

1 **An empirical test of the role of value certainty in decision making**

2 Douglas Lee¹, Giorgio Coricelli¹

3 ¹ University of Southern California, Department of Economics

4

5

6

7

8

9

10 Address for correspondence:

11

12

13 Tel:

14 E-mail: DouglasGLee@gmail.com

15

16 Keywords:

17 Word count:

18

19

20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

Abstract

Most contemporary models of value-based decisions are built on value estimates that are typically self-reported by the decision maker. Such models have been successful in accounting for choice accuracy and response time, and more recently choice confidence. The fundamental driver of such models is choice difficulty, which is almost always defined as the absolute value difference between the subjective value ratings of the options in a choice set. Yet a decision maker is not necessarily able to provide a value estimate with the same degree of certainty for each option that he encounters. We propose that choice difficulty is determined not only by absolute value distance of choice options, but also by their value certainty. In this study, we first demonstrate the reliability of the concept of an option-specific value certainty using three different experimental measures. We then demonstrate the influence that value certainty has on choice, including accuracy (consistency), choice confidence, response time, and choice-induced preference change (i.e., the degree to which value estimates change from pre- to post-choice evaluation). We conclude with a suggestion of how popular contemporary models of choice (e.g., race model, drift-diffusion model) could be improved by including option-specific value certainty as one of their inputs.

38 INTRODUCTION

39 For many decades, researchers in economics, psychology, and cognitive neuroscience have
40 studied the concept of subjective value, and how implicit subjective value influences explicit
41 choices. In more recent years, decision researchers have frequently relied on self-reported
42 estimates of subjective value (value ratings) to support their theories and models (see Rangel et
43 al, 2008, for a review). Value ratings collected by self-report have served as the fundamental
44 input of most contemporary models of choice. One key component of many models is what is
45 referred to as “choice difficulty”, which is most often defined as the distance between the
46 subjective value ratings of the options in the choice set (where difficulty declines with distance).
47 Choice difficulty has been shown to reliably predict both choice (i.e., the probability of choosing
48 the higher-rated option decreases with difficulty) and reaction time (i.e., deliberation time
49 increases with difficulty) (e.g., Hunt et al, 2012; Palmer et al, 2005). Yet until recently, most
50 models did not explicitly incorporate the possibility that a decision maker (DM) might not be fully
51 certain about the value ratings that he reports. For example, a DM might have an estimate about
52 the (subjective) value of an option, yet simultaneously have a belief about the accuracy of his
53 value estimate. Said another way, sometimes we might feel that we like something a certain
54 amount, but also feel more or less sure about that amount—sometimes we might feel that we
55 know for sure precisely how much we value something, other times we might feel unsure about
56 our value estimates. Consequently, when choosing between items, certainty about the values of
57 the individual options directly impacts choice confidence, which can be defined as a feeling of
58 certainty about which of the options has a higher value. The more uncertain the individual option
59 values are, the more difficult it is to determine which one has the higher value. It would therefore

60 follow that a more complete definition of choice difficulty should account for such a concept as
61 subjective certainty.

62 Some recent studies have indeed included value certainty as an independent variable in their
63 models. DeMartino et al (2013) examined the effects of what they called “bid confidence” on
64 the dependent variables in their model. The authors suggested that bid certainty should
65 moderate the impact of choice difficulty (defined as value rating distance) on choice accuracy,
66 choice confidence, and choice response time (RT). Lee and Daunizeau (2020a) took this a step
67 further, demonstrating in their model how value certainty might be explicitly involved in choice.
68 In that study, the authors suggested that when value ratings are uncertain, a DM will be less
69 accurate and less confident in his choice, but he will also be more inclined to invest mental effort
70 (for which RT could serve as a proxy) in order to decrease the value uncertainty and enable him
71 to confidently choose his preferred option. In fact, in that model, one of the principal drivers of
72 the proposed effort-confidence tradeoff is the desire on the part of the DM to increase the
73 certainty that he has about his value estimates. The model builds upon previous work that
74 demonstrated the same principle in a less formal manner, by showing that both RT and
75 preference change (difference between post- and pre-choice ratings) were decreasing functions
76 of value certainty (for which the authors used the average value certainty of the options being
77 compared), while choice confidence was an increasing function of value certainty (Lee and
78 Daunizeau, 2020a). In their model, the authors explain that lower value certainty impairs the
79 ability of the DM to distinguish the options. This dampens choice confidence, which the DM
80 attempts to boost through mental effort (proxied by RT). In turn, the effort allocation leads to
81 value estimate refinements, and potentially changes of mind (preference reversals).

82 Recent work in other, non-subjective value-based domains also suggests that a measure of
83 certainty about the options is conceptually important when examining choice behavior (Pouget
84 et al, 2016). Frydman and Jin (2019) invoke the principle of efficient coding to suggest a link
85 between value certainty and choice behavior. In this study, certainty spawns from repeated
86 exposure, which causes greater precision in neural representation (i.e., efficient coding). Higher
87 certainty, thus defined, leads to higher choices accuracy (i.e., consistency with value estimates).
88 Padoa-Schioppa and Rustichini (2014) illustrate a similar concept based on *adaptive* coding,
89 where neural activity is normalized according to the range of option values in the current
90 environment, thus causing choice stochasticity to increase as representation precision decreases.
91 Along similar lines, Woodford (2019) explains how many important aspects of economic choice
92 (e.g., choice stochasticity, risk aversion, decoy effects) could result from noisy neural
93 representations of value. Although the author does not directly refer to a subjective feeling of
94 certainty, it is no far stretch to relate the precision of neural representations to a subjective
95 feeling of certainty. In Polania et al (2019), the authors refer to efficient coding as well as
96 Bayesian decoding principles to explain how choice behavior is influenced by value certainty.
97 Interestingly, the authors do not record any self-reports of value certainty to validate their model.
98 Instead, certainty in this study is captured by rating consistency (i.e., similarity of repeatedly self-
99 reported value estimates). The basic idea here is that greater precision in the neural encoding of
100 value will lead to greater consistency across multiple interrogations of the value-encoding neural
101 population. This precision thus leads to both more consistent ratings and more choices
102 consistent with those ratings.

103 In spite of recent theoretical and empirical evidence that value certainty plays an important role
104 in choice behavior, most popular models of value-based decision making still do not include this
105 variable. In particular, so-called accumulation-to-bound models such as the race model (e.g.,
106 Tajima et al, 2019; Kepecs et al, 2008) and the drift-diffusion model (e.g., Kiani and Shadlen, 2009;
107 Krajbich and Rangel, 2011) do not include option-specific certainty. Tajima et al (2019) model
108 noise in the accumulation process at the system level, where all options have the same degree
109 of uncertainty imposed upon them by the environment, rather than at the option level (although
110 the authors themselves suggest that future studies should explore the various sources of value
111 uncertainty). Other studies have similarly included systemic, but not option-specific, uncertainty
112 (e.g., Louie et al, 2013). Krajbich and Rangel (2011), along with most other published versions of
113 the drift-diffusion model (DDM), fail to include a variable to represent value (un)certainty.

114 Perhaps more researchers would be willing to include value certainty in their models if there was
115 more available evidence demonstrating that certainty could be reliably measured. In this study,
116 we hope to provide some such evidence. In line with Lee and Daunizeau (2020a, 2020b), we
117 explicitly ask decision makers to report their subjective feelings of certainty about the ratings
118 they provide about each of a large set of options. In line with Polania et al (2019), we also
119 implicitly capture value certainty by calculating consistency across multiple ratings of the same
120 options. We then show that the explicit and implicit measurements of value certainty are highly
121 correlated for each individual DM, which suggests that they are both reliably expressing the same
122 internal representation precision. We show that RT during value estimation for each item is also
123 strongly correlated with both measures of certainty, which suggests a link between the
124 representation precision and the cognitive effort required to decode the value signal. Finally, we

125 show that self-reported estimates of certainty generally increase across repeated value
126 estimations. This should be expected if contemplating the value of an option is tantamount to
127 constructing its internal representation, because as the cumulative total of processed
128 information rises (i.e., across multiple rating sessions), the precision of the representation should
129 also increase.

130 In a second study, which is essentially a replication of the rating-choice-rating paradigm of Lee
131 and Daunizeau (2020a, 2020b), we reproduce the previous findings that value estimate certainty
132 positively influences choice consistency and choice confidence, and negatively influences
133 response time and choice-induced preference change. We also show an interesting novel
134 result—we can predict which option a DM will choose based on the relative value estimate
135 certainty of the options being compared, even while ignoring the options' value estimates
136 themselves.

137 In sum, we suggest that researchers should no longer neglect the concept of value certainty when
138 building their models. We provide confirmatory and novel evidence that value certainty plays an
139 important role in the cognitive process of decision making. In particular, we suggest that of the
140 three different measures of value estimate certainty that we examined (self-reports, rating
141 consistency, rating RT), self-reported ratings are ideal, but that consistency could work well if
142 multiple ratings are available for each item. With respect to using rating RT as a proxy for value
143 certainty, we suggest that this could suffice if no other measures were available, but that caution
144 should be exercised when interpreting the results.

145 Note: for clarity, we explicitly use different terminology throughout this paper for subjective
146 beliefs about value estimates and about choices: "certainty" refers to the subjective feeling of
147 certainty about a value estimate rating; "confidence" refers to the subjective feeling of
148 confidence that the chosen item is the better one. We never use these terms interchangeably.

149 **METHODS**

150 We conducted a pair of behavioral experiments with the intention of demonstrating the
151 reliability of various measures of value certainty, which include: self-reports, rating consistency,
152 and response time. Participants considered a set of 200 options and provided three separate
153 value ratings for each option, as well as three separate certainty reports about those ratings. We
154 also recorded response time for each evaluation. Participants also made choices between pairs
155 of options, as well as choice confidence reports. Detailed task descriptions can be found below.

156 In Study 1, we asked participants to rate the value of each of a series of items. In addition to the
157 standard subjective value question, we also asked participants to rate their subjective certainty
158 regarding each subjective value judgment. Because we are interested in assessing the
159 consistency of value ratings and how that relates to subjective certainty ratings, we repeated the
160 value and certainty ratings three times during the experiment. This allows us to assess each of
161 our hypotheses, in particular: the correlation between value rating consistency and certainty
162 rating; the increase in certainty across repeated value ratings; the decrease in response time
163 across repeated value ratings.

164 In Study 2, we asked a different group of participants to rate the value of each of a series of items,
165 as well as to make choices between pairs of items. In addition to the standard subjective value
166 question, we also asked participants to rate their subjective certainty regarding each subjective
167 value judgment, and their subjective confidence regarding each choice.

168 *Materials*

169 We built our experiment in Gorilla (gorilla.sc). The experimental stimuli consisted of 200 digital
170 images, each representing a distinct snack item food item. The stimulus set included a wide
171 variety of items. Prior to commencing the experiment, participants received a written description
172 about the tasks and detailed instructions on how to perform them.

173 *Ethics statement*

174 Our analysis involved de-identified participant data and was approved by the ethics committee
175 of the University of Southern California (USC). In accordance with the Helsinki declaration, all
176 participants gave informed consent prior to commencing the experiment.

177 **Study 1**

178 *Participants*

179 A total of 37 people participated in this study (22 female; age: mean=21.4, stdev=3.6, min=18,
180 max=35). All participants were recruited from the undergraduate population at USC using
181 SonaSystems. Each participant received course credit as compensation for one hour of time.

182 *Experimental Design*

183 The experiment was divided into four sections. There was no time limit for the overall
184 experiment, nor for the different sections, nor for the individual trials.

185 In the first section (Exposure), participants merely observed as all individual items were displayed
186 in a random sequence for 750ms each. The purpose of the Exposure section was to familiarize
187 the participants with the full set of items that they would later evaluate, allowing them to form
188 an impression of the range of subjective value for the set. This would diminish the possibility
189 that value ratings would become more accurate across time merely due to a dynamic adaptation
190 of the range as the participants viewed new items.

191 The second through fourth sections (Rating1, Rating2, Rating3) were identical, except for the
192 sequence of trials. In each rating section, all stimuli were displayed on the screen, one at a time,
193 in a random sequence (randomized across participants and across sections for each participant).

194 At the onset of each trial, a fixation cross appeared at the center of the screen for 750ms. Next,
195 a solitary image of a food item appeared at the center of the screen. Participants responded to
196 the question, “How pleased would you be to eat this?” using a horizontal slider scale. The
197 leftmost end of the scale was labeled “Not at all.” and the rightmost end was labeled “Very
198 much!” The scale appeared to participants to be continuous, and the data was captured in
199 increments of 1 (ranging from 1 to 100). Participants could revise their rating as many times as
200 they liked before finalizing it. Participants clicked the “Enter” button to finalize their value rating
201 response and proceed to the next screen. Participants then responded to the question, “How
202 sure are you about that?” by clicking on a horizontal Likert scale (left to right: “not at all”,

203 “slightly”, “somewhat”, “fairly”, “very”, “extremely”) to indicate their level of subjective certainty
204 regarding the preceding value judgment. At that time, the next trial began. Participants were
205 not aware that there were three different rating sections, as the design technically only included
206 one rating section. Within that section, the sequence of trials included a random ordering of all
207 items, followed by another random ordering of all items, followed by another random ordering
208 of all items. From the perspective of the participant, the study consisted of a long series (600) of
209 item evaluations, where repetition could (and indeed, always did) occur. This design, in which
210 participants could not anticipate that each item was going to be rated multiple times, helped to
211 preclude participants from explicitly trying to remember and replicate their evaluations across
212 repetitions. A typical within-trial event sequence is shown in Figure 1.



213

214 **Figure 1:** an example of a within-trial even sequence for the rating tasks (study 1 and study 2)

215 **Study 2**

216 *Participants*

217 A total of 50 people participated in this study (18 female; age: mean=30.5, stdev=11.4, min=18,

218 max=64; 8 missing gender info; 2 missing age info). All participants were recruited from the online

219 subject pool using Prolific (prolific.co). Each participant received \$5 as compensation for 45
220 minutes of time.

221 *Experimental Design*

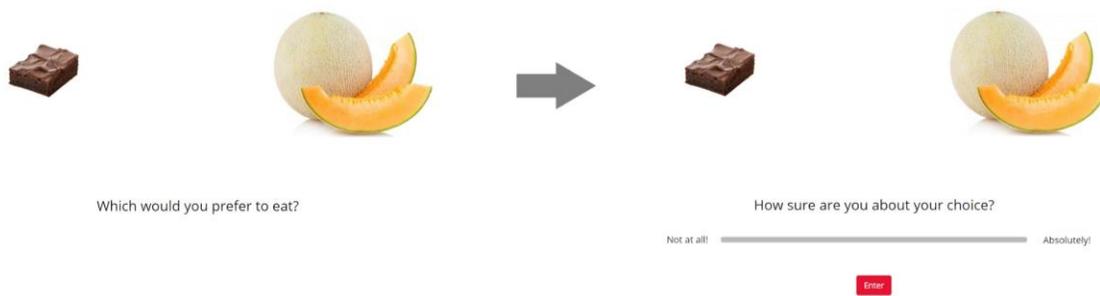
222 The experiment contained three different tasks: exposure, rating, and choice. There was no time
223 limit for the overall experiment, nor for the different tasks, nor for the individual trials.

224 The Exposure task was structurally identical to that described above for Study 1.

225 The Rating1 and Rating2 tasks were identical, except for the sequence of trials. For each rating
226 task, all of the stimuli to which the participant had initially been exposed were again displayed
227 on the screen, one at a time, in a random sequence (randomized across participants and across
228 sections for each participant). The structure of the rating trials was identical to that described
229 above for Study 1. Prior to commencing Rating2, participants were reminded that they should
230 report what they felt at that time and not try to remember what they reported during Rating1.
231 This helped to preclude participants from explicitly trying to remember and replicate their
232 evaluations across repetitions. A typical within-trial event sequence is shown in Figure 1 above.

233 For the Choice task, all stimuli were displayed on the screen, two at a time, in a random sequence
234 (randomized across participants and across sections for each participant). At the onset of each
235 trial, a fixation cross appeared at the center of the screen for 750ms. Next, a pair of images of
236 food items appeared on the screen, one towards the left, one towards the right. Participants
237 responded to the question, “Which would you prefer to eat?” by clicking on the image of their
238 preferred item. Participants then responded to the question, “Are you sure about your choice?”

239 using a horizontal slider scale. The leftmost end of the scale was labeled “Not at all!” and the
240 rightmost end was labeled “Absolutely!” Participants could revise their confidence report as
241 many times as they liked before finalizing it. Participants clicked the “Enter” button to finalize
242 their confidence report and proceed to the next screen. A typical within-trial event sequence is
243 shown in Figure 2.



244

245 **Figure 2:** an example of a within-trial even sequence for the choice task (study 2)

246

247 The pairings of items for each choice trial were created in a deliberate manner. Specifically, we
248 wanted to maximize the number of difficult choices that participants would be faced with. Here
249 we define difficulty as the similarity of the value ratings between choice pair items. Because our
250 simplified online experimental design did not allow for choice pairs to be created dynamically
251 based on each participant’s personal subjective value ratings, we relied on our data from Study
252 1. That data provided us with value ratings for 200 items across 37 participants, which we used
253 to calculate population statistics (median and variance of value estimate ratings) for each item.
254 We first calculated the population value estimate variability, which was the variance of the value
255 estimate ratings for each item across all 37 participants. Because we only wanted 150 items for

256 Study 2, we sorted the original 200 items from lowest to highest population variance and
257 removed the 50 highest-variability items (i.e., the items for which different participants had
258 provided the most variable value estimate ratings) from our set. We thought that this would
259 improve our chances that a new participant would rate the items similarly to the population
260 average ratings. Next, we calculated the population median value for each item. We used the
261 median instead of the mean so as not to be unduly influenced by extreme ratings. Sorting the
262 item set from highest to lowest value, we created triplets of items (i.e., [item1 item2 item3],
263 [item4 item5 item6], ...). We created 50 choice pairs for by selecting the first and second
264 elements from each triplet. We created an additional 50 choice pairs by selecting the first and
265 third elements from each triplet. We thus had a total of 100 choice pairs, all of which should be
266 difficult trials based on population statistics. (The reason why we created two separate sets in
267 the manner described was to allow us to pilot test a hypothesis for a future study, but is irrelevant
268 to this current study.) Obviously, individual ratings deviate from population ratings, which would
269 naturally cause many of the choice pair trials to be more or less difficult for individual
270 participants.

271 The usage of population median ratings in this way was solely to create choice pairs that would
272 *a priori* be likely to be difficult for most participants—it had no impact on the choice data itself,
273 which was based on the individual value estimate ratings of the participants who would actually
274 make the choices. Therefore, although choice pairs were created based on population value
275 estimate ratings from Study 1, the actual choice difficulty analyzed in the data for Study 2 was
276 determined entirely by the personal value ratings provided by each participant in that
277 study. Fortunately, this technique did indeed result in each participant in Study 2 facing a large

278 number of difficult choices (defined by their own personal value ratings). Note: the validity of
279 our analysis would not have been impacted either way, but the effects of interest would likely
280 have diminished.

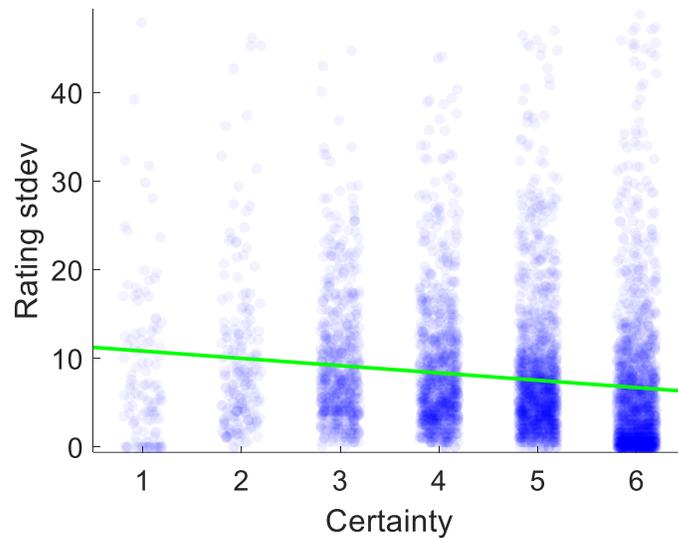
281 **RESULTS**

282 Before conducting our main analyses, we first validated that our data were reliable (see
283 Supplementary Material for details). We determined that the data were generally reliable,
284 although we decided to exclude 11 participants from Study 1 and six participants from Study 2,
285 for failing to perform the tasks properly for the duration of the experiment (see Supplementary
286 Material for details).

287 **Study 1**

288 *Hypothesis 1: Certainty should negatively correlate with rating inconsistency.*

289 Certainty reports were provided by the participants during the study, but we needed to obtain a
290 measure of rating inconsistency. For each participant, we thus calculated the within-item across-
291 section variance of value ratings (i.e., $V[\text{Rating1}_i \text{Rating2}_i \text{Rating3}_i]$ for $i=1:200$). We deemed
292 that variance is a measure of inconsistency, because perfect consistency would yield a variance
293 of zero and higher degrees of inconsistency would yield higher variance. For each participant,
294 we used the average certainty for each item across the three rating sections as our measure of
295 certainty. The correlation between certainty and inconsistency was negative and significant, as
296 expected (median Spearman's $\rho = -0.245$, $p < 0.001$). (See Figure 3.)



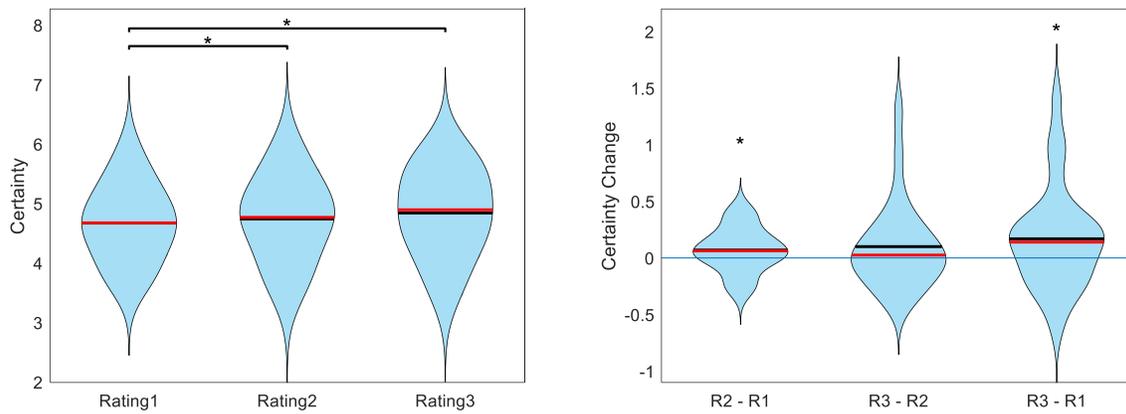
297

298 **Figure 3:** Scatter plot of the relationship between value estimate certainty and the standard deviation of ratings
299 across rounds 1-3, pooled across all participants. Each dot represents one trial.

300

301 *Hypothesis 2: Certainty should increase with repeated ratings.*

302 We first calculated the within-participant mean of certainty reports separately for Rating1,
303 Rating2, and Rating3. We then calculated the group averages for these values. The across-
304 participant across-item mean certainty for Rating1, Rating2, and Rating3 was 4.67, 4.74, and 4.84,
305 respectively (see Figure 4). The increase in average certainty between Rating1 and Rating2 and
306 between Rating1 and Rating3 were marginally significant ($p=0.086$, $p=0.068$; two-sided t-tests),
307 but the increase in certainty between Rating2 and Rating3 was not significant ($p=0.198$; two-
308 sided t-test).

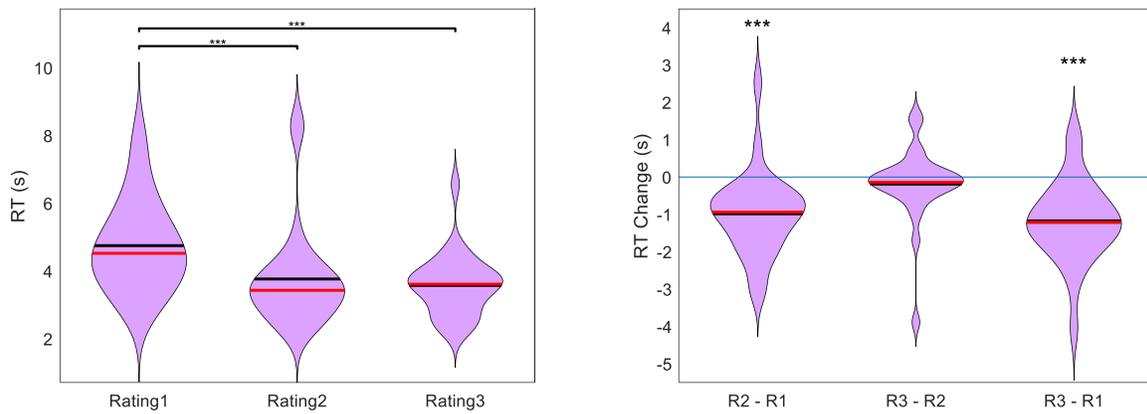


309

310 **Figure 4:** Across participants, value estimate certainty (across items) increased with each additional rating. Figure
311 shows the cross-participant mean of within-participant mean certainty ratings across trials, separate for each
312 rating task. (black lines indicate means, red lines indicate medians, significance stars indicate: * $p < 0.05$, ** $p < 0.01$,
313 *** $p < 0.001$)

314

315 In addition to the gradual increase in average certainty from Rating1 to Rating3, we also checked
316 to see if there was a gradual decrease in average response time (RT). Because online testing is
317 often plagued by distractions that cause some trials to have exceptionally long response times,
318 we first removed all outlier trials. We defined an outlier as any trial in which RT was greater than
319 the within-participant median RT plus three times the within-participant median average
320 deviation. This resulted in an average of 9 trials being removed per subject (out of 100). After
321 cleaning the data in this way, we indeed found that average RT decreased from one rating section
322 to the next. The across-participant across-item mean RT for Rating1, Rating2, and Rating3 was
323 4.75s, 3.77s, and 3.58s, respectively (see Figure 5). The decrease in RT between Rating1 and
324 Rating2 as well as between Rating1 and Rating3 was significant (both $p < 0.001$), but the decrease
325 between Rating2 and Rating3 was not ($p = 0.348$).



326

327 **Figure 5:** Across participants, RT (across items) decreased with each additional rating. Figure shows the cross-
328 participant mean of within-participant mean RT across trials, separate for each rating task. (black lines indicate
329 means, red lines indicate medians, significance stars indicate: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

330

331 We wondered if the decrease in RT over the course of the experiment (i.e., from Rating1 to
332 Rating2 to Rating3) could simply be due to a habituation to the task. If this were true, RT would
333 not only decrease across sections, but also within sections. We thus tested for a correlation
334 between RT and trial number, within each rating section for each participant. Across participants,
335 we found a mean correlation of -0.183 for Rating1 ($p < 0.001$) and -0.079 for Rating2 ($p = 0.001$),
336 but no reliable correlation for Rating3 (mean = -0.033, $p = 0.238$). In order to determine whether
337 the decrease in RT from Rating1 to Rating2 (reported above) was actually due to certainty gains
338 and not merely habituation, we split Rating1 into first half and second half trials. We then tested
339 for a change in RT from Rating1 to Rating2 separately for first and second half trials, for each
340 participant. There remained a significant decrease for both halves (first half mean RT change = -
341 1.20s, $p < 0.001$; second half mean RT change = -0.78s, $p < 0.001$).

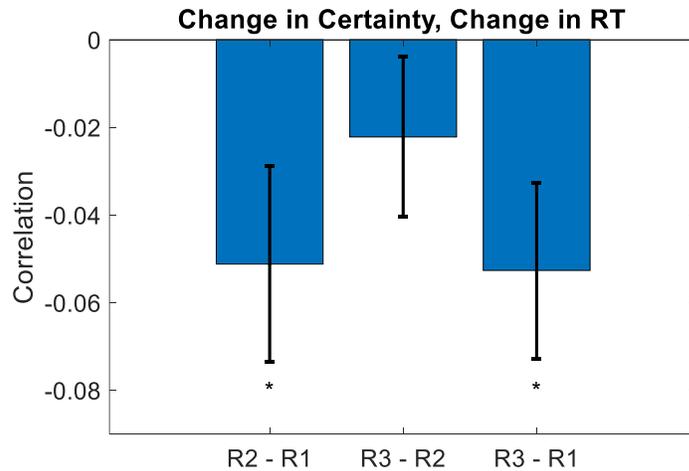
342

343 *Hypothesis 3: Certainty should negatively correlate with response time.*

344 For Rating1, the across-participant mean correlation between certainty and RT was negative, as
345 predicted (mean Spearman's $\rho = -0.104$, $p=0.004$). For Rating2, there was no statistically
346 significant correlation between certainty and RT (mean Spearman's $\rho = 0.016$, $p=0.539$). For
347 Rating3, there was actually a positive correlation between certainty and RT (mean Spearman's
348 $\rho = 0.082$, $p=0.009$).

349 Recalling that overall RT decreased across rating sections, we thought that this might have hidden
350 the inherent relationship between certainty and RT. The idea is as follows. Initially (i.e., for
351 Rating1), some items are evaluated with high certainty, others with low certainty. The high
352 certainty evaluations are reported faster than the low certainty evaluations, thereby establishing
353 the negative correlation between certainty and RT. Eventually (i.e., for Rating2 and Rating3),
354 low-certainty evaluations become more certain (and thus more quickly evaluated). But, high-
355 certainty evaluations remain certain, and there is not much room for an increase in certainty for
356 these evaluations. Therefore, when averaging across the entire set of items, this would cause an
357 overall increase in certainty as well as an overall decrease in RT. This could deteriorate the initial
358 relationship between certainty and RT, as the set of items in effect shifts towards similarity (i.e.,
359 high certainty and low RT). To test this idea, we examined the evolution of certainty and RT on
360 an item-by-item basis. We first calculated, for each participant and each item, the change in both
361 the certainty and the RT from Rating1 to Rating2, and then calculated the within-participant
362 correlation between those variables across items. Across participants, the correlation between
363 Certainty Change and RT Change from Rating1 to Rating2 was indeed negative (mean Spearman's

364 rho = -0.051, p=0.031). We then repeated this same analysis using the differences from Rating2
365 to Rating3, and from Rating1 to Rating3. These correlations were both negative as well, although
366 the former was not significant (mean Spearman's rho = -0.022, p=0.235; mean Spearman's rho =
367 -0.053, p=0.015; two-sided t-tests). (see Figure 6).



368

369 **Figure 6:** Across participants, the change in value rating certainty negatively correlated with the change in value
370 rating RT. The idea here is that on an item-by-item basis, as certainty increases (across repeated ratings), it takes
371 the DM less time to decide upon a rating estimate. (error bars represent s.e.m., significance stars represent: *
372 p<0.05, ** p<0.01, *** p<0.001)

373

374 **Study 2**

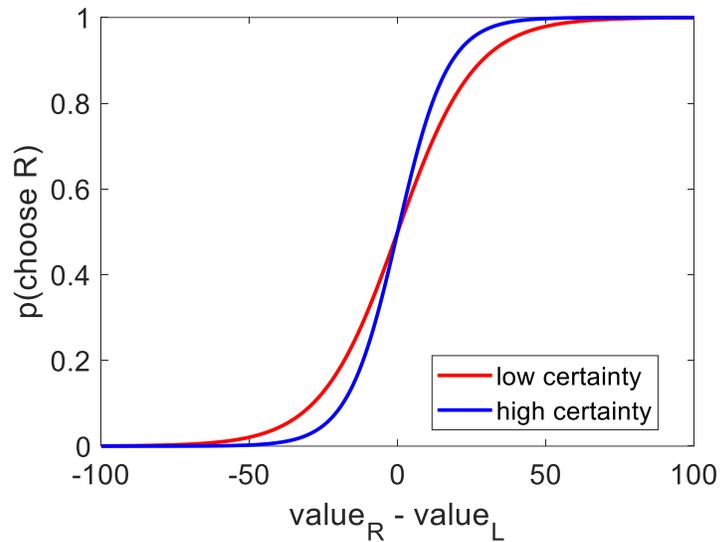
375 With Study 1, we demonstrated the reliability of our experimental measures of certainty
376 regarding subjective value estimates. With Study 2, we seek to further demonstrate the
377 importance of such measures by establishing their instrumental role in the decision making
378 process.

379

380 *Hypothesis 1: Choices will be more stochastic when value certainty is lower*

381 Value-based choice is primarily a function of the difference in the value estimates of the different
382 options in the choice set. The farther apart the value estimates are, the more likely it is that the
383 higher-rated item will be chosen; the closer together the value estimates are, the more likely it
384 is that the choice will appear to be random. Indeed, our data follow this pattern. For each
385 participant, we performed a logistic regression of choices against the difference in value ratings
386 of the paired options (choice = $\beta_0 + \beta_1 \cdot dV + \epsilon$). We found that this function fit the data
387 well above chance level, with a cross-participant average balanced accuracy of 77% ($p < 0.001$,
388 two-sided t-test). Across participants, there was no inherent bias for one side over the other
389 (mean $\beta_0 = -0.036$, $p = 0.350$) and there was a significant positive inverse temperature parameter
390 (mean $\beta_1 = 0.077$, $p < 0.001$).

391 What would be more interesting, however, would be to see how value estimate certainty
392 influences this choice model. We thus performed a similar logistic regression, for each
393 participant, except this time we also included an indicator variable that took the value of 1 if the
394 value certainty of a particular choice pair was greater than the median for that participant, and
395 0 otherwise (choice = $\beta_0 + \beta_1 \cdot dV + \beta_2 \cdot I \cdot dV + \epsilon$). Balanced accuracy remained at 77%
396 ($p < 0.001$). As with the previous model, there was no bias (mean $\beta_0 = -0.035$, $p = 0.356$) and the
397 inverse temperature for value difference was positive and significant (mean $\beta_1 = 0.077$,
398 $p < 0.001$). Notably, the regression coefficient for the interaction of value difference and the high
399 certainty indicator (i.e., the increase in choice precision between low and high value certainty
400 trials) was positive but only marginally significant (mean $\beta_2 = 0.042$, $p = 0.108$). (See Figure 7.)



401

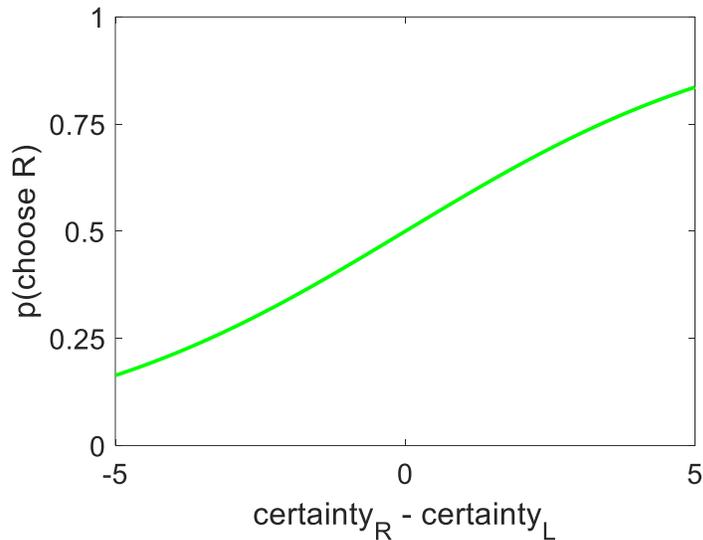
402 **Figure 7:** Across participants, the probability of choosing the option on the right increased as a function of the
403 value estimate difference (right option – left option). In particular: choices that were made between options with
404 low value certainty (red curve, within subject median split) were more stochastic than choices that were made
405 between options with high value certainty (blue curve) (left plot).

406

407 *Hypothesis 2: Options with higher value certainty will be chosen more often*

408 We posited that choices might be partially determined by how certain the individual value
409 estimates for each option were. We thus wondered how well choice could be predicted using
410 the difference in value certainty alone, without considering the difference in value estimates
411 themselves. For each participant, we ran a logistic regression of choices against the difference in
412 value estimate certainty (choice = $\beta_0 + \beta_1 \cdot dC + \epsilon$). Balanced accuracy was lower under
413 this model, as expected, but it was still well above chance level (cross-participant mean = 59%,
414 $p < 0.001$). Again, there was no bias (mean $\beta_0 = -0.044$, $p = 0.134$). The inverse temperature for
415 value certainty difference was positive and significant, as expected (mean $\beta_1 = 0.326$,
416 $p < 0.001$) (see Figure 8). This shows that choices can indeed be predicted by the difference in the

417 value certainty of the options under consideration, without directly examining the difference in
418 the value estimates themselves.



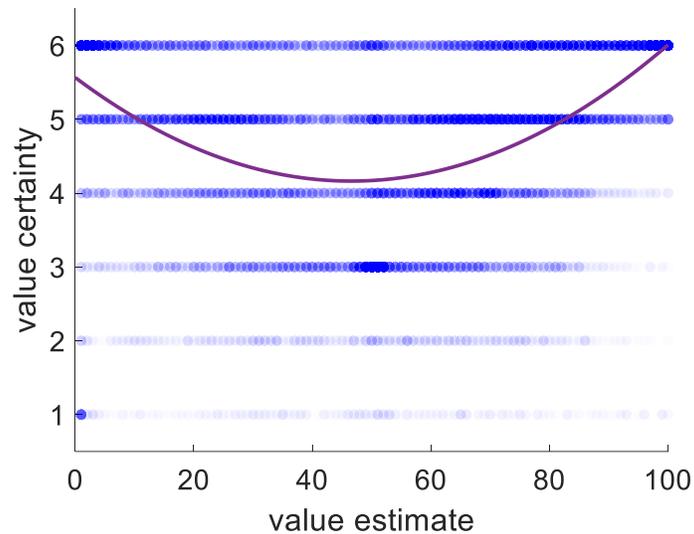
419

420 **Figure 8:** Across participants, the probability of choosing the option on the right increased as a function of the
421 value certainty difference (right option – left option).

422

423 Although we showed that choice could be predicted by value certainty even without considering
424 value estimate, we realized that there is generally a strong relationship between those two
425 variables. Supporting this notion, we found that value certainty correlated positively with value
426 estimate (mean Spearman's rho = 0.254, $p < 0.001$). Moreover, there was a clear u-shaped
427 relationship between value estimate and value certainty (see Figure 9). We note, however, that
428 the value certainty reports carried additional information beyond the value estimate ratings
429 themselves. The data clearly show that whereas very high or very low value estimates almost
430 always correspond to very high certainty, mid-range value estimates do not necessarily
431 correspond to relatively low certainty. It seems that sometimes participants estimated an item's

432 value to be mid-range because they were not certain about its true value, but other times they
433 were quite certain that its value was mid-range. This shows that value certainty is partially
434 constrained, but not fully determined, by value estimate itself.



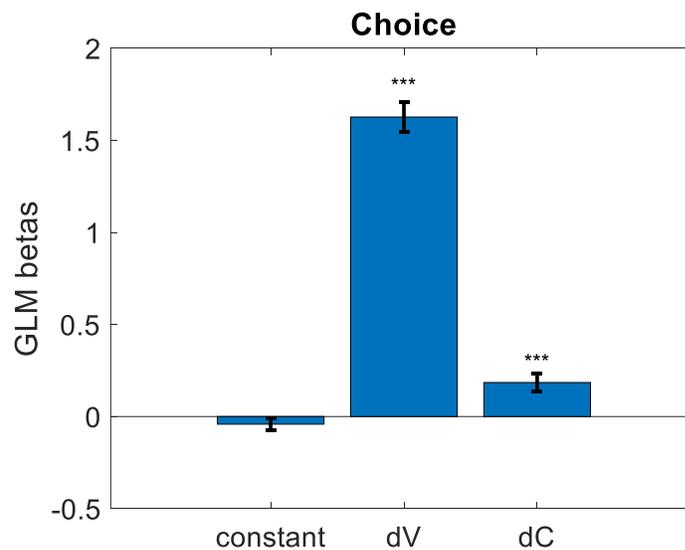
435

436 **Figure 9:** Value certainty is related to value estimate, with both a linear effect and a quadratic one. Blue dots
437 represent one item for one participant for one rating session. Purple curve represents the best linear + quadratic
438 fit across all trials and all participants.

439

440 Exploring further, we wondered if the predictive power of value certainty might be driven entirely
441 by its relationship with value estimate. That is, we wanted to check if the information contained
442 in the value certainty reports beyond what they convey about the value estimates themselves
443 would be useful in predicting choice. We predicted that, all else equal on a particular trial, the
444 option with the higher value estimate certainty would be the chosen option. To test this, we ran,
445 for each participant, a logistic regression of choices against the difference in value estimate
446 ratings as well as the difference in value certainty reports (choice = $\beta_0 + \beta_1 \cdot dV + \beta_2 \cdot dC$
447 + ϵ). Prior to running the regression, we first z-scored value estimate ratings and value certainty

448 reports separately for each participant. Balanced accuracy remained at 77% ($p < 0.001$). As with
449 the previous models, there was no bias (mean $\beta_0 = -0.042$, $p = 0.231$), and the inverse
450 temperature for value difference was positive and significant (mean $\beta_1 = 1.626$, $p < 0.001$).
451 Notably, the inverse temperature for certainty difference was also positive and significant (mean
452 $\beta_2 = 0.185$, $p < 0.001$). (See Figure 10.) The regression function we used orthogonalizes
453 regressors sequentially when calculating beta weights (i.e., here dC was orthogonalized to dV),
454 so the impact of dC was truly separate from the impact of dV. This suggests that not only did the
455 participants consider the difference in value estimates when choosing their preferred options,
456 but they also considered the difference in value certainty irrespective of the value estimates.



457

458 **Figure 10:** Cross-participant mean beta weights from GLM logistic regression of trial-by-trial value estimate
459 difference (dV) and value certainty difference (dC) onto choice. (error bars represent s.e.m.)

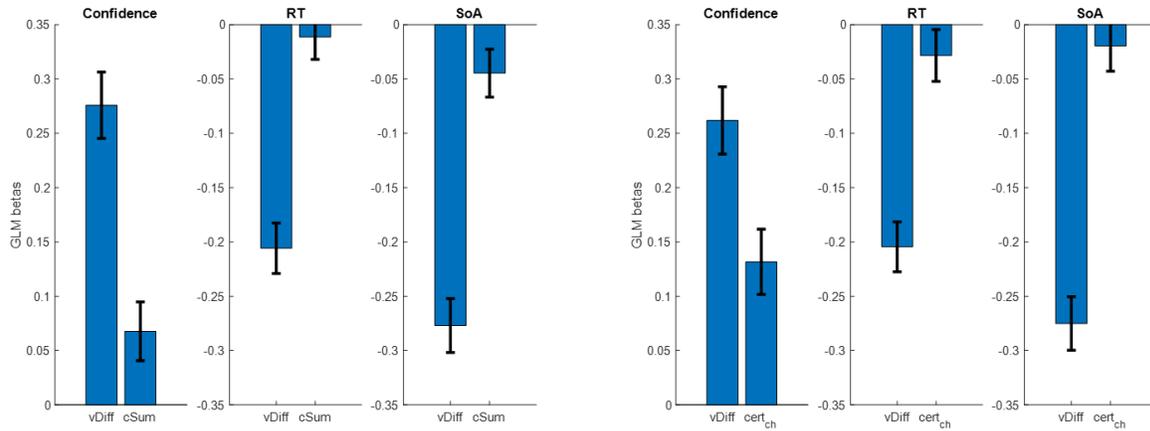
460

461

462

463 *Replication Results*

464 After testing our hypotheses, we next performed a series of analyses to try to replicate previously
465 reported results showing how value estimate certainty impacts a variety of dependent variables
466 during choice (Lee and Daunizeau, 2020a, 2020b). Specifically, we checked whether choice
467 confidence, response time, or choice-induced preference change changed as a function of value
468 estimate certainty. For our measure of choice-induced preference change, we used the
469 spreading of alternatives, defined as the post- minus pre-choice rating for the chosen option
470 minus the post- minus pre-choice rating for the rejected option. For each of the above dependent
471 variables, we ran a linear regression using absolute value estimate difference (vDiff) and summed
472 value estimate certainty (cSum) as regressors. We removed an average of five trials per
473 participant for having outlier RT ($RT > \text{median} + 3 * \text{MAD}$), and z-scored all variables within
474 participant. For choice confidence, we found that both independent variables had positive beta
475 weights, as predicted (mean for vDiff = 0.276, $p < 0.001$; mean for cSum = 0.068, $p = 0.016$; two-
476 sided t-tests). For response time, we found that both independent variables had negative beta
477 weights, as predicted, although only vDiff was significant (mean beta for vDist = -0.206, $p < 0.001$;
478 mean beta for cSum = -0.011, $p = 0.586$; two-sided t-tests). For spreading of alternatives, we
479 found that both independent variables had negative beta weights, as predicted (mean beta for
480 vDiff: -0.277, $p < 0.001$; mean beta for cSum: -0.045, $p = 0.048$; two-sided t-tests). (See Figure 11.)



481

482 **Figure 11:** Cross-participant mean beta weights from GLM regressions of trial-by-trial absolute value estimate
483 difference (vDiff) and summed value estimate certainty (cSum) onto choice confidence, choice response time, and
484 spreading of alternatives (left set of plots); same, using certainty of only the chosen option (right set of plots).
485 (error bars represent s.e.m.)

486

487 The use of cSum to represent the relevant aspect of value estimate certainty during choice
488 deliberation was somewhat arbitrary. We therefore decided to examine other measures in the
489 place of cSum, specifically: certainty of the chosen option (cert_{ch}), certainty of the rejected option
490 (cert_{rej}), difference in certainty between the chosen and rejected options (cert_{ch-rej}). For each
491 participant, we repeated the same GLM regression as described above, replacing cSum with
492 cert_{ch}, cert_{rej}, and cert_{ch-rej} in turn. We started with cert_{ch}. For choice confidence, we found that
493 both independent variables had positive beta weights, as predicted (mean for vDiff = 0.262,
494 p<0.001; mean for cert_{ch} = 0.132, p<0.001; two-sided t-tests). For response time, we found that
495 both independent variables had negative beta weights, as predicted, although only vDist was
496 significant (mean beta for vDist = -0.205, p<0.001; mean beta for cert_{ch} = -0.028, p=0.240; two-
497 sided t-tests). For spreading of alternatives, we found that both independent variables had
498 negative beta weights, as predicted, although only vDiff was significant (mean beta for vDiff: -

499 0.275, $p < 0.001$; mean beta for cert_{ch} : -0.020, $p = 0.408$; two-sided t-tests). (See Figure 10 above.)
500 The regression analyses using cert_{rej} did not yield significant beta weights for the certainty term
501 (confidence: mean beta for vDiff : 0.275, $p < 0.001$; mean beta for cert_{rej} : -0.032, $p = 0.180$; RT:
502 mean beta for vDiff : -0.205, $p < 0.001$; mean beta for cert_{rej} : 0.010, $p = 0.659$; SoA: mean beta for
503 vDiff : -0.281, $p < 0.001$; mean beta for cert_{rej} : -0.039, $p = 0.108$; two-sided t-tests). The regression
504 analyses using $\text{cert}_{\text{ch-rej}}$ yielded similar results as when using cert_{ch} , though with slightly lower
505 beta weights and slightly larger p-values (confidence: mean beta for vDiff : 0.263, $p < 0.001$; mean
506 beta for $\text{cert}_{\text{ch-rej}}$: 0.116, $p < 0.001$; RT: mean beta for vDiff : -0.204, $p < 0.001$; mean beta for $\text{cert}_{\text{ch-}}$
507 rej : -0.021, $p = 0.402$; SoA: mean beta for vDiff : -0.282, $p < 0.001$; mean beta for $\text{cert}_{\text{ch-rej}}$: 0.021,
508 $p = 0.460$; two-sided t-tests).

509

510 **DISCUSSION**

511 In this study, we have demonstrated the reliability of multiple measures of subjective value
512 estimate certainty, including self-reports, rating consistency, and response time. We have also
513 demonstrated the important role that value estimate certainty plays in choice itself, including its
514 positive impact on choice consistency and choice confidence, as well as its negative impact on
515 response time and choice-induced preference change. We might suggest that any contemporary
516 or future model of value-based decision making (and arguably, all types of decision making)
517 should consider including some measure of value estimate certainty for each of the options in
518 the choice set. At the present time, the only choice model that we are aware of that explicitly
519 includes a variable to represent value estimate certainty is the Metacognitive Control of Decisions

520 (MCD) presented by Lee and Daunizeau (2020a). This feature alone sets the MCD model apart
521 from the plethora of alternative models that abound in the literature. Yet it would not be
522 reasonable to claim that one class of model is inherently better than another simply because the
523 alternative failed to consider an important variable. Rather, we propose that the popular models
524 that already exist in the literature should be expanded to include value rating certainty. Only
525 then can a more complete and fair model comparison be made, and only then will we be able to
526 reach a better understanding of the cognitive mechanisms of decision making.

527 In particular, we call upon proponents of the so-called accumulation-to-bound models, such as
528 the Race Model (RM) and the Drift-Diffusion Model (DDM), to strongly consider revising their
529 models to include value estimate certainty. As it stands, most such models completely exclude
530 the possibility of item-specific certainty. These models typically (or exclusively) account for
531 stochasticity in the choice deliberation process at the system level, rather than at the option
532 level. This means that such models can explain or predict variations in observed behavior that
533 are dependent on choice context (e.g., clarity of perception, mental workload), but not on the
534 composition of the choice set itself. Given that stochasticity is one of the fundamental
535 components of evidence accumulation models (i.e., the diffusion parameter), it begs the question
536 as to why the nature of the stochasticity has not been more thoroughly explored. A related line
537 of work has indeed explored this question, concluding that uncertainty could spawn from noise
538 in sensory processing, stochasticity in response selection, or imperfections in probabilistic
539 inference (Drugowitsch et al, 2016). However, they did not discuss the possibility that choice
540 options might have different degrees of certainty intentionally represented in the brain.

541 Recent work has proven that an accumulation-to-bound process such as that represented by the
542 RM or DDM is an optimal policy, at least when optimality is defined as the maximization of reward
543 in a series of sequential decisions with a limited amount of time (Tajima et al, 2016, 2019). These
544 authors do indeed acknowledge the importance of certainty in their work, although it is not quite
545 of the same nature as that which we described in our study. In the work of Tajima et al (2016,
546 2019), pre-choice certainty about an option refers to the prior belief that a DM has about the
547 value distributions from which each option originates, rather than a belief about the value
548 estimates of the options themselves. However, we have shown that item-specific pre-choice
549 certainty is an important input to the choice deliberation process. Without a measure of item-
550 specific certainty, such a model cannot account for variations in choice behavior when the
551 different options originate from the same categorical set (e.g., snack foods). Tajima et al (2016)
552 suggest that evidence accumulation serves to increase the certainty about the option values, but
553 that the momentary evidence itself is uncertain. According to the authors, noise in the
554 momentary evidence itself could originate both externally (e.g., the stochastic nature of stimuli,
555 perceptual noise, ambiguity, incomplete knowledge) or internally (e.g., uncertain memory, value
556 inference that extends over time) (Tajima et al, 2016). Here, the authors seem to pave the way
557 for certainty measures that vary on an option-by-option basis, although they do not make this
558 explicit in their work.

559 Other recent work has suggested that the evidence accumulation process illustrated by a DDM is
560 influenced by attention (Sepulveda et al, 2020; Krajbich and Rangel, 2011; Krajbich et al, 2010).
561 Specifically, it has been proposed that during choice deliberation, evidence accumulates at a
562 higher rate for the option that is currently being gazed at, relative to the other option(s). This

563 evidence might support value estimation directly (Krajbich and Rangel, 2011; Krajbich et al, 2010)
564 or a more general goal-relevant information estimation (Sepulveda et al, 2020). However,
565 neither of these models include value estimate certainty. Indeed, these models explicitly assume
566 that both the prior uncertainty (i.e., variability in the environment from which the options
567 originate) and the evidence uncertainty (i.e., stochasticity in samples drawn from probability
568 distributions with fixed means) are identical across options. We have shown that such an
569 assumption is not reasonable, and thus likely impedes model performance.

570 Furthermore, these authors do not specifically examine the role of value certainty in gaze
571 allocation. If gaze fixation is what focuses information processing on each option in turn, it could
572 be that a DM will be inclined to gaze more at options whose values are less certain. Similar to
573 the concept of exploration in the classic exploration/exploitation dilemma (e.g., Daw et al, 2006),
574 this gaze bias could be instrumental for the DM to make more informed choices (Callaway et al,
575 2020). Or, it could be that a DM will be inclined to gaze more at options whose values are *more*
576 certain. Value certainty could be a measure of how reliable the information about the option is,
577 which could bias gaze towards options with higher certainty. Further studies will be required to
578 demonstrate the direction of the influence of value certainty on gaze patterns, but it is likely that
579 it plays some important role in the decision process. Regardless of the direction, we hold that
580 gaze duration should correlate with an increase in value estimate certainty.

581

582 **ACKNOWLEDGEMENTS**

583 This manuscript has been released as a pre-print at

584 <https://www.biorxiv.org/content/10.1101/2020.06.16.155234v1>.

585

586 **BIBLIOGRAPHY**

587

588 Callaway, F., Rangel, A., & Griffiths, T. L. (2020). Fixation patterns in simple choice are consistent
589 with optimal use of cognitive resources. PsyArXiv preprint PsyArXiv: <https://doi.org/10.31234/osf.io/57v6k>.

591 Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for
592 exploratory decisions in humans. *Nature*, 441(7095), 876-879.

593 De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based
594 choice. *Nature neuroscience*, 16(1), 105.

595 Drugowitsch, J., Wyart, V., Devauchelle, A. D., & Koechlin, E. (2016). Computational precision of
596 mental inference as critical source of human choice suboptimality. *Neuron*, 92(6), 1398-1411.

597 Frydman, C., & Jin, L. J. (2019). Efficient coding and risky choice. Available at SSRN 3270773.

598 Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F., & Behrens, T. E. (2012).
599 Mechanisms underlying cortical activity during value-guided choice. *Nature neuroscience*,
600 15(3), 470-476.

601 Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and
602 behavioural impact of decision confidence. *Nature*, 455(7210), 227-231.

603 Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by
604 neurons in the parietal cortex. *science*, 324(5928), 759-764.

605 Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship
606 between visual fixations and choice in value-based decisions. *Proceedings of the National
607 Academy of Sciences*, 108(33), 13852-13857.

608 Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison
609 of value in simple choice. *Nature neuroscience*, 13(10), 1292.

610 Lee, D., & Coricelli, G. (2020). An empirical test of the role of value certainty in decision making.
611 *bioRxiv*, 155234.

612 Lee, D., & Daunizeau, J. (2020a). Trading Mental Effort for Confidence: The Metacognitive Control
613 of Value-Based Decision-Making. *bioRxiv*, 837054.

614 Lee, D., & Daunizeau, J. (2020b). Choosing what we like vs liking what we choose: How choice-
615 induced preference change might actually be instrumental to decision-making. *PLoS One*,
616 15(5), e0231081.

617 Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism
618 for context-dependent decision making. *Proceedings of the National Academy of Sciences*,
619 110(15), 6139-6144.

620 Padoa-Schioppa, C., & Rustichini, A. (2014). Rational attention and adaptive coding: a puzzle and
621 a solution. *American Economic Review*, 104(5), 507-13.

622 Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and
623 accuracy of a perceptual decision. *Journal of vision*, 5(5), 1-1.

- 624 Polania, R., Woodford, M., & Ruff, C. C. (2019). Efficient coding of subjective value. *Nature*
625 *neuroscience*, 22(1), 134-142.
- 626 Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic
627 quantities for different goals. *Nature neuroscience*, 19(3), 366.
- 628 Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of
629 value-based decision making. *Nature reviews neuroscience*, 9(7), 545-556.
- 630 Sepulveda, P., Usher, M., Davies, N., Benson, A., Ortoleva, P., & De Martino, B. (2020). Visual
631 attention modulates the integration of goal-relevant evidence and not value. *bioRxiv*.
- 632 Tajima, S., Drugowitsch, J., Patel, N., & Pouget, A. (2019). Optimal policy for multi-alternative
633 decisions. *Nature neuroscience*, 22(9), 1503-1511.
- 634 Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making.
635 *Nature communications*, 7(1), 1-12.
- 636 Woodford, M. (2019). Modeling imprecision in perception, valuation and choice (No. w26258).
637 National Bureau of Economic Research.
638

639 **Supplementary Material**

640 *Data Quality Check*

641 **Study 1**

642 Before testing our hypotheses, we performed a number of simple data quality checks. First, we
643 assessed the test-retest reliability of value ratings. For each participant, we thus measured the
644 correlation between first rating (Rating1) and second rating (Rating2), across items. We found
645 that ratings were generally consistent (median Spearman's rho = 0.817). Most participants
646 showed a correlation of greater than 60%. We then measured, for each participant, the
647 correlation between Rating1 and Rating3. We found that ratings were generally consistent
648 (median Spearman's rho 0.818).

649 Next, we performed a similar assessment of the test-retest reliability of certainty reports. Before
650 examining the certainty data, we first converted the qualitative reports to numbers ("not at all"
651 = 1, "slightly" = 2, "somewhat" = 3, "fairly" = 4, "very" = 5, "extremely" = 6). For each participant
652 we then measured the correlation between certainty for Rating1 (Certainty1) and certainty for
653 Rating2 (Certainty2), across items. We found that certainty reports were generally consistent
654 (median Spearman's rho = 0.352), although much less so than value ratings. We then measured,
655 for each participant, the correlation between Certainty1 and Certainty3, across items. We found
656 that certainty reports were generally consistent (median Spearman's rho = 0.353), although
657 much less so than value ratings.

658 Because rating certainty is a key variable for testing our hypotheses, we needed to be sure that
659 participants responded meaningfully to the rating certainty question. An analysis of how
660 certainty correlates with our other variables of interest would not be possible where there is
661 insufficient variability in the certainty data. For this reason, we calculated a score for how much
662 variance each participant had across certainty reports. The median certainty report variance was
663 0.935, 0.882, and 0.599 for Rating1, Rating2, and Rating3, respectively. We deemed that there
664 were no participants who were obvious outliers based on this score.

665 For each of the test-retest reliability measures described above, we searched for population
666 outliers. We defined an outlier for a specific measure as a participant whose score was more
667 than three median average deviations (MAD) away from the population median. This technique
668 yielded three outlier participants based on the Rating1-Rating2 test-retest reliability scores, and
669 eight outlier participants based on the Rating1-Rating3 test-retest reliability scores. There was
670 one outlier with respect to Certainty1-Certainty2, and six with respect to Certainty1-Certainty3.
671 Seven of the outlier participants were caught by more than one filter, which left us with a set of
672 11 total outlier participants for Study 1. We excluded these participants from our reported
673 analyses.

674 **Study 2**

675 Before testing our hypotheses, we performed a number of simple data quality checks. First, we
676 assessed the test-retest reliability of value ratings. For each participant, we thus measured the
677 pairwise linear correlation between first rating (Rating1) and second rating (Rating2), across

678 items. We found that ratings were generally consistent (median Spearman's $\rho = 0.803$,
679 $p < 0.001$).

680 Next, we performed a similar assessment of the test-retest reliability of certainty reports. Before
681 examining the certainty data, we first converted the qualitative reports to numbers ("not at all"
682 = 1, "slightly" = 2, "somewhat" = 3, "fairly" = 4, "very" = 5, "extremely" = 6). For each participant
683 we then measured the pairwise linear correlation between certainty for Rating1 (Certainty1) and
684 certainty for Rating2 (Certainty2), across items. We found that certainty reports were generally
685 consistent (median Spearman's $\rho = 0.344$, $p < 0.001$), although much less so than value ratings.

686 Because rating certainty is a key variable for testing our hypotheses, we needed to be sure that
687 participants responded meaningfully to the rating certainty question. An analysis of how
688 certainty correlates with our other variables of interest would not be possible where there is
689 insufficient variability in the certainty data. For this reason, we calculated a score for how much
690 variance each participant had across certainty reports. The median certainty report variance was
691 0.971 and 0.818 for Rating1 and Rating2, respectively. We deemed that there were no
692 participants who were obvious outliers based on this score.

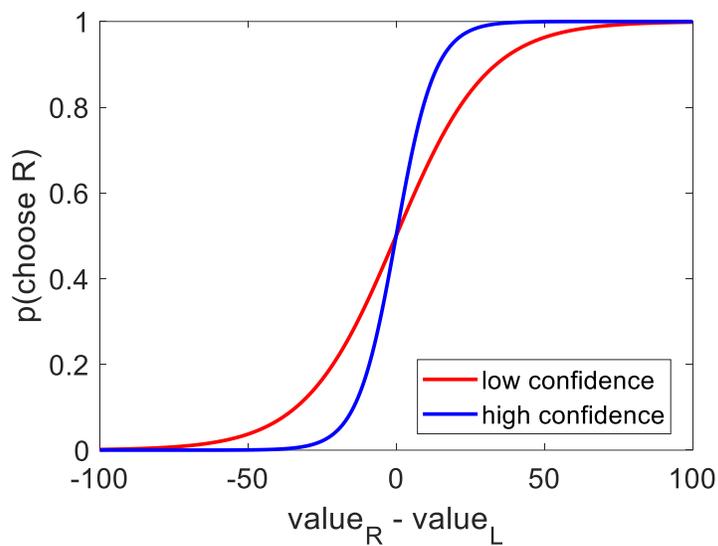
693 Finally, we checked whether choices were consistent with pre-choice ratings. For each
694 participant, we performed a logistic regression of choices against the difference in value ratings
695 of the paired options. We found that the balanced prediction accuracy was beyond chance level
696 (mean 77%), indicating participants were performing the choice task properly.

697 For each of the test-retest reliability measures described above, we searched for population
698 outliers. We defined an outlier for a specific measure as a participant whose score was more
699 than three median average deviations (MAD) away from the population median. This technique
700 yielded three outlier participants based on the Rating1-Rating2 test-retest reliability scores. We
701 excluded these participants from our reported analyses.

702 *Effect of confidence on choice consistency*

703 We then explored a step further, postulating that choice confidence should modulate choice
704 consistency (often referred to as accuracy). The idea is that for high confidence choices, the DM
705 would more consistently distinguish the items, relative to low confidence choices. We thus
706 performed a similar logistic regression as we did in our main analysis (see Figure 6 in Results), for
707 each participant, except this time the indicator represented high choice confidence (within-
708 participant median split) instead of value certainty (choice = $\text{logistic}[\beta_0 + \beta_1 \cdot dV +$
709 $\beta_2 \cdot \text{Ind} \cdot dV]$). Under this model, balanced accuracy was also 77% ($p < 0.001$). Again, there was
710 no bias (mean $\beta_0 = -0.028$, $p = 0.466$), and the inverse temperature parameter remained
711 positive and significant (mean $\beta_1 = 0.065$, $p < 0.001$). Notably, the regression coefficient for the
712 interaction of value difference and the high confidence indicator (i.e., the increase in choice
713 precision between low and high confidence trials) was positive and significant (mean $\beta_2 =$
714 0.088 , $p < 0.001$) (see Figure S1). We thus confirmed a common observation that choice
715 confidence and choice accuracy are closely linked.

716



717

718 **Figure S1:** Across participants, the probability of choosing the option on the right increased as a function of the
719 value estimate difference (right option – left option). In particular: choices that were made with low confidence
720 (red curve, within subject median split) were more stochastic than choices that were made with high confidence
721 (blue curve).

722