Running Online Behavioral Experiments Using R:

Implementation of a Response-Time Decision Making Task as an R-Shiny App

Agustín Pérez Santangelo^{1,2} and Guillermo Solovey³

¹ Instituto de Investigación en Ciencias de la Computación, Universidad de Buenos Aires (UBA), Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET)

² Laboratorio de Neurociencia, CONICET, Universidad Torcuato Di Tella

³ Instituto de Cálculo, Facultad de Ciencias Exactas y Naturales, UBA-CONICET

Correspondence concerning this article should be addressed to Guillermo Solovey, Instituto de Cálculo, Facultad de Ciencias Exactas y Naturales, UBA-CONICET. Intendente Güiraldes 2160 - Ciudad Universitaria - C1428EGA - Buenos Aires, Argentina . Email: gsolovey@gmail.com

Abstract

Online experiments make it possible to reach a large number of participants in a short time at a relatively low cost. However, uncontrolled conditions in online experiments can be problematic, particularly when it is paramount to measure participant's response-time (RT) with millisecond-level precision. To evaluate the reliability of RT data collected online, we implemented a rapid two-digit number comparison task in a mobile-friendly open-source application using R-Shiny, a popular R package for web app development. In the task, subjects were briefly presented with a two-digit target number and had to decide if it was larger or smaller than a standard fixed number (65). A total of $N=169$ participants (109 with a mobile device, 60 on a desktop computer) completed 116 trials over a ~7-minute session. Using generalized linear mixed models estimated with Bayesian inference methods, we observed a classical numerical distance effect: RT decreases with the logarithm of the absolute difference between the target and the standard. Our results support the use of R-Shiny for RT-data collection. Furthermore, we report systematic delays in RTs induced by different OS, web-browsers, and devices, highlighting the relevance of this information to dissect the technological component from the psychological component of RTs. Our work provides the reader an example of a seamless and robust implementation of an RT decision making task running online over desktop and mobile devices. It further paves the ground for the design of simple cognitive tasks using only R, a widely popular programming framework among cognitive scientists.

Keywords: online experiments, R-Shiny, response time, numerical cognition

Introduction

Online experiments are increasingly common among behavioral and cognitive science researchers as they allow collecting large datasets from far and diverse participants at a relatively low cost in a short period of time. Moreover, during the COVID-19 pandemic, the only available choice to keep projects running is to resort to some type of online data collection. A particularly valuable source of information in experimental psychology are participant's response times (RTs). While RTs are crucial for inference of many psychological processes (e.g. Stroop interference effect), they are extremely sensitive. Oftentimes, RT effects only show up at the millisecond level. Thus, it is important to assess the reliability of RT data obtained in web-based experiments for further inference of underlying psychological processes.

To get a sense of what an online setting really means, picture yourself riding a bus on your way back home when a message pops up on your phone inviting you to participate in a decision making task. The bus swings constantly, the ambient is noisy and you are a bit tired, but you decide to give it a try anyway. After all, the task requires only 5 minutes to make a bunch of simple and quick decisions. Though it may be an extreme example, it serves to illustrate the point: uncontrolled conditions in online studies contrast sharply with the typical atmosphere during laboratory experiments in which participants are welcomed by a researcher, receive detailed instructions and complete hundreds of trials during ~60 minutes in front of a calibrated computer in a quiet room.

Online behavioral experiments are not new. Attempts to use the web as a tool for psychological research date back to the late 1990s when the popularity of the web started to grow (Krantz & Dalal, 2000; Reips, 2002). The reliability of data collected online has always been a focus of attention (Anwyl-Irvine et al., 2020; Brown et al., 2014; Crump et al., 2013;

Dandurand et al., 2008; Dufau et al., 2011; Garaizar & Reips, 2019; Germine et al., 2012; Gosling et al., 2004; Grootswagers, 2020; Huber et al., 2019; Kochari, 2019; Krantz & Dalal, 2000; Pronk et al., 2020; Semmelmann & Weigelt, 2017; Simcox & Fiez, 2014). As technology develops and new instruments become available, there is an increasing need to verify the compatibility between online and lab experiments.

Some labs developed their own native mobile apps that users could download from an app store, such as the Google Play or the iOS Apple Store, to run gamified versions of cognitive behavioral experiments (Brown et al., 2014; Dufau et al., 2011; Zimmerman et al., 2016). The results demonstrated the feasibility of using smartphones to measure RTs at the millisecond-level and the ability to replicate classical psychological findings in different domains. Later on, with the surge of services such as Amazon Mechanical Turk (AMT) where people can perform web tasks for small sums of money, Crump et al. reproduced several classical cognitive behavioral experiments using AMT participants, including four tasks that required to detect within-subject differences as short as 50ms in reaction times (Crump et al., 2013).

Recent studies have shown that online experiments are useful in numerical cognition research. An online study built in HTML/JavaScript (jsPsych) replicated classical RT effects in a task in which participants had to compare a single-digit number with a predetermined standard (Kochari, 2019). Huber et al. designed a two-digit number comparison task implemented using Qualtrics software (Qualtrics, Provo, UT) and also replicated previous findings such as the decade distance effect (Huber et al., 2019). To avoid variability on RTs introduced by different devices, Huber et al. discarded data from participants that ran the experiment on a mobile device. Two other recent numerical cognition studies also limited data collection and analysis to participants using desktop computers (Cipora et al., 2019; Gökaydin et al., 2018). Regrettably, this methodological approach severely constraints the

overall number of potential participants of an online experiment given that mobile devices are much more used than computers, especially in emerging economies (Silver et al., 2019).

Taken together, important lessons can be learned from past experiences using online tools for behavioral experiments. However, we believe there are currently new challenges to take upon. First, there is an urgent call for adopting open science practices (Munafò et al., 2017). We therefore decided to run an experiment with open-source tools and share the code online so that anyone can use it or build upon ours. Second, we advocate for an integrated workflow that goes from designing and building a behavioral experiment to modeling and data analysis. Third, to fully capitalize the power of online experiments collecting data from as many participants as possible, we need a fluid -and robust- user interface on mobile devices, and an appropriate statistical model to account for the effect of participants' devices on data. These conditions are hardly met simultaneously in most previous frameworks for online studies. Native mobile apps are good to gamify experiments but are usually task-specific and hard to adapt to new tasks. Moreover, programming these apps, as well as most web-based experiments, require skilled coding in HTML, CSS, and JavaScript, which is not what most cognitive scientists use on a daily basis for modeling and data analysis.

Shiny (Chang et al., 2018) is a powerful yet straight-forward package in R for the design and development of interactive apps that fits perfectly to our goals. As it is well integrated in RStudio (a popular IDE to write and run R code), we can code the experiment, collect data, do data analysis, modeling and -eventually- even create a Markdown report within a unified framework. Shiny is also a growing -easy to learn- framework that has already shown interesting potential for behavioral experiments (Kaufman, 2020; Steiner et al., 2020), although its power as an experimental psychology tool has yet to be fully harnessed. Last but not least, it is fairly simple to deploy a Shiny app online within RStudio

(R most popular IDE). Using a single command or button-press, a newly developed app can be online in a matter of seconds for free on shinyapps.io.

Here, our main goal is to implement a behavioral task on a Shiny app capable of measuring precise RTs at a millisecond level for participants using different devices (i.e. on mobile devices and desktop computers). A two-digit number comparison task is well suited for our purpose. In this task (depicted in Fig. 1), participants have to decide whether a number (target) is smaller or larger than a reference number (standard). Typically, RTs are shorter and error rates are lower when the absolute difference (D) between the numbers being compared is larger (i.e., comparing 61 vs. 65 is more difficult than comparing 34 vs. 65). This long-studied phenomenon, coined the numerical distance effect, has shed light on how the mind represents numbers (Dehaene, 2011; Dehaene et al., 1990; Hinrichs et al., 1981; Krajcsi & Kojouharova, 2017; R. S. Moyer & Landauer, 1973; Robert S. Moyer & Landauer, 1967; Nuerk et al., 2001; Verguts & De Moor, 2005). Several reasons support our choice to implement this task for our purpose. First, even though there are different versions of this task, mainly changing the standard number (e.g., 55, 65, 75), a robust converging result that emerge from previous studies is that $RT \sim \beta_{\log D} \log(D)$ with $\beta_{\log D} \cong -100ms$, probably revealing a universal feature of human numerical cognition (Dehaene, 2011). As such, it is sensible to expect that this effect holds true in online experiments. Second, to infer the effect of numerical distance from RT data it is necessary to measure reliable RTs at a millisecond level. Last, it is an arguably engaging and short task, which makes it even more suited for running online. We therefore hypothesize that:

H1: RT decreases linearly with the logarithm of D , consistent with previous results (Dehaene, 1989, 2011; Dehaene et al., 1990; Hinrichs et al., 1981; Krajcsi & Kojouharova, 2017; R. S. Moyer & Landauer, 1973; Robert S. Moyer & Landauer, 1967; Nuerk et al., 2001; Verguts & De Moor, 2005).

Importantly, a typical confound in online experiments are device-specific delays related, for instance, to the device's OS and web-browser (Anwyl-Irvine et al., 2020). Mobile communication delays are larger than desktop delays (Holden et al., 2019; Pronk et al., 2020; Reimers & Stewart, 2008). To control for device-specific variability, we will collect information of the device (OS and web-browser) that each participant uses to complete the task, incorporate this information into a combined variable in a statistical model, and determine its impact on RT. We hypothesize that:

H2: Participants device, OS, and web-browser will systematically affect RT delays. Specifically, we expect that mobile devices introduce a larger delay on RTs than desktop devices.

Method

Participants

A total of 169 participants (65 males, age 31.94 ± 11.69 y.o.) completed the online task either with a mobile device ($N = 109$) or on a desktop computer ($N = 60$). All participants gave informed consent, and the local ethics committee approved the study (“Centro de Educación Médica e Investigaciones Clínicas” protocol ID 435). No monetary compensation was awarded to participants.

Task

The task was adapted from Dehaene et al., 1990. In this two-digit number comparison task, subjects had to decide whether a target number was smaller or greater than 65. We presented 58 target numbers: 29 were lower than the standard (from 36 to 64), and 29 were greater than the standard (from 66 to 94) (Fig. 1). Each of these targets was displayed twice, rendering a total of 116 trials. The order of the targets was pseudo-random, as no target was shown two times in a row and correct-response location was never the same for more than two consecutive trials (i.e. we never showed more than two targets that were either lower or greater than the standard). We constructed 19 different lists of numbers with these features, from which one was randomly chosen at the beginning of each task-session. Trials were interleaved by an intertrial interval (ITI) that was randomly drawn from a uniform distribution: $ITI \sim U(700 \text{ ms}, 1000 \text{ ms})$.

In each trial, a target number was presented in bold Arial font (the actual size in pixels depended on the device) for a maximum of 2 seconds. Responses were made either using the keyboard arrows (for desktop users) or tapping the screen on one of two button boxes presented down and to the left and right of the target (for mobile device users). The opacity of these boxes changed on-response as visual feedback. A left arrow keypress (box tap)

indicated a “smaller than 65” response and a right arrow keypress (box tap) indicated a “greater than 65” response. For each decision, we recorded response time (RT) and type of response (smaller or greater).

Programming in Shiny

The implementation of the task was largely based on ShinyPsych (Steiner et al., 2020) a recently developed R package for behavioral experiments using Shiny (Chang et al., 2018). The main rationale behind this implementation is that each element in the flow of the app (i.e. instructions, training, experimental phase, demographics survey and debriefing) is built as a separate “page” that is rendered upon user interactions with the app. This package provides a straight-forward way to specify elements within a page and comes with off-the-shelf examples that can be used as templates. However, as many specific aspects of our task (including a flexible implementation to allow functionality across devices and systems) were not yet fully supported by ShinyPsych, we used R/HTML/CSS/JavaScript code to further customize the app code and address these aspects. The main features we included were: collection of device type (mobile or desktop), web-browser and operating system of participants’ devices; variable inter-trial intervals within which a fixation cross was displayed; scoring system; task progress bar; device-dependent visual parameters’ specification (e.g. text size); device-dependent responses (arrow-keys for desktop and tapping for mobile users), and RT collection with a temporal deadline (upon which a custom timed-out message popped-out for a fixed time). The reliability of the latter is important, as one of our main goals was to infer the numerical-distance effect from RT data. Our app (as any Shiny app) uses JavaScript under the hood to record events and user-interactions within an HTML page. Specifically, RT is defined as the interval between two specific events: stimulus (number) onset and user response. To represent this logic, we coded JavaScript

functions that -on each trial- counted time elapsed from number onset to user-response action (tap or keypress, for mobile and desktop devices, respectively) to obtain RT.

Procedure

Subjects accessed the app online via a shortened URL link (that we distributed across social media: Facebook, Twitter, WhatsApp) which redirected them to a Shiny server where the app was hosted (running in a Donweb (donweb.com) cloud server with 2GB of memory). Upon connection, instructions were displayed and informed consent was prompted. As in Dehaene et al., 1990, we instructed subjects to be as accurate and fast as possible. However, to further enforce these instructions we applied a time deadline (2 seconds) and a scoring system (+2 points for correct answers, -1 point for errors and timeouts). This scoring system also served as a way to “gamify” the app and to promote engagement in the task. We also included a progress bar that allowed users to track their progress, thus reducing participant-dropouts due to not knowing how much more time they would have to spend in the task.

The experimental phase (comprising 116 trials) started after 4 training trials. Afterwards, participants were asked age, sex, and education level with a short in-app survey. Finally, performance feedback (total points) was provided and a short debriefing text was shown.

In total, a typical app session lasted approximately 7 minutes (and it could be completed in less than 5 minutes).

Data Analysis & Statistical Models

To determine whether our results support the numerical distance effect, rather than relying on ANOVA and OLS regressions, we based our analyses on generalized linear mixed models (GLMM) that were estimated with Bayesian methods. GLMMs extend the scope of

general linear models to all the exponential family of probability distributions, which include the (shifted) log-Normal and Bernoulli probability density functions which we used to model RT and error rate, respectively. Moreover, random-effects (multi-level) modeling (the “mixed” part of GLMM) allowed us to represent within-subject data-dependency through a grouping variable which, at the same time, increased model parsimony and allowed for generalization in inference and interpretation. We created our GLMMs within the Stan computational framework (Carpenter et al., 2017) accessed with *brms* (Bürkner, 2017), which uses a No-U-Turn Sampler (NUTS, an extension of the Hamiltonian Monte Carlo sampling method) for estimation of the posterior probability. To improve convergence and guard against overfitting, we used weakly informative -yet regularizing- priors (Gelman et al., 2008).

Our modelling strategy followed a sequential rationale: model specification based on hypothesis and descriptive data-analysis, prior predictive checking, estimation, diagnosis, validation, and inference (i.e. hypothesis testing).

Next, we report model specification for the RT and error rate models.

Response times

Previous studies on the numerical distance have showed that the main predictor of mean RT is the log of the absolute difference between the target (T) and the standard (S) numbers (Dehaene et al., 1990; Hinrichs et al., 1981; Nuerk et al., 2001), $D = abs(T - S)$, i.e. $RT \sim \log(D)$. We hereby follow this mapping. However, given that RTs are positive by definition, have a right-skewed distribution and the minimum RT is necessarily greater than the minimum time used to encode stimuli, we model RT as a shifted-logNormal random variable (Wagenmakers & Brown, 2007):

$$\log(RT - \theta) \sim Normal(\mu, \sigma)$$

This probability density function is defined by 3 parameters: μ (mean), σ (SD) and θ (non-decision time or shift). The latter accounts for the delay in recorded RT from stimulus perceptual-encoding and from other sources that may systematically induce a delay, such as propagation delays in computer networks or data communication. To render a linear relationship between mean RT and $\log(D)$, we model μ as the logarithm of a linear predictor (ν), which includes, among others -as we will detail below-, the coefficient for $\log(D)$.

$$RT \sim e^{Normal(\log(\nu), \sigma)} + \theta$$

This way, mean RT maps linearly onto $\log(D)$, i.e., $E[RT] = \nu$, and the coefficients from the linear predictor of ν can be interpreted analogously to the previous studies.

Although the main focus of our analysis was to answer whether we could observe a numerical distance effect, we also modeled other aspects of our data. Specifically, we aimed to understand whether the participant's device type (e.g., mobile, desktop computer), defined by its operating system (OS, e.g. Android, Linux) and web-browser (e.g. Chrome, Firefox) had an effect on the mean, standard deviation and/or shift of RT. Thus, we included in the model a predictor variable that collectively represents these features of the participant's interface that we call *system*. Altogether, the predictor variables of ν , σ , and θ are:

logD, a numerical variable indicating the log of the absolute numerical distance between the target for each trial and the standard (the minimum $\log D$ is $\log(1)=0$ and the maximum is $\log(29) \cong 3.37$).

smaller, a categorical variable indicating whether the target in a given trial was smaller (1) or greater (0) than the standard (this is relevant because previous studies show that numerical comparisons are slower for smaller than for greater numbers (Dehaene et al., 1990)).

system, a categorical variable indicating the participant device's OS and web-browser. In our data we found 10 categories: 4 mobile-device systems (Android.Chrome, iPhone.Mozilla, iPhone.Safari and iPad.Mozilla) and 6 desktop systems (Windows.Chrome, Windows.Firefox, Linux.Chrome, Linux.Firefox, Macintosh.Chrome and Macintosh.Safari). These categories are indexed with sub-index j ($j=1,...,9$) and we used "Android.Chrome" as the reference category.

trial.number, a numerical variable indicating the trial number (centered and scaled) within each participant session, to control for chronological effects (i.e. fatigue, learning, hurry-to-finish) (the minimum is 1 and the maximum is 116).

age, a numerical variable (centered and scaled) to control for participants' age effects (the minimum is 18 and the maximum 72).

Next, we report the specification for the linear predictors of ν , σ , and θ . Note that we use a log-link function for the linear predictors of σ and θ to aid model convergence.

$$\begin{aligned} \nu_i = & \beta_{Intercept, id_i} + \\ & \beta_{trial.number, id_i} \times \mathbf{trial.number}_i + \\ & \beta_{smaller, id_i} \times \mathbf{smaller}_i + \\ & \beta_{\log D, id_i} \times \mathbf{logD}_i + \\ & \beta_{age} \times \mathbf{age}_i + \\ & \beta_{system, j} \times \mathbf{system}_{i,j} + \\ & \beta_{\log D \times smaller} \times \mathbf{logD}_i \times \mathbf{smaller}_i \end{aligned} \quad (1)$$

$$\begin{aligned} \log(\sigma_i) = & \gamma_{Intercept} + \\ & \gamma_{system, j} \times \mathbf{system}_{i,j} + \\ & \gamma_{age} \times \mathbf{age}_i \end{aligned} \quad (2)$$

$$\begin{aligned} \log(\theta_i) = & \delta_{Intercept, id_i} + \\ & \delta_{system, j} \times \mathbf{system}_{i,j} + \\ & \delta_{age} \times \mathbf{age}_i \end{aligned} \quad (3)$$

Note, id : participant ID; i : trial index; j : system index.

We represented the hierarchical structure of the data (i.e. the within-subjects design) by modelling random effects by participant ID (id) over the intercept of the linear predictor for v and θ , and over the coefficients for the effects of $\log D$, *smaller*, and *trial.number* on v . Moreover, we modeled the correlation between the random effects over these coefficients, and between the intercepts of v and θ . Next, we show how we represented these relations.

$$\begin{aligned} \begin{bmatrix} \beta_{Intercept, id} \\ \delta_{Intercept, id} \end{bmatrix} & \sim MVNormal\left(\begin{bmatrix} \beta_{Intercept} \\ \delta_{Intercept} \end{bmatrix}, S_{Intercept}\right) \\ S_{Intercept} & = \begin{pmatrix} \sigma_{\beta_{Intercept}} & 0 \\ 0 & \sigma_{\delta_{Intercept}} \end{pmatrix} \times R_{Intercept} \times \begin{pmatrix} \sigma_{\beta_{Intercept}} & 0 \\ 0 & \sigma_{\delta_{Intercept}} \end{pmatrix} \\ \begin{bmatrix} \beta_{trial.number, id} \\ \beta_{smaller, id} \\ \beta_{\log D, id} \end{bmatrix} & \sim MVNormal\left(\begin{bmatrix} \beta_{trial.number} \\ \beta_{smaller} \\ \beta_{\log D} \end{bmatrix}, S_{\beta}\right) \\ S_{\beta} & = \begin{pmatrix} \sigma_{\beta_{trial.number}} & 0 & 0 \\ 0 & \sigma_{\beta_{smaller}} & 0 \\ 0 & 0 & \sigma_{\beta_{\log D}} \end{pmatrix} \times R_{\beta} \times \begin{pmatrix} \sigma_{\beta_{trial.number}} & 0 & 0 \\ 0 & \sigma_{\beta_{smaller}} & 0 \\ 0 & 0 & \sigma_{\beta_{\log D}} \end{pmatrix} \end{aligned}$$

Note, *MVNormal*: Multivariate Normal; S : covariance matrix; R : correlation matrix.

Next, we report prior specifications for all the estimated model parameters.

$$\beta_{Intercept} \sim Normal(0.6, 0.3)$$

$$\left(\beta_{trial.number}, \beta_{smaller}, \beta_{\log D}, \beta_{age}, \beta_{system,j}, \beta_{\log D \times smaller} \right) \sim Normal(0, 1)$$

$$\gamma_{Intercept} \sim Normal(-1, 0.2)$$

$$\left(\gamma_{system,j}, \gamma_{age} \right) \sim Normal(0, 0.1)$$

$$\delta_{Intercept} \sim Normal(-1.3, 0.2)$$

$$\left(\delta_{system,j}, \delta_{age} \right) \sim Normal(0, 0.1)$$

$$\left(\sigma_{\beta_{Intercept}}, \sigma_{\beta_{trial.number}}, \sigma_{\beta_{smaller}}, \sigma_{\beta_{\log D}} \right) \sim Normal(0.2, 0.2)$$

$$\sigma_{\delta_{Intercept}} \sim Student(3, 0, 2.5)$$

$$\left(R_{Intercept}, R_{\beta} \right) \sim LKJcorr(1)$$

Error rate

We modeled errors as a Bernoulli random variable (i.e. a binary categorical variable where 1 = Error and 0 = Correct) with a logit link to the linear predictor of the probability π of making an error (Equation (4)), which included the predictor variables *trial.number* and *logD*. That is, for each trial i , we estimated the probability π_i of making an error on that trial, conditional on these predictor variables.

$$Error_i \sim Bernoulli(\pi_i)$$

$$\begin{aligned}
 \text{logit}(\pi_i) = & \varphi_{Intercept, id_i} + \\
 & \varphi_{trial.number, id_i} \times \text{trial.number}_i + \\
 & \varphi_{smaller, id_i} \times \text{smaller}_i + \\
 & \varphi_{\log D, id_i} \times \log D_i + \\
 & \varphi_{age} \times \text{age}_i + \\
 & \varphi_{system, j} \times \text{system}_{i,j} + \\
 & \varphi_{\log D \times smaller} \times \log D_i \times \text{smaller}_i
 \end{aligned} \tag{4}$$

Similarly to the RT model, we represented the hierarchical structure of the data by modelling random effects by participant ID (*id*) over the intercept and the coefficient for *trial.number* of the linear predictor for π , and modeled the correlation between these random effects. Next, we show how we represented these relations.

$$\begin{bmatrix} \varphi_{Intercept, id} \\ \varphi_{trial.number, id} \\ \varphi_{smaller, id} \\ \varphi_{\log D, id} \end{bmatrix} \sim MVNormal \left(\begin{bmatrix} \varphi_{Intercept} \\ \varphi_{trial.number} \\ \varphi_{smaller} \\ \varphi_{\log D} \end{bmatrix}, S \right)$$

$$S = \begin{pmatrix} \sigma_{\varphi_{Intercept}} & 0 & 0 & 0 \\ 0 & \sigma_{\varphi_{trial.number}} & 0 & 0 \\ 0 & 0 & \sigma_{\varphi_{smaller}} & 0 \\ 0 & 0 & 0 & \sigma_{\varphi_{\log D}} \end{pmatrix} \times R \times \begin{pmatrix} \sigma_{\varphi_{Intercept}} & 0 & 0 & 0 \\ 0 & \sigma_{\varphi_{trial.number}} & 0 & 0 \\ 0 & 0 & \sigma_{\varphi_{smaller}} & 0 \\ 0 & 0 & 0 & \sigma_{\varphi_{\log D}} \end{pmatrix}$$

Next, we report prior specifications for all the estimated model parameters.

$$(\varphi_{Intercept}, \varphi_{trial.number}, \varphi_{smaller}, \varphi_{\log D}, \varphi_{age}, \varphi_{system, j}, \varphi_{\log D \times smaller}) \sim Normal(0, 1)$$

$$(\sigma_{\varphi_{Intercept}}, \sigma_{\varphi_{trial.number}}, \sigma_{\varphi_{smaller}}, \sigma_{\varphi_{\log D}}) \sim Student(3, 0, 2.5)$$

$$R \sim LKJcorr(1)$$

To diagnose model estimation, we evaluated traces autocorrelation and number of effective samples, visually inspected chains with trace plots for each model parameter, and assessed whether $\hat{R}$ values were smaller than 1.01 (denoting that chains mixed well) using *launch_shinystan* (from *rstanarm* (Goodrich et al., 2018)). For both models, we estimated 20000 samples (4 chains of 6000 traces each with the first 1000 for warm up) with an adaptive delta of 0.99 to minimize divergent transitions after warmup. Indeed, there were zero divergent transitions during sampling of both models.

We validated the models with posterior predictive checks (PPC) by simulating 2000 new datasets from the estimated parameters' posterior distributions, and overlaying the simulated distributions over the observed data distribution. For both models, the posterior predictive distributions were qualitatively indistinguishable from the observed distributions. To quantify this qualitative assessment, we computed -for the error model- the proportion of errors and -for the RT model- the 0.1, 0.3, 0.5 (median), 0.7, and 0.9 quantiles and SD of the conditional (by *smaller*, i.e. smaller or greater than 65, and by system) RT distributions of each set. We then obtained the median and the quantile-based credible intervals (containing 95% of the probability density, CI_{95}) of the distribution of these summary-values and compared it to the observed-data summaries (assessing whether the observed value was contained in the CI_{95}). This approach not only allowed us to assign credibility to the model, but also to have a measure of uncertainty derived from the distribution of simulated summary statistics. For the error rate model, the CI_{95} of the conditional summary statistics included all the summary statistics of the observed data. Similarly, for the RT model, this was generally true, except that -mainly- for the 0.9 quantile for some categories of system, the predictions were off by -at most- 10ms from the observed summary statistic.

Additionally, we computed a measure of in-sample goodness-of-fit for each model. For the RT model, we computed a Bayesian approximation to R^2 (Gelman et al., 2019) and

found that $R^2 = 0.438$ ($CI_{95} = [0.429, 0.447]$). For the error rate model, we computed the area under the Receiver Operating Characteristic curve (AUC-ROC), which reflects the classifying ability of the model predictions (as a rule-of-thumb, AUC-ROC = 0.8 denotes a good model performance (Hosmer Jr et al., 2013)), and found that AUC-ROC = 0.813.

Finally, we tested our hypotheses directly over the posterior probability distribution by computing the CI_{95} of the coefficient-of-interest distribution and assessing whether this interval contained (or for the logD predictor, was smaller than) 0. If the CI_{95} did not include 0 (or if 95% of the density was smaller than 0, for logD), the results support the hypothesis that the corresponding coefficient is different from (smaller than) that null value. To aid interpretation of the error rate model, we first exponentiated the coefficient-of-interest posterior values to obtain odds ratios, and then computed the summary values of these posterior values (median and CI_{95}).

Results

We found evidence consistent with a numerical-distance effect on RT. As expected, the greater the absolute difference between target and standard, the faster the responses (Fig. 2A). Regression coefficients for our model (Equations 1-3) show that on average RTs decrease 70ms for every unit increase in logarithmic numerical distance, e.g. roughly the distance from 66 ($\log_{(D=1)} = 0$) to 68 ($\log_{(D=3)} \cong 1$) ($\beta_{\log D} = -70.3ms$, $CI_{95} = [-75.4, -65.6]ms$). Strikingly, our estimated value of $\beta_{\log D}$ closely resembles that estimated from the data shown in Dehaene et al., 1990¹ ($\beta_{\log D \text{ (Dehaene et al., 1990)}} = -78.3ms$), even when we used a different statistical model for RT and different technology to implement the task. Also consistent with previous studies (Dehaene, 1989; Dehaene et al., 1990; Hinrichs et al., 1981), we found that RTs for numbers smaller than 65 were slower than RTs for numbers greater than 65 (

$\beta_{smaller} = 55.8ms$, $CI_{95} = [39.7, 72.3]ms$). Again, our estimated value for $\beta_{smaller}$ closely matches that from Dehaene et al., 1990 ($\beta_{smaller (Dehaene et al., 1990)} = 53.8ms$).

We detected chronological effects on RT by which for every 35 completed trials, correct responses were made on average 20ms faster (

$\beta_{trial.number} = -20.6ms$, $CI_{95} = [-26.1, -15.3]ms$). Further, we found that older participants were predominantly slower than younger ones (

$\beta_{age} = 10.1ms$, $CI_{95} = [-2.18, 22.4]ms$). This value corresponds to an average increment of ~9ms for every 10 years of age. Indeed, the impact of age over RTs is a well-documented effect in binary RT decision-making tasks (Ratcliff et al., 2004).

Our modeling strategy also allows us to understand whether and how the participants' system (i.e. device, OS, and web-browser) affected their RTs. We found that the system indeed affected RT delays (Fig. 3) but we did not find evidence for an effect of the system on the mean nor on the SD of RTs (Table 1). The main pattern we observed is that RTs from mobile devices (e.g., Android.Chrome) had larger delays than RTs from desktop computers (e.g. Windows.Chrome) ($\Delta RT_{delay(mobile-deskto)} = 163ms$, $CI_{95} = [142, 185]ms$), which is consistent with previous reports (Holden et al., 2019; Pronk et al., 2020; Reimers & Stewart, 2008).

Error rates were also revealing of a numerical distance effect, by which the greater the numerical distance, the smaller the rate of errors. Consistent with previous studies (Dehaene et al., 1990; Hinrichs et al., 1981), we observed that participants made ~10% errors for the hardest numerical comparisons (e.g., 64 vs 65, 66 vs 65), and that the error rates dropped as the numerical distance increased (Fig. 2B). Specifically, we found that for a unit increase in

logarithmic numerical distance, the odds of making an error decreased by 50% on average (

$$\varphi_{\log D} = -0.71, CI_{95} = [-0.80, -0.62]).$$

Although the task was rather short (116 trials), we found a relevant chronological effect ($\varphi_{\text{trial.number}} = -0.15, CI_{95} = [-0.27, -0.03]$). This implies that for -roughly- every 35 completed trials, the odds of making an error decrease by 15%. This effect possibly represents a learning process by which less errors are made as the task unfolds.

To summarize, our results show that the numerical-distance effect on both RT and error rate can be observed with data collected in an online experiment. Notably, the magnitude of the effects of numerical distance closely match previous results. Moreover, we found that the system that the participant used to complete the task (i.e. the OS and web-browser of their devices) reliably affected RT delays, but we did not find evidence for changes in the mean nor on the SD of RTs. We quantified the change in delay for each different system used by the participants and found that mobile devices generally induced larger delays than desktop systems.

Discussion

Collecting behavioral data online is attractive, but doing it wisely warrants precaution. In particular, many paradigms in perceptual and cognitive decision-making rely on measuring RTs at the millisecond-level. It is far from obvious that reliable RTs can be obtained using an online setup. Moreover, we need not forget that individuals will participate using different OS, web-browsers, and devices, and will possibly get distracted during the task. In a sense, it is like running an experiment with different setups in an uncontrolled environment.

In this paper, we show that -beyond all limitations- it is indeed possible to achieve high-standard data quality (including precise response-timing) in binary decisions using an

R-Shiny application running on desktop systems and mobile devices. As a case study, we replicated the numerical distance effect, a classical numerical cognition finding. A long-standing research on this phenomenon robustly shows that RTs and error rate decrease when the absolute difference between the numbers to be compared increases (Dehaene et al., 1990; Hinrichs et al., 1981; Krajcsi & Kojouharova, 2017; R. S. Moyer & Landauer, 1973; Verguts & De Moor, 2005, p.).

Behavioral studies performed 35-40 years ago found that RTs map linearly onto the log of the numerical distance ($\log D$) with a regression coefficient of

$\beta_{\log D \text{ (Dehaene et al., 1990)}} = -78.3ms$ when the comparison task uses 65 as the standard (Dehaene et al., 1990) with small variations when the standard is set to 55 or 75 (Dehaene, 1989; Dehaene et al., 1990; Hinrichs et al., 1981). Considering that huge changes in technology occurred in the last four decades and that new statistical methods are becoming standard, it was far from obvious that we would find consistent evidence in an online study run mostly on mobiles in fairly uncontrolled conditions.

Although there are discussions regarding the precise cognitive mechanism that underlies this phenomenon (Dehaene, 1989; Dehaene et al., 1990; Hinrichs et al., 1981; Krajcsi & Kojouharova, 2017; R. S. Moyer & Landauer, 1973; Nuerk et al., 2001; Verguts & De Moor, 2005), our work was not aimed to discern between these accounts. Rather, we sought to provide a descriptive proof of consistency of the numerical distance effect inferred from online data collected with an R-Shiny app.

Importantly, our results also highlight the relevance of identifying features of the participant's device such as web-browser and OS, to dissect the technological component from the psychological component of RTs. Here, we were able to collect data from 10 different combinations of OS, web-browser, and type of device (mobile/desktop) that we

collectively called *systems*. However, it is important to note that these categories were not equally represented in our data. In fact, most participants completed the task using Chrome browser (running either on Windows (desktop) or Android (mobile) OS). Nonetheless, our interpretation of the results stems from the comparison between the effects of these two more-frequent systems on RT-delays. These differences might reflect OS processing times, the response mode (finger tap for mobile devices; keypress for desktops), and/or features related to the use of these devices (e.g. desktops likely have a WiFi connection, while mobile devices might be connected to a slower 4G network).

In the last decade, many experimental tools have been developed with the aim of providing a robust way to perform experiments online (e.g., jsPsych (de Leeuw, 2015), Cognition (www.cognition.run), JATOS (Lange et al., 2015)). Although it is certainly healthy to have a wide array of alternatives at disposal, all come with an inherent trade-off between flexibility (i.e. easiness for customization) and ease of learning/use. Among these options, Shiny hits a sweet spot, as it is highly *and* easily customizable. Even more, Shiny-app development can be readily integrated into the workflow of experimental psychologists that use R for data analysis without the need to switch back and forth between frameworks and/or programming languages.

Altogether, our study provides a proof of concept of the implementation of simple binary RT decision-making tasks in online experiments with Shiny and paves the way for the development of online experiments with reliable response-timing data collection. We show that online experiments are not only a fruitful avenue for survey-style data-collection, but also for RT-tasks and thus enclose a greater (and still unharnessed) power for experimental psychologists. This is central in the face of the current COVID pandemic that makes in-lab data-collection impossible.

Declarations

Ethics Approval

This study was performed in line with the principles of the Declaration of Helsinki. Approval was granted by the Ethics Committee of Centro de Educación Médica e Investigaciones Clínicas (protocol ID 435). The authors complied with APA ethical standards in the treatment of their sample.

Consent to Participate

Informed consent was obtained from all individual participants included in the study.

Data Accessibility

The datasets generated during and/or analysed during the current study (containing data from 19604 trials) are available in the OSF repository, <https://osf.io/6gmzs>.

Code Accessibility

All code for the app and analyses is available in the OSF repository, <https://osf.io/6gmzs>. A demo version of the app is accessible in ShinyApps, <https://2exp3.shinyapps.io/numbers/>.

Competing interests

The authors have no competing interests to declare.

Authors' contributions

A.P.S. and G.S. designed and performed research, analyzed data, and wrote the paper.

Acknowledgments

This research was supported by the Science and Technology Council of the University of Buenos Aires (UBACYT, Grant 20020170100330BA) and by the National Scientific and Technical Research Council (CONICET grant # PIP 11220150100787CO). We would like to thank Pablo Barttfeld for sharing cloud hosting for our app.

Endnotes

1. Although in this paper the authors report the correlation coefficient between RT and logD, the regression coefficient is not provided. Since the data from this study is not publicly available, we scraped the data from the corresponding figure using WebPlotDigitizer (<https://github.com/ankitrohatgi/WebPlotDigitizer>) and employed the same statistical method described in the paper to obtain a rough estimate of the logD regression coefficient.

References

- Anwyl-Irvine, A., Dalmaijer, E. S., Hodges, N., & Evershed, J. K. (2020). Realistic precision and accuracy of online experiment platforms, web browsers, and devices. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-020-01501-5>
- Brown, H. R., Zeidman, P., Smittenaar, P., Adams, R. A., McNab, F., Rutledge, R. B., & Dolan, R. J. (2014). Crowdsourcing for Cognitive Science – The Utility of Smartphones. *PLOS ONE*, 9(7), e100662. <https://doi.org/10.1371/journal.pone.0100662>
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1), 1–32. <https://doi.org/10.18637/jss.v076.i01>
- Chang, W., Cheng, J., Allaire, J., Xie, Y., & McPherson, J. (2018). shiny: Web application framework for r, 2015. *R Package Version*, 1(0), 14.
- Cipora, K., Soltanlou, M., Reips, U.-D., & Nuerk, H.-C. (2019). The SNARC and MARC effects measured online: Large-scale assessment methods in flexible cognitive effects. *Behavior Research Methods*, 51(4), 1676–1692. <https://doi.org/10.3758/s13428-019-01213-5>
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon’s Mechanical Turk as a Tool for Experimental Behavioral Research. *PLOS ONE*, 8(3), e57410. <https://doi.org/10.1371/journal.pone.0057410>
- Dandurand, F., Shultz, T. R., & Onishi, K. H. (2008). Comparing online and lab methods in a problem-solving experiment. *Behavior Research Methods*, 40(2), 428–434.

<https://doi.org/10.3758/BRM.40.2.428>

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12.

<https://doi.org/10.3758/s13428-014-0458-y>

Dehaene, S. (1989). The psychophysics of numerical comparison: A reexamination of apparently incompatible data. *Perception & Psychophysics*, 45(6), 557–566.

<https://doi.org/10.3758/BF03208063>

Dehaene, S. (2011). *The number sense: How the mind creates mathematics* (Rev. and updated ed). Oxford University Press.

Dehaene, S., Dupoux, E., & Mehler, J. (1990). Is numerical comparison digital? Analogical and symbolic effects in two-digit number comparison. *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 626–641.

<https://doi.org/10.1037/0096-1523.16.3.626>

Dufau, S., Duñabeitia, J. A., Moret-Tatay, C., McGonigal, A., Peeters, D., Alario, F.-X., Balota, D. A., Brysbaert, M., Carreiras, M., Ferrand, L., Ktori, M., Perea, M., Rastle, K., Sasburg, O., Yap, M. J., Ziegler, J. C., & Grainger, J. (2011). Smart Phone, Smart Science: How the Use of Smartphones Can Revolutionize Research in Cognitive Science. *PLOS ONE*, 6(9), e24974. <https://doi.org/10.1371/journal.pone.0024974>

Garaizar, P., & Reips, U.-D. (2019). Best practices: Two Web-browser-based methods for stimulus presentation in behavioral experiments with high-resolution timing requirements. *Behavior Research Methods*, 51(3), 1441–1453.

<https://doi.org/10.3758/s13428-018-1126-4>

Gelman, A., Goodrich, B., Gabry, J., & Vehtari, A. (2019). R-squared for Bayesian Regression Models. *The American Statistician*, 73(3), 307–309.

<https://doi.org/10.1080/00031305.2018.1549100>

- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *Annals of Applied Statistics*, 2(4), 1360–1383. <https://doi.org/10.1214/08-AOAS191>
- Germine, L., Nakayama, K., Duchaine, B. C., Chabris, C. F., Chatterjee, G., & Wilmer, J. B. (2012). Is the Web as good as the lab? Comparable performance from Web and lab in cognitive/perceptual experiments. *Psychonomic Bulletin & Review*, 19(5), 847–857. <https://doi.org/10.3758/s13423-012-0296-9>
- Gökaydin, D., Brugger, P., & Loetscher, T. (2018). Sequential Effects in SNARC. *Scientific Reports*, 8(1), 10996. <https://doi.org/10.1038/s41598-018-29337-2>
- Goodrich, B., Gabry, J., Ali, I., & Brilleman, S. (2018). rstanarm: Bayesian applied regression modeling via Stan. *R Package Version*, 2(4), 1758.
- Gosling, S. D., Vazire, S., Srivastava, S., & John, O. P. (2004). Should We Trust Web-Based Studies? A Comparative Analysis of Six Preconceptions About Internet Questionnaires. *American Psychologist*, 59(2), 93–104. <https://doi.org/10.1037/0003-066X.59.2.93>
- Grootswagers, T. (2020). A primer on running human behavioural experiments online. *Behavior Research Methods*, 52(6), 2283–2286. <https://doi.org/10.3758/s13428-020-01395-3>
- Hinrichs, J. V., Yurko, D. S., & Hu, J. (1981). Two-digit number comparison: Use of place information. *Journal of Experimental Psychology: Human Perception and Performance*, 7(4), 890–901. <https://doi.org/10.1037/0096-1523.7.4.890>
- Holden, J., Francisco, E., Lensch, R., Tommerdahl, A., Kirsch, B., Zai, L., Dennis, R., & Tommerdahl, M. (2019). Accuracy of different modalities of reaction time testing: Implications for online cognitive assessment tools. *BioRxiv*, 726364. <https://doi.org/10.1101/726364>

Hosmer Jr, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied Logistic Regression*.

John Wiley & Sons.

Huber, S., Nuerk, H.-C., Reips, U.-D., & Soltanlou, M. (2019). Individual differences

influence two-digit number processing, but not their analog magnitude processing: A large-scale online study. *Psychological Research*, 83(7), 1444–1464.

<https://doi.org/10.1007/s00426-017-0964-5>

Kaufman, A. R. (2020). Implementing novel, flexible, and powerful survey designs in R

Shiny. *PLOS ONE*, 15(4), e0232424. <https://doi.org/10.1371/journal.pone.0232424>

Kochari, A. R. (2019). Conducting Web-Based Experiments for Numerical Cognition

Research. *Journal of Cognition*, 2(1). <https://doi.org/10.5334/joc.85>

Krajcsi, A., & Kojouharova, P. (2017). Symbolic Numerical Distance Effect Does Not

Reflect the Difference between Numbers. *Frontiers in Psychology*, 8.

<https://doi.org/10.3389/fpsyg.2017.02013>

Krantz, J. H., & Dalal, R. (2000). Validity of Web-based psychological research. In

Psychological experiments on the Internet (pp. 35–60). Academic Press.

<https://doi.org/10.1016/B978-012099980-4/50003-4>

Lange, K., Kühn, S., & Filevich, E. (2015). "Just Another Tool for Online Studies" (JATOS):

An Easy Solution for Setup and Management of Web Servers Supporting Online

Studies. *PLOS ONE*, 10(6), e0130834. <https://doi.org/10.1371/journal.pone.0130834>

Moyer, R. S., & Landauer, T. K. (1973). Determinants of reaction time for digit inequality

judgments. *Bulletin of the Psychonomic Society*, 1(3), 167–168.

<https://doi.org/10.3758/BF03334328>

Moyer, Robert S., & Landauer, T. K. (1967). Time required for judgements of numerical

inequality. *Nature*, 215(5109), 1519–1520. <https://doi.org/10.1038/2151519a0>

Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du

- Sert, N., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., & Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1(1), 1–9.
<https://doi.org/10.1038/s41562-016-0021>
- Nuerk, H.-C., Weger, U., & Willmes, K. (2001). Decade breaks in the mental number line? Putting the tens and units back in different bins. *Cognition*, 82(1), B25–B33.
[https://doi.org/10.1016/S0010-0277\(01\)00142-1](https://doi.org/10.1016/S0010-0277(01)00142-1)
- Pronk, T., Wiers, R. W., Molenkamp, B., & Murre, J. (2020). Mental chronometry in the pocket? Timing accuracy of web applications on touchscreen and keyboard devices. *Behavior Research Methods*, 52(3), 1371–1382.
<https://doi.org/10.3758/s13428-019-01321-2>
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A Diffusion Model Analysis of the Effects of Aging in the Lexical-Decision Task. *Psychology and Aging*, 19(2), 278–289. <https://doi.org/10.1037/0882-7974.19.2.278>
- Reimers, S., & Stewart, N. (2008). Using Adobe Flash Lite on mobile phones for psychological research: Reaction time measurement reliability and interdevice variability. *Behavior Research Methods*, 40(4), 1170–1176.
<https://doi.org/10.3758/BRM.40.4.1170>
- Reips, U.-D. (2002). Standards for Internet-Based Experimenting. *Experimental Psychology*, 49(4), 243–256. <https://doi.org/10.1026//1618-3169.49.4.243>
- Semmelmann, K., & Weigelt, S. (2017). Online psychophysics: Reaction time effects in cognitive experiments. *Behavior Research Methods*, 49(4), 1241–1260.
<https://doi.org/10.3758/s13428-016-0783-4>
- Silver, L., Smith, A., Johnson, C., Taylor, K., Jiang, J., Anderson, M., & Rainie, L. (2019). Mobile connectivity in emerging economies. *Pew Research Center*, 7.
- Simcox, T., & Fiez, J. A. (2014). Collecting Response Times using Amazon Mechanical Turk

and Adobe Flash. *Behavior Research Methods*, 46(1), 95–111.

<https://doi.org/10.3758/s13428-013-0345-y>

Steiner, M., Phillips, N., & Trutmann, K. (2020). *ShinyPsych: An easy way to program psychology experiments using Shiny*.

Verguts, T., & De Moor, W. (2005). Two-digit Comparison. *Experimental Psychology*, 52(3), 195–200. <https://doi.org/10.1027/1618-3169.52.3.195>

Wagenmakers, E.-J., & Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review*, 114(3), 830–841. <https://doi.org/10.1037/0033-295X.114.3.830>

Zimmerman, F., Shalom, D., Gonzalez, P. A., Garrido, J. M., Heduan, F. A., Dehaene, S., Sigman, M., & Rieznik, A. (2016). Arithmetic on Your Phone: A Large Scale Investigation of Simple Additions and Multiplications. *PLOS ONE*, 11(12), e0168431. <https://doi.org/10.1371/journal.pone.0168431>

Figures and Tables

Figure 1. Number-comparison task.

Participants had to decide whether a target number (49, in this example) was smaller or greater than a standard (65 in this study). We presented 58 target numbers: 29 were lower than the standard (from 36 to 64), and 29 were greater than the standard (from 66 to 94). Each of these targets was displayed twice, rendering a total of 116 trials. Trials were interleaved by an intertrial interval (ITI) that was randomly drawn from a uniform distribution ($ITI \sim U(700 \text{ ms}, 1000 \text{ ms})$). Responses were made either using the keyboard arrows (for desktop users) or tapping the screen on one of two button boxes presented down and to the left and right of the target (for mobile device users). A left-arrow keypress (box tap) indicated a “smaller than 65” response and a right-arrow keypress (box tap) indicated a “greater than 65” response. For

each decision, we recorded response time (RT) and type (smaller or greater). Participants were instructed to be as accurate and fast as possible. However, to further enforce these instructions we applied a time deadline (each target was displayed for a maximum of 2 seconds) and a scoring system (+2 points for correct answers, -1 point for errors and timeouts). Last, we included a progress bar (in light-blue (grey) at the top) to minimize dropouts due to uncertainty about how long the task would take.

Figure 2. Numerical distance effect on RT and error rate.

(A) Expected and observed values for RT as a function of the numerical distance between the target and the standard (65, dashed vertical line). Green (gray) dots are averages of by-participant medians of RT for each target number. The black line and shaded area represent the median and CI_{95} of the RT expected values (for an average participant), respectively, obtained from 2000 model posterior samples. (B) Expected and observed values of error rate (i.e., probability of making an error) as a function of the numerical distance. Green (gray) dots are averages of by-participant proportions of errors for each target number. The black line and shaded area represent the median and CI_{95} of the expected values (for an average participant), respectively, obtained from 2000 model posterior samples. Note that, as

not all systems were equally represented in our data (e.g., there was only 1 session from an iPad device), we obtained and represented the weighted (by the frequency of the participants systems) marginal means for each target number in both panels. All means were computed for mean participants' age and at mean trial number (i.e., at the middle of the task).

Figure 3. RT delays by participants' system.

We computed the conditional means by system (i.e., for each of the 10 categories of systems, see Method) using 2000 posterior samples from the estimated coefficients of the linear predictor of θ (Equation(3)). The colored points' position and grey horizontal intervals represent the median and CI_{95} of the posterior samples for each conditional mean, respectively. We used dark pink (dark gray) and blue (gray) points to represent mobile and desktop devices, respectively, to show that mobile devices induce larger delays than desktop devices. The annotation above each point reports the number of participants (N) that completed the task using that specific system. All means were computed for mean participants' age.

Table 1. *Estimated Effects of the Participants' System Over ν and σ .*

Distributional parameter	Coefficient	Median	Lower	Upper
ν	$\beta_{\text{system, Windows.Chrome}}$	-0.009	-0.037	0.019
ν	$\beta_{\text{system, Windows.Firefox}}$	-0.030	-0.086	0.025
ν	$\beta_{\text{system, Linux.Chrome}}$	-0.041	-0.097	0.015
ν	$\beta_{\text{system, Linux.Firefox}}$	0.006	-0.130	0.141
ν	$\beta_{\text{system, Macintosh.Chrome}}$	0.010	-0.126	0.148
ν	$\beta_{\text{system, Macintosh.Safari}}$	0.025	-0.125	0.178
ν	$\beta_{\text{system, iPhone.Safari}}$	-0.001	-0.043	0.042
ν	$\beta_{\text{system, iPhone.Mozilla}}$	-0.012	-0.057	0.033
ν	$\beta_{\text{system, iPad.Mozilla}}$	-0.061	-0.210	0.095
σ	$\gamma_{\text{system, Windows.Chrome}}$	-0.010	-0.059	0.039
σ	$\gamma_{\text{system, Windows.Firefox}}$	0.037	-0.052	0.123
σ	$\gamma_{\text{system, Linux.Chrome}}$	-0.054	-0.140	0.031
σ	$\gamma_{\text{system, Linux.Firefox}}$	-0.090	-0.229	0.047
σ	$\gamma_{\text{system, Macintosh.Chrome}}$	-0.049	-0.202	0.100
σ	$\gamma_{\text{system, Macintosh.Safari}}$	0.092	-0.051	0.234
σ	$\gamma_{\text{system, iPhone.Safari}}$	0.044	-0.027	0.111
σ	$\gamma_{\text{system, iPhone.Mozilla}}$	-0.007	-0.088	0.069
σ	$\gamma_{\text{system, iPad.Mozilla}}$	0.119	-0.059	0.283

Note, Lower and Upper refer to the lower and upper bounds of the CI_{95} of each coefficient posterior distribution. Importantly, all these intervals contain 0, suggesting that -at least from our data- the participants' system did not induce any systematic change on either the mean or the SD of RTs.