an effect size of 0.5; this is the "knowledge-based prior" proposed by Bem et al. (submitted).[10] We will compute the Bayes factor for H_0 vs. H_1 (BF_{01}) and for H_0 vs. H_2 (BF_{02})."

The details of how the two alternative hypotheses were specified are not important here, save for the fact that these hypotheses were constructed a priori, based on general principles (the default prior) or substantive considerations (the knowledge-based prior).

Next, we outlined a series of six hypotheses to test. For instance, the second analysis was specified as follows:

> "(2) Based on the data of session 1 only: Does performance for erotic pictures differ from chance (in this study 50%)? To address this question we compute a one-sample t-test and monitor BF_{01} and BF_{02} as the data come in."

And the sixth analysis was specified as follows:

> "(6) Same as (2), but now for the combined data from sessions 1 and 2."

Readers curious to know whether people can look into the future are invited to examine the results for all six hypotheses in the online appendix at http://pps.sagepub.com/supplemental.[11] In this article, we only present the results from our sixth hypothesis. Figure 2 shows the development of the Bayes factor as the data accumulate. It is clear that the evidence in favor of H_0 increases as more participants are tested and the number of sessions increases. With the default prior, the data are 16.6 times more likely under H_0 than under H_1; with the "knowledge-based prior" from Bem, Utts, and Johnson (2011), the data are 6.2 times more likely under H_0 than under H_1. Because our analysis uses the Bayes factor, we did not have to indicate in advance that we were going to test 100 participants. We calculated the Bayes factor two or three times as the experiment was running, and after 100 participants we inspected

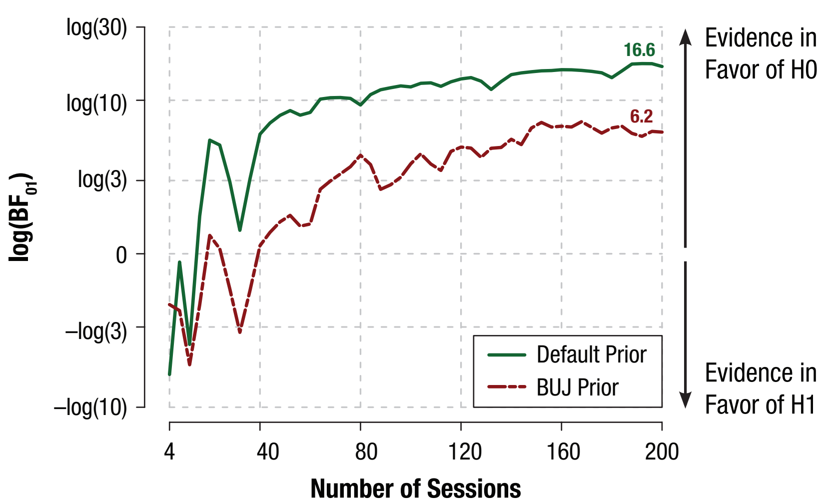

Fig. 2. Results from a purely confirmatory replication test for the presence of precognition. The intended analysis was specified online in advance of data collection. The evidence (i.e., the logarithm of the Bayes factor) supports H_0 ("performance for erotic stimuli does not differ from chance"). Note that the evidence may be monitored as the data accumulate.