The Anchor Integration Model: A Descriptive Model of Anchoring Effects

Brandon M. Turner^{a,*}, Dan R. Schley^b

^a*Department of Psychology, The Ohio State University*

^b*Rotterdam School of Management, Erasmus University*

Abstract

Few experimental effects in the psychology of judgment and decision making have been studied as meticulously as the anchoring effect. Although the existing literature provides considerable insight into the psychological processes underlying anchoring effects, extant theories up to this point have only generated qualitative predictions. While these theories have been productive in advancing our understanding of the underlying anchoring process, they leave much to be desired in the interpretation of specific anchoring effects. In this article, we introduce the Anchor Integration Model (AIM) as a descriptive tool for the measurement and quantification of anchoring effects. We develop two versions the model: one suitable for assessing between-participant anchoring effects, and another for assessing individual differences in anchoring effects. We then fit each model to data from two experiments, and demonstrate the model's utility in describing anchoring effects.

Keywords: Anchoring, Judgments, Modeling, Estimation

*Corresponding Author

Email address: `turner.826@gmail.com` (Brandon M. Turner)

This research was supported by a National Science Foundation Graduate Research Fellowship to the second author, and a National Science Foundation Research Grant SES-1047757 to Ellen Peters. We thank Ellen Peters, Michael DeKay, Gretchen Chapman, and Duane Wegener for helpful comments on earlier versions of this manuscript.

Preprint submitted to Cognitive Psychology

December 15, 2015

1. Introduction

As technology and innovation have brought a new age of intellectual prosperity, we find ourselves with a multitude of quantitative information. For years, psychologists have tried to understand how incidental numbers influence judgments. Anchoring effects were brought to the forefront of psychological research by Kahneman and Tversky (1974), although earlier examples of a similar phenomenon were demonstrated in previous research on preference reversals (Lichtenstein and Slovic, 1971; Slovic, 1967; Slovic and Lichtenstein, 1968). In their original article, Kahneman and Tversky first asked participants to indicate whether the percentage of African nations in the United Nations was greater than or less than an arbitrary number, derived from spinning a wheel of fortune. This arbitrary number, the anchor, was set to be either high (e.g., 65%) or low (e.g., 10%). Participants were then asked to give their best estimate of the percentage of African nations in the United Nations. Results indicated that participants' estimates assimilated to the provided anchor value, such that the mean estimate of participants in the high-anchor condition was 45% and the mean estimate of participants in the low-anchor condition was 25%.

Over the past forty years, anchoring research has received extensive attention because of its robust effect sizes and its broad applicability in a myriad of domains. One such example was provided by Northcraft and Neale (1987), in which students and real estate appraisers toured a home and appraised its value. Although each participant had relevant experiential information about the home, both groups showed a significant correlation between their appraisal and the anchor value. In another example, Plous (1989) demonstrated that students' estimates of the likelihood of nuclear war were significantly influenced by arbitrary anchors. Other common paradigms include the pricing of gambles (Carlson, 1990; Johnson and Schkade, 1989; Chapman and Johnson, 1994, 1999), self efficacy (Cervone and Peake, 1986), negotiations (Galinsky and Mussweiler, 2001), judicial verdicts (Englich and Mussweiler, 2001; Englich et al., 2006; Chapman and Bornstein, 1996), consumer decisions and willingness-to-pay (Ariely et al., 2003; Wansink et al., 1998; Stewart, 2009; Simonson and Drolet, 2004; Green et al., 1998), and general knowledge questions (Jacowitz and Kahneman, 1995; Strack and Mussweiler, 1997; McElroy and Dowd, 2007).

1.1. The Problem with the Anchoring Effect

Although it is clear that anchoring effects have widespread implications, less is known about the cognitive mechanisms that drive these effects. Over the past few decades, theories of the anchoring-effect mechanism have risen, but unfortunately, extant theories have not fallen. Currently, there are five major theories – reviewed below – some with additional minor offshoots, that provide insight about how and why individuals assimilate judgments to presented anchor values. These theories provide psychological mechanisms that explain the processes underlying the anchoring effect and potentially relevant boundary conditions for the effect (for a review, see Furnham and Boo, 2011). Nevertheless, after 4 decades of research, what *exactly* do we know about the anchoring effect and how to predict it?

Anchoring effects often involve the presentation of a quantitative anchor, from which participants produce a quantitative response (although there are exceptions, e.g., Oppenheimer et al., 2008). Because the anchoring effect refers to the assimilation of a quantitative judgment to a quantitative anchor, qualitative theories are incapable of fully articulating the effect. In particular, extant theories provide considerable insight into the psychological processes that ostensibly produce the assimilation of individuals’ judgments toward a presented anchor value, though they offer little insight into the magnitude of said assimilation. For example, consider a question in which participants are asked to estimate the length of the Mississippi river. Current theories provide qualitative predictions about how the presence of a high or low anchor will affect the shift of the judgments toward the anchor value, but they do not describe the distribution of these judgments. Hence, current theories of anchoring effects are incapable of making quantitative predictions for say, an anchor value of 3,500 miles versus an anchor value of 1,500 miles.

The extensive anchoring literature provides considerable evidence for the robustness of the anchoring effect across substantive domains. If the anchoring effect can reliably influence important decisions like home prices (Northcraft and Neale, 1987) and judicial decisions (Englich and Mussweiler, 2001; Englich et al., 2006; Chapman and Bornstein, 1996), then the ability to make quantitative predictions (i.e., the amount of assimilation for any given anchor) would be extremely useful. In this article, we attempt to provide a long overdue tool for the study of anchoring effects by introducing a descriptive quantitative model.

Considerable attention has recently focused on the contribution of behavioral research as a tool to allow policy makers to “nudge” individuals’ judg-

ments and decisions in a manner that facilitates socially desirable outcomes (Allcott and Mullainathan, 2010; Johnson and Goldstein, 2003; Thaler and Sunstein, 2008). Take for instance the proportion of income that American consumers save. The savings rate has large macroeconomic consequences, with deleterious outcomes associated with both savings rates that are too low (i.e., the aggregate costs of interest decreases consumers' subsequent buying power, and as a consequence, gross domestic product) and too high (i.e., lower rates of consumer consumption generally decrease gross domestic product). Here, policy recommendations intended to nudge consumer spending would act as anchors (e.g., telling individuals that the ideal savings rate is 7% for 25-40 year-olds). Extant anchoring research offers intuition about what types of anchors produce relatively larger and smaller anchoring effects, but offer little insight into which anchors might produce the most optimal outcomes. Because there are more and less optimal saving rates, tools that quantify important behavioral effects offer policy makers the ability to potentially introduce optimal nudges (Urminsky and Goswami, 2015).

In subsequent sections we will begin by reviewing extant theories as a means of guiding the development of our model. Next, we present two versions of the Anchor Integration Model (AIM) for the prediction of quantitative judgments. We will first present a version of AIM developed to model anchoring effects on a between-participant level (e.g., providing a group of participants a particular anchor value and estimating the distribution of corresponding judgments). Following, we will present a within-participant version of AIM meant to characterize how an individual anchor influences an individual's judgment. As a first endeavor into quantitatively evaluating anchoring effects, the goal of the present article is to develop a descriptive model of the data, rather than a mechanistic model of the underlying psychological process.

2. Current Theories of Anchoring and a General Framework for AIM

Although AIM is not intended to be a process model, reviewing current theories of the psychological mechanism at play will be useful in guiding the development of our descriptive model.

107 2.1. *Anchoring-and-Adjustment Theory*

108 The first theory in this line of research – Anchoring-and-adjustment the-
109 ory – assumes that participants anchor their judgments on the provided an-
110 chor value and insufficiently adjust their estimates away from the anchor
111 (Kahneman and Tversky, 1974). Consider the length of the Mississippi River.
112 According to Anchoring-and-adjustment, individuals will begin at a given an-
113 chor value (e.g., 3,500 miles) and make a series of incremental adjustments,
114 pausing after each adjustment to decide whether or not to continue adjusting.
115 The anchoring effect is thought to occur because individuals insufficiently ad-
116 just away from the anchors, resulting in estimates that are skewed toward
117 the anchor value (Epley and Gilovich, 2006).

118 Within the literature, two primary accounts are used to explain this insuf-
119 ficient adjustment. The first account argues that individuals have a plausible
120 set of values for all judgments and that estimates are derived from the first
121 value within their plausible interval (Quattrone, 1982; Quattrone et al., 1984;
122 Epley and Gilovich, 2006). The second account of Anchoring-and-adjustment
123 focuses on the attention-demanding nature of adjustment. According to this
124 account, adjustment is an elaborative process in which an individual must
125 consider each adjustment and choose whether to continue elaborating (i.e.,
126 adjusting to a new value) or settle on the current value (Gilbert, 2002; Ep-
127 ley, 2004; Epley and Gilovich, 2006). Because attention is a limited resource,
128 insufficient adjustment is a product of the cumulative costs of adjusting.

129 2.2. *Numeric and Magnitude Priming*

130 Initially set forth by Wilson et al. (1996), the Numeric-priming account
131 suggests that *any* number present in the judgment environment can act as an
132 anchor, independent of context or relevance. For example, participants in one
133 study expected that the price of a meal would be higher at a restaurant named
134 “Studio 97” than at one named “Studio 17” (Critcher and Gilovich, 2008). In
135 addition, Oppenheimer et al. (2008) demonstrated that simply drawing a long
136 line before producing an estimate resulted in reliably larger estimates than
137 those produced after drawing a short line. That is, drawing a line produces
138 a sense of “bigness” or “smallness” which primes magnitude information in
139 a subsequent judgment. The primary proposition of this account is that
140 any magnitude information available during the focal judgment is conflated
141 with the judgment-relevant information. For example, if an individual sees
142 the number 3,500 before being asked to judge the length of the Mississippi

143 River, the “3,500” will be added to judgment-relevant information, regardless
144 of its actual informativeness.

145 2.3. *The Selective Accessibility Model*

146 The Selective-accessibility model postulates that anchoring effects are me-
147 diated by a selective increase in the accessibility of anchor-consistent knowl-
148 edge about the target item (Strack and Mussweiler, 1997; Mussweiler and
149 Strack, 1999; Chapman and Johnson, 1999, 1994, 2002; Mussweiler et al.,
150 2000). Specifically, Selective Accessibility supposes that when making the
151 comparative judgment in an anchoring task (e.g., is the Mississippi River
152 greater than or less than 3,500 miles?) individuals test the hypothesis that
153 the anchor is equal to the true value of the target item. To perform this con-
154 firmatory hypothesis test, individuals *selectively* retrieve information that is
155 consistent with the Mississippi River being 3,500 miles long (see Klayman
156 and Ha, 1991; Wason, 1960). When the estimate is subsequently elicited, this
157 retrieved information is more *accessible*, biasing estimates toward the anchor
158 value. In support of this model, Mussweiler and Strack (1999, 2000, 2001)
159 conducted a series of studies measuring the increase in available informa-
160 tion around the anchor value. In one study, participants who had indicated
161 whether the average price for a German automobile was greater than or
162 less than the high anchor of 40,000 German Marks, were faster at identify-
163 ing words associated with expensive cars (e.g., Mercedes-Benz) than words
164 associated with cheaper cars (e.g., Volkswagen) in a lexical decision task
165 (Mussweiler and Strack, 2000). Thus, information around the anchor value
166 is activated in memory in a manner consistent with Selective Accessibility.

167 2.4. *The Attitudinal Perspective of Anchoring*

168 A recent perspective of numeric anchoring has arisen from the attitudes
169 and persuasion literature. The Attitudinal Perspective contends that indi-
170 viduals process numeric anchors and persuasive messages in similar ways
171 (Wegener et al., 2010). In particular, the Attitudinal Perspective argues
172 that different types of anchoring can occur between thoughtful and non-
173 thoughtful processes, in a manner consistent with the elaboration likelihood
174 model (Petty and Cacioppo, 1986). An efficient summary was provided by
175 Wegener et al. (2010):

176 “Effortful use of anchor-consistent background knowledge pro-
177 vides a reasonable account for high-elaboration anchoring, but

178 there may be a number of relatively non-thoughtful processes
179 that contribute to low-elaboration anchoring. Numeric (magni-
180 tude) priming is one potential process, and it also seems plausible
181 that people take at least some anchors as ‘hints’ as to the correct
182 (direction of) target judgment.” (Wegener et al., 2010, p. 11).

183 Simply put, when individuals are elaborating on the judgment, they pro-
184 cess anchors in a manner consistent with Selective Accessibility, but when
185 they are not elaborating, they process anchors in a manner consistent with
186 the Numeric-priming account. One additional advantage of the Attitudinal
187 Perspective is that it allows for anchors to serve as “hints” (Schwarz, 1994;
188 Wegener et al., 2001, 2010). In particular, participants may assume that
189 the presented anchor provides implicit information regarding the judgment
190 (Grice, 1975; Jacowitz and Kahneman, 1995; Schwarz, 1994); a feature not
191 available in prior accounts. Consistent with this notion, Zhang and Schwarz
192 (2013) found that anchors with precise numbers (e.g., 3,478) resulted in larger
193 anchoring effects than comparable round anchors (e.g., 3,500) because indi-
194 viduals use the precision of the presented anchor as a means of assessing the
195 quality of the anchor information. As noted by Wegener et al. (2010), an
196 anchor from a credible source should result in a larger anchoring effect than
197 an anchor from a less credible source.

198 2.5. *The Scale Distortion Theory of Anchoring*

199 The most recent addition to the growing collection of anchoring theories is
200 the Scale-distortion theory of anchoring proposed by Frederick and Mochon
201 (2012). According to this account, individuals map judgments to an underly-
202 ing response scale. The presentation of an anchor does not shift individuals’
203 subjective representation of the target item (Mussweiler, 2003), but informs
204 the individual about the scale used to make the judgment. For example,
205 when presented a low anchor for the length of the Mississippi River (e.g.,
206 70 miles), individuals typically produce lower estimates for the river’s length
207 (e.g., 1,000 miles), compared to a control condition without an anchor (e.g.,
208 2,000 miles). Scale Distortion suggests that individuals do not perceive the
209 Mississippi River as any shorter, rather the scale used to make the judgment
210 distorts around the anchor value (Frederick and Mochon, 2012; Mochon and
211 Frederick, 2013). When judging the river’s length, given the presence of a
212 70 mile anchor, 1,000 miles seems sufficient to convey that the Mississippi
213 River is very long, but without the 70 mile anchor to serve as a cue, 2,000

214 miles seems necessary to convey the “largeness” of the length of the river.
 215 For example, Frederick and Mochon (2012) presented participants with a list
 216 of 15 animals ordered by weight. Participants were tasked with choosing
 217 an animal whose average adult weight was closest to 1,000 pounds. Before
 218 doing so, half of the participants estimated the average weight of an adult
 219 wolf, the other half did not. Results indicated that participants who first
 220 judged the weight of a wolf selected examples of a 1,000-pound animal that
 221 were substantially larger than did those who did not first judge the weight of
 222 the wolf. According to the authors, because wolves are much less than 1,000
 223 pounds, if anchors (i.e., judging the weight of a wolf) distort the scale of
 224 animal weights around the the anchor, then a much larger animal would be
 225 chosen as an exemplar of a 1,000-pound animal. Simply put, animals smaller
 226 than the anchor would assimilate to values higher along the scale and animals
 227 larger than the anchor would assimilate to values lower along the scale.

228 *2.6. Integrating Across Theories*

229 To the reader, the existence of five non-mutually-exclusive theories may
 230 cause concern about the efficacy of these theories or the anchoring litera-
 231 ture more generally. We suggest that the persistence of this ever increasing
 232 number of theories is due to a misnomer within this literature. This assim-
 233 ilation of judgments to presented anchor values is commonly referred to as
 234 the “anchoring effect,” despite a range of judgment domains and types of
 235 anchors. We believe that it would be helpful to conceptualize this not as “an
 236 effect” but as a class of effects with a common antecedent (i.e., the presence
 237 of an anchor) and a common consequence (i.e., the assimilation of judgments
 238 toward the anchor value). To facilitate clarity, we would suggest referring
 239 to “anchoring effects” as opposed to “the anchoring effect.” If the body of
 240 experiments in this literature represent a class of effects, and the existing
 241 theories provide insight into the domain-general processes across paradigms,
 242 then our model may be informed by considering commonalities across the
 243 range of theories and anchoring effects.

244 The theories reviewed above differ considerably with regard to how the
 245 anchor influences judgments : According to Anchoring-and-adjustment, the
 246 anchor provides a meaningful starting point to begin the elaborative itera-
 247 tion process. Numeric Priming suggest that the anchor acts as an uncontext-
 248 tualized information node that is assimilated into the judgment. Selective
 249 Accessibility suggests that the anchor acts as a cue that elicits a confirmatory
 250 hypothesis test; increasing the accessibility of anchor-consistent knowledge.

251 The Attitudinal Perspective asserts that the anchor is treated as a persuasive
252 message, such that the individual treats the anchor as a conversational hint.
253 Lastly, the Scale-distortion theory of anchoring proposes that the anchor in-
254 forms the scale with which the judgment is made. Despite these differences,
255 the above theories share the proposition that the anchor is a source of infor-
256 mation, in one way or another. If it is fair to assume that the theories agree
257 that the anchor is an informational source that influences judgments, then
258 how might we use this insight to develop our model?

259 While AIM is novel in its development and execution, it is not the first
260 model intended to capture anchoring effects. Past research on belief updating
261 in order effects (Hogarth and Einhorn, 1992) and preference reversals (John-
262 son and R., 2005; Goldstein and Einhorn, 1987) have used an anchoring-and-
263 adjustment mechanism as a means of describing a particular effect (as an
264 alternative, see Birnbaum and Zimmermann, 1998); though few models have
265 focused primarily on anchoring (e.g., Bhatia and Chaudhry, 2013). To vary-
266 ing extents, these models relate in that they treat anchoring as a process of
267 belief updating and information integration. Complementing the intuition of
268 these models, AIM develops a more comprehensive structure for quantifying
269 this updating process.

270 This notion of information updating is neither new nor is it exclusive to
271 anchoring and related effects. Within other areas of cognition, researchers
272 have long studied how the mind integrates the presentation of a novel stim-
273 ulus with their existing knowledge (for a review, see Anderson, 1990). If
274 anchoring effects involve the integration of a novel stimulus with individuals’
275 prior information, in a manner similar to other areas of cognition, then intu-
276 ition from these other areas may be particularly useful in developing AIM.
277 Within cognitive psychology and related fields, researchers not only attempt
278 to develop theories for how the mind processes information, but also develop
279 quantitative models that 1) allow for the prediction of effects, and 2) pro-
280 vide a framework for testing theoretical predictions regarding psychological
281 processes; we will return to the latter point in the General Discussion.

282 There are many dynamic models of perception that have demonstrated
283 the usefulness of assuming the integration of novel stimuli into existing knowl-
284 edge representations. How these models regulate the integration of informa-
285 tion provides an important distinction between them. A special class of mod-
286 els assume that the integration occurs within the Bayesian framework and are
287 often referred to as “Bayesian models of cognition.” Bayesian models of cog-
288 nition incorporate both prior information and the information provided by

the stimuli to generate a response. These models have become very popular because they provide a convenient way to not only model the cognitive mechanisms underlying the behavior of interest, but they also provide a method for rule-based generalization of cognitive processes. Bayesian models have been used to explain lower-level cognitive processes such as visual scene perception (Yuille and Kersten, 2006), semantic memory (Steyvers et al., 2006), and language acquisition (Chater and Manning, 2006). For decades, research in judgment and decision making has, in large, assumed more elaborative cognitive processes (e.g., judgment and choice) to be simplifying heuristics rather than an algorithmic procedure (for an exhaustive review of this literature, see Gilovich et al., 2002). In contrast to these simplifying heuristics, recent research has focused on developing Bayesian models to account for more elaborative judgments such as human inductive learning (Tenenbaum et al., 2006), social cognition (Baker et al., 2007), and future predictions (Griffiths and Tenenbaum, 2011).

Given the discussion above, we would suggest that anchoring effects involve two primary informational inputs: 1) individuals' prior representations about a target judgment (i.e., the information the individual would have available in the absence of the anchor value), and 2) the information provided by the anchor. In the versions of AIM presented subsequently, we remain agnostic as to *how* the anchor provides information; how the anchor provides information is discussed extensively within the existing anchoring theories. We generally assume that the integration of the anchor information with one's prior information will occur in a manner similar to Bayesian updating. Although Bayesian updating models are one way of conceptualizing the integration, there also exists relevant alternate models (e.g., Turner et al., 2011).

The basic intuition of AIM is illustrated in Figure 1. AIM assumes that individuals have some prior representation of any target judgment. For instance, consider an individual who believes the river to be roughly between 1,000 and 3,000 miles long, with their best guess at around 2,000 miles. AIM conceptualizes this representation as the gray information density function in Figure 1A. For the sake of this example we use simple Gaussian distributions, but we discuss distributional assumptions for AIM when we present the first version in the subsequent section. Because individuals can have different amounts of information, or confidence, in their prior representation, differences are captured in terms of how diffuse the representations are. For example, an individual with more information about the Mississippi River

327 being 2,000 miles, or having more confidence in their information, would
328 have a less diffuse prior representation (Figure 1B). Conversely, individuals
329 with less information or confidence would have a more diffuse representation
330 (Figure 1C).

331 Now consider that the individual is presented an anchor about the Mis-
332 sissippi River being 3,500 miles long. The green distributions in Figure 1
333 illustrate the anchor information. For the sake of this illustration, consider
334 the anchor information as an information density function centered on 3,500
335 miles. This function is diffuse to illustrate the weakness of the anchor in-
336 formation. For instance, an anchor can provide more or less information,
337 which is captured in the variance of the green information density functions.
338 If an anchor provides relatively less information, then it will have a more
339 diffuse representation (e.g., Figure 1D). If an anchor provides relatively more
340 information, then it will have a less diffuse representation (e.g., Figure 1E).

341 The focal component of AIM is the integration of anchor information
342 with prior information to form a posterior representation of the target item.
343 In the model derivation sections below we formalize this integration. In the
344 meanwhile, to offer some insight into the information integration consider
345 again Figure 1. The posterior representation is an amalgamation of the prior
346 and anchor information. The anchoring effects are characterized by the shift
347 from the prior to the posterior representation. Given the AIM framework,
348 anchoring effects are determined by the relative amount of prior information
349 to anchor information. This intuition is partly informed by the observation
350 that less informative anchors produce smaller anchoring effects (e.g., Wil-
351 son et al., 1996) and also by the observation that greater prior knowledge
352 attenuates anchoring effects (e.g., Smith et al., 2013; Smith and Windschit,
353 2015). Thus, compared to Figure 1A, smaller anchoring effects occur when
354 individuals have more prior information (Figure 1B) or less informative an-
355 chors (Figure 1D). Similarly, larger anchoring effects occur when individuals
356 have less prior information (Figure 1C) or more informative anchors (Figure
357 1E).

358 The above example is intended as a useful illustration for how AIM con-
359 ceptualizes the prior and anchor information and their integration in the pos-
360 terior representation of the target judgment. In the following sections we for-
361 malize this intuition and derive, first a version of AIM intended for between-
362 participant data, then a version of AIM intended for within-participant data.

Figure 1: Illustrative example of how anchor information and prior information are integrated when updating the posterior representation of a target item. Compared to Panel A, Panels B and D have a smaller proportion of anchor information to prior information resulting in a smaller influence of anchor information on the posterior representation (i.e., a smaller anchoring effect). Compared to Panel A, Panels C and E have a larger proportion of anchor information to prior information resulting in a larger influence of anchor information on the posterior representation (i.e., a larger anchoring effect).

363 3. A Between-participants Version of AIM

364 We now present the explicit details of AIM so that it can be fit to empirical
 365 data. There are three major components to the model. First, we assume that
 366 for every possible stimulus, individuals have some representation of the target
 367 of interest. For the purposes of this article, the stimuli consist of general
 368 knowledge questions such as “What is the length of the Mississippi River?”
 369 The true solution to these stimuli is called the target, and we denote it Θ_0 .
 370 AIM assumes that each individual has a particular representation of what
 371 they believe to be the true state of Θ_0 , prior to the presentation of the anchor.
 372 We refer to this representation of Θ_0 as the prior representation, and its
 373 *support* is denoted Θ . The second component of AIM quantifies the amount
 374 of *perceived* information the anchor value α provides about Θ_0 , and we refer
 375 to this component as the information function. When presented an anchor
 376 value α , AIM assumes that individuals combine their prior representation
 377 with the information provided by the anchor to form a new representation
 378 about the target Θ_0 . This new representation is the third component of AIM,
 379 and we refer to it as the posterior representation. We will now describe each
 380 of these components in turn.

381 3.1. The Prior Representation

382 The first component of our model is the prior representation. The prior
 383 representation serves as the observer’s internal representation of Θ_0 , which
 384 we will temporarily denote $\text{Prior}(\Theta)$. The prior representation is a function
 385 over Θ that provides the probability that each value of Θ is equal to the true
 386 state Θ_0 (i.e., the density of information at any given value). This represen-
 387 tation could take any form. It could be a single point on the continuum of
 388 possible values (e.g., $\text{Prior}(\Theta)=2,320$), or it could be a distribution over an
 389 entire continuum (e.g., $\text{Prior}(\Theta) = \text{Uniform}[0, 10,000]$). The information an
 390 observer has about Θ_0 is twofold. First, the prior representation provides in-
 391 formation about the most plausible value for Θ_0 , which can be determined by
 392 finding the point along the continuum Θ that corresponds to the highest den-
 393 sity in the prior representation. Second, the prior representation quantifies
 394 the amount of uncertainty in the individual’s information about Θ_0 . Prior
 395 representations with larger variances will reflect larger degrees of uncertainty
 396 whereas narrower prior representations will reflect a strong tendency toward
 397 a particular range along the continuum Θ (i.e., indicates greater prior infor-
 398 mation about the target item).

399 For this initial instantiation of AIM, we assume a convenient form for the
400 prior representation. Specifically, we assume that the prior representation can
401 be well approximated by a normal distribution with mean parameter θ and
402 standard deviation parameter δ . To justify this assumption, it will sometimes
403 be necessary to transform the space of Θ so that a normal distribution is
404 appropriate. Because the prior representation will depend on both θ and δ ,
405 we will henceforth make this explicit in the notation by denoting the prior
406 representation as $\text{Prior}(\Theta|\theta, \delta)$.

407 Although an assumption of normality may be a naïve one, we argue that
408 the normal distribution offers compelling psychological plausibility. First,
409 the normal distribution easily quantifies a “best estimate” in the form of the
410 mean parameter θ . Second, the normal distribution also provides a parame-
411 ter for the dispersion of the representation δ , which reflects the uncertainty
412 about Θ_0 . Larger values of δ will spread the prior over a much larger region
413 of Θ , reflecting larger uncertainty about the target item. In addition, the
414 choice of a normal distribution offers psychological plausibility by applying
415 a symmetric, exponentially decreasing weight to values as a function of their
416 distance from the mean θ – a common assumption in cognitive modeling
417 with respect to both time and similarity (e.g., Rubin and Wenzel, 1996; Liu
418 and Aitkin, 2008; Wixted, 1990; Lee, 2004; Lee and Wagenmakers, 2010;
419 Treisman and Williams, 1984; Turner et al., 2011; Nosofsky, 1986; Kruschke,
420 1992). Whereas the assumption of a convenient distribution works well for
421 the current data, we advise that other more complicated prior representations
422 may be necessary for some data (e.g., Lewandowsky et al., 2009).²

423 3.2. *The Influence of the Anchor*

424 We now turn to the quantification of the influence of the anchor. Al-
425 though there are a myriad of supporting materials to justify the use of a
426 normal prior representation, there are far fewer to aid us in selecting how
427 the anchor might influence an individual’s estimate. We draw from assump-
428 tions commonly made in categorization and similarity models. These models
429 assume that as stimuli are presented, they activate internal representations
430 of categories as a function of their distance from the perceived stimuli (e.g.,
431 Nosofsky, 1986; Kruschke, 1992). The further these perceptions are from

²We are currently investigating more complicated, infinite dimensional prior represen-
tations using nonparametric Bayesian alternatives.

the category representations, the less active the internal category becomes, which results in a decrease in the likelihood of that particular classification. Unlike categorization models, AIM assumes a continuous internal representation for the “classification” of the true state of Θ_0 . AIM assumes that the anchor possesses information about Θ_0 that influences the internal representation. Clearly, some anchors will have a stronger influence on an individual’s representation than others. To accommodate this feature, the information function can take many forms depending on the value of an “anchor influence” parameter λ . Because of the explicit dependence of the information function on the anchor α and the anchor influence parameter λ , we denote the information function $\text{Info}(\Theta|\alpha, \lambda)$.

We will again assume that the influence of the anchor has a convenient functional form. The form of the information function will take one of two distributions, depending on whether or not the anchor provides directional information. For anchors that do not provide directional information, we assume that the information function is normally distributed with mean equal to the anchor α and standard deviation equal to λ , such that

$$\text{Info}(\Theta|\alpha, \lambda) = \frac{1}{\sqrt{2\pi}\lambda} \exp \left[-\frac{(\Theta - \alpha)^2}{2\lambda^2} \right].$$

It is important to note that some anchors can contain directional information (Simmons et al., 2010). For example, whereas an anchor about the Mississippi river being 3,500 miles long does not necessarily contain directional information, an anchor like “the length of the Mississippi river is less than 3,500 miles long” provides specific information about the range of values of Θ_0 . We argue that directional and regular anchors have similar psychological properties but directional anchors will carry additional information provided by the item. If the anchor α carries with it a negative directional influence (e.g., the length of the Mississippi river is less than 3,500 miles long), then we assume that the information function will be normal in form but will have zero density in the region to the right of the anchor, such that

$$\text{Info}(\Theta|\alpha, \lambda) = \begin{cases} \frac{1}{\sqrt{2\pi}\lambda} \exp \left[-\frac{(\Theta - \alpha)^2}{2\lambda^2} \right] & \Theta \leq \alpha \\ 0 & \Theta > \alpha. \end{cases}$$

Similar functions can be formed to accommodate positive directional influences by substituting Θ with $-\Theta$, and anchors that are negative in value by

replacing α with $-\alpha$.

AIM’s assumptions about the influence of the anchor are made mostly out of convenience. Specifically, by using a normal information function concurrently with a normal prior representation, we can provide closed form expressions for the posterior representation, which is convenient for estimation purposes. Notice that for directional anchors, the information function is not a probability distribution (i.e., it does not integrate to one). However, as we show below, we do not require the information function to be a probability distribution due to the integration step used to derive the posterior representation. Specifying the information function in this way allows us to compare the parameter λ across a variety of anchor types (e.g., regular and directional).

3.3. The Posterior Representation

The final component of the model is the posterior representation. Similar to the prior representation, the posterior representation is a distribution over Θ , and reflects the underlying beliefs about Θ_0 following the presentation of the anchor value α . The posterior representation serves as the final result of the integration of *a priori* information (i.e., the prior representation) and information provided by the anchor. There are many ways in which the prior representation could be combined with the information function to form this new posterior representation.³ One convenient way to combine two distributions is through Bayes’ rule.

We denote the posterior representation as $\text{Post}(\Theta|\theta, \delta, \alpha, \lambda)$, and we obtain it by multiplying the prior representation $\text{Prior}(\Theta|\theta, \delta)$ by the information function $\text{Info}(\Theta|\alpha, \lambda)$ and marginalizing, such that

$$\text{Post}(\Theta|\theta, \delta, \alpha, \lambda) = \frac{\text{Prior}(\Theta|\theta, \delta)\text{Info}(\Theta|\alpha, \lambda)}{\int \text{Prior}(\Theta|\theta, \delta)\text{Info}(\Theta|\alpha, \lambda)d\Theta}. \quad (1)$$

We again note that this posterior representation is a probability density function that provides the probability that each Θ equals the individual’s estimate of the true state Θ_0 .

³Indeed, the combination of these two sources of information may vary across tasks, context, anchor types, and individuals. However, for our purposes, we will not allow the combination process to vary.

490 The marginalization constant in the denominator of Equation 1 can be
 491 difficult to evaluate. However, our assumptions about how the prior repre-
 492 sentation is transformed into the posterior representation allow us to derive
 493 analytic expressions for both the prior and posterior representations. We use
 494 the terms prior and posterior to refer to the representations available before
 495 and after the presentation of an anchor. That is, a prior representation is the
 496 information an individual would have available to make the target judgment
 497 when no anchor is present. A posterior representation is the information an
 498 individual has which includes the prior and anchor information. Within the
 499 standard anchoring paradigm, only the posterior representation is actually
 500 constructed to form a response (i.e., a judgment is not elicited before the
 501 presentation of the anchor). Therefore, the prior representation in our model
 502 represents the hypothetical prior representation that would have been con-
 503 structed had a judgment been elicited without an anchor. Hence, for most
 504 applications, the prior representation is a latent construct to be inferred from
 505 the data.

506 3.3.1. *Eliciting a Response*

507 Although AIM generally assumes that individuals have access to an en-
 508 tire (posterior) representation (i.e., not a single point), participants in an
 509 experiment are typically asked to form a response of a single value. As a
 510 consequence, to form a testable model we must specify an elicitation mecha-
 511 nism that converts the posterior representation $\text{Post}(\Theta|\theta, \delta, \alpha, \lambda)$ to the par-
 512 ticipant’s estimate $\hat{\Theta}$. When the estimate $\hat{\Theta}$ is a single value, we could simply
 513 adopt the rule $\max_{\Theta} \{\text{Post}(\Theta|\theta, \delta, \alpha, \lambda)\} = \hat{\Theta}$. However, this particular rule
 514 may not be what the individual had intended. For example, this rule may not
 515 be the best choice if $\text{Post}(\Theta|\theta, \delta, \alpha, \lambda)$ is a bi-modal or uniform distribution.

516 Rather than assuming that an individual always produces an estimate
 517 at the mode of their representation, AIM assumes that individuals sample
 518 evidence probabilistically from their posterior representation. For example,
 519 when an individual provides an estimate, they randomly sample a single
 520 value from their posterior representation, and as a consequence, they are more
 521 likely to sample from higher-density regions. We assume random sampling for
 522 convenience, but acknowledge that individuals more frequently respond with
 523 rounded numbers. Although the actual process may be a pseudo-random
 524 sampling of the posterior representation, we assume random sampling within
 525 the model for simplicity.

526 The amount of sampled evidence n may be treated as another param-

527 eter in the model, and it performs similar to other “threshold” parameters
 528 common in models of evidence accumulation, such as models of response
 529 time (e.g., Ratcliff, 1978; Usher and McClelland, 2001; Brown and Heath-
 530 cote, 2005, 2008). In this article, we will not explore the role n plays and will
 531 instead assume that $n = 1$ throughout. Thus, a single sample of evidence
 532 is drawn from the posterior representation, and this sample serves as the
 533 participant’s best estimate for Θ_0 .

534 3.4. Summary

535 AIM has three parameters that control how each source of information
 536 is used in generating the observed data. Figure 1 shows some hypothetical
 537 model predictions under a few settings of the model parameters. In all pan-
 538 els, the mean of the prior representation θ is held constant to exhibit the
 539 influence of the other parameters δ and λ . The prior representation, anchor
 540 information function, and posterior representation are illustrated by the gray,
 541 green, and red density functions, respectively. For example, Figures 1B and
 542 1C illustrate the effects of the parameter δ on the posterior representation.
 543 Under strong prior information (i.e., a small value of δ), the prior repre-
 544 sentation is less influenced by the anchor information, and as a result, the
 545 posterior representation is nearer to the prior. However, under weak prior
 546 information (i.e., a large value of δ), the prior representation is susceptible
 547 to the influence of the anchor, and so the posterior representation is nearer
 548 to the anchor value.

549 Figures 1D and 1E show the influence of the anchor influence parameter
 550 λ under weak and strong influence settings, respectively. Under weak anchor
 551 information (i.e., large values of λ ; Figure 1D), the prior representation is
 552 less influenced, resulting in only a small shift from the prior representation
 553 to the posterior representation. Under strong anchor information (i.e., small
 554 values of λ ; Figure 1E), the prior representation is highly influenced, result-
 555 ing in a relatively large shift from the prior representation to the posterior
 556 representation.

557 Having discussed all the major components of AIM, we can now turn
 558 to evaluating the model on the basis of empirical data. Experiment 1 was
 559 designed to demonstrate AIM’s ability to describe the effects of an assortment
 560 of anchor values. We first describe the details of the experiment, and all of
 561 the essential effects present in the data. We then discuss the implementation
 562 details of the between-participants version of AIM and fit it to data. We
 563 close the next section by discussing the modeling results.

Table 1: Outline of between-participant conditions for Experiment 1. Key: CL is Cognitive Load

Condition	Description/Example
Prior	Unanchored judgment
Regular	Is the Mississippi River greater than or less than 3,500 miles?
Directional	The Mississippi River is less than 3,500 miles.
Irrelevant	Is the distance between Cairo and Istanbul greater than or less than 3,500 miles?
CL	Memorized 9-letter string before anchor then provided string after the judgment.

564 4. Experiment 1

565 Experiment 1 was intended to test the model across a range of judgment
566 targets and a wide variety of anchoring effects. The following sections will
567 demonstrate the etiologies of regular anchoring effects, directional anchors,
568 irrelevant anchors, and the influence of cognitive load on anchoring effects.

569 4.1. Methods

570 Participants were recruited from Amazon’s Mechanical Turk website (see
571 Paolacci et al., 2010; Buhrmester et al., 2011, for a review of Mechanical
572 Turk as a viable source of experimental data) and completed the task in
573 exchange for \$0.50. A total of 589 participants (sub sample information
574 will be provided by condition) completed one of the five between-participant
575 conditions: prior, regular anchors, directional anchors, irrelevant anchors,
576 and cognitive load.

577 4.1.1. Materials

578 Participants in every condition were presented the same 6 anchoring prob-
579 lems taken from Simmons et al. (2010) and Jacowitz and Kahneman (1995).
580 The six target questions were as follows: 1) the length of the Mississippi
581 River, 2) the average annual rainfall in Philadelphia, 3) the maximum speed
582 of a house cat, 4) the average annual temperature of Phoenix, 5) the pop-
583 ulation of Chicago, and 6) the height of the tallest redwood tree. Once
584 the correct answer was determined, for some questions we set anchors that
585 were either higher or lower than the true value. Questions 1, 3, and 5 used
586 high anchors (3,500 miles, 55 miles per hour, and 5 million people, respec-
587 tively) whereas Questions 2, 4, and 6 used low anchors (25 inches, 52 degrees
588 Fahrenheit, and 65 feet, respectively).

4.1.2. Procedure

The five between-participant conditions are outlined in Table 1. In all conditions participants estimated each of the six target items. Thus, Experiment 1 employed a 5 between-participant (Anchoring Condition) x 6 within-participant (Target Questions) mixed design. In addition to the prior (i.e., unanchored control condition) and regular anchor conditions, we included three additional anchoring paradigms that have received notable attention within the anchoring literature. Because anchors often provide directional information we included a condition with directional anchors (Simmons et al., 2010). The literature on irrelevant anchors and numeric primes has been controversial to say the least, with evidence toward (Adaval and Monroe, 2002; Critcher and Gilovich, 2008; Englich, 2008; Oppenheimer et al., 2008; Mussweiler and Englich, 2005; Wilson et al., 1996; Wong and Kwong, 2000) and against (Brewer and Chapman, 2002; Chapman and Johnson, 1994; Frederick and Mochon, 2012; Mochon and Frederick, 2013; Mussweiler and Strack, 2001) such effects. As such, we included a condition where the anchors were in reference to items different than those of the target judgments. Lastly, given the research that used cognitive load as a means of studying anchoring effects (Epley and Gilovich, 2001, 2006; van Rooijen and Daamen, 2006; Blankenship et al., 2008), we included a condition where participants completed the regular anchoring task under cognitive load. We included this range of conditions to test AIM’s ability to model a range of anchoring paradigms.

Borrowing terminology from above, we refer to estimates made by participants in the prior condition as prior estimates and estimates made by participants in the anchoring conditions as posterior estimates. Roughly 3% of the data were deleted due to either extremely implausible answers or because they attempted to look up the answers on-line. We were able to identify the most accessible answers through Wikipedia, Google, and Yahoo searches. If participants had longer response times (e.g., greater than 10 seconds) and answers identical to the ones we identified through these websites, those observations were deleted. Also as part of the 3% of deletions, estimates in the directional anchoring condition that were on the incorrect side of the anchor (e.g., an estimate of 4,500 miles to an anchor that was stated to be less than 3,500 miles) were deleted. This method of deletions is consistent with Epley and Gilovich (2001, 2006) and Simmons et al. (2010).

624 4.2. Technical Details

625 To appropriately constrain the model, we assumed that responses elicited
 626 in the prior condition were reflective of the general population prior for each
 627 question. In this particular version of AIM, we are not concerned with how
 628 anchors affect *individual* behavior. Instead, we are only concerned with the
 629 effects particular anchors have on *groups* of individuals. Later we will extend
 630 the basic model to incorporate individual differences.

631 We first applied a log transformation to the data for each question, and
 632 fit each question independently. We denote data from the prior condition as
 633 x and data from each of the $K = 4$ anchoring conditions as y_k . We assumed
 634 that the data x arise from a normal prior representation with mean θ and
 635 standard deviation δ , such that

$$\begin{aligned} x &\sim \text{Prior}(\theta, \delta) \\ &\sim \mathcal{N}(\theta, \delta), \end{aligned} \tag{2}$$

636 where $\mathcal{N}(a, b)$ denotes a normal distribution with mean a and standard de-
 637 viation of b .

638 We then assume that, upon presentation of the anchor value α_k in the k th
 639 condition, the prior representation is updated through multiplication of the
 640 information function, which is governed by the anchor influence parameter
 641 λ_k . Here, we assume that the information function is normal in form, such
 642 that

$$\text{Info}(\Theta|\alpha_k, \lambda_k) = \frac{1}{\sqrt{2\pi}\lambda_k} \exp \left[-\frac{(\alpha_k - \Theta)^2}{2\lambda_k^2} \right]. \tag{3}$$

643 Following the presentation of the anchor, the participant updates their prior
 644 representation by the information function to arrive at the posterior repre-
 645 sentation

$$\text{Post}(\Theta|\theta, \delta, \lambda_k, \alpha_k) = \frac{\text{Prior}(\Theta|\theta, \delta)\text{Info}(\Theta|\alpha_k, \lambda_k)}{\int \text{Prior}(\Theta|\theta, \delta)\text{Info}(\Theta|\alpha_k, \lambda_k)d\Theta}.$$

646 The assumption of normality in the prior representation and the normal
 647 information function enables us to provide closed form expressions for the
 648 posterior representation. By combining Equations 2 and 3, the posterior

649 representation can be expressed as

$$\begin{aligned}
\text{Post}(\Theta|\theta, \delta, \lambda_k, \alpha_k) &= \frac{\exp\left[-\frac{(\Theta - \theta)^2}{2\delta^2}\right] \exp\left[-\frac{(\alpha_k - \Theta)^2}{2\lambda_k^2}\right]}{\int \exp\left[-\frac{(\Theta - \theta)^2}{2\delta^2}\right] \exp\left[-\frac{(\alpha_k - \Theta)^2}{2\lambda_k^2}\right] d\Theta} \\
&\propto \exp\left[-\frac{1}{2}\left(\frac{1}{\delta^2} + \frac{1}{\lambda_k^2}\right)\left(\Theta - \left(\frac{\frac{\theta}{\delta^2} + \frac{\alpha_k}{\lambda_k^2}}{\frac{1}{\delta^2} + \frac{1}{\lambda_k^2}}\right)\right)^2\right] \\
&= \frac{1}{\sqrt{2\pi}b_k} \exp\left[-\frac{(\Theta - a_k)^2}{2b_k^2}\right], \tag{4}
\end{aligned}$$

650 which is a normal distribution with mean a_k and standard deviation b_k , where

$$\begin{aligned}
a_k &= \frac{\frac{\theta}{\delta^2} + \frac{\alpha_k}{\lambda_k^2}}{\frac{1}{\delta^2} + \frac{1}{\lambda_k^2}}, \text{ and} \\
b_k &= \left(\frac{1}{\delta^2} + \frac{1}{\lambda_k^2}\right)^{-1/2}. \tag{5}
\end{aligned}$$

651 Thus, we assume that the data from the k th anchoring condition comes from
652 the distribution

$$\begin{aligned}
y_k &\sim \text{Post}(\theta, \delta, \lambda_k, \alpha_k) \\
&\sim \mathcal{N}(a_k, b_k).
\end{aligned}$$

653 Because the data y_k are sampled from a posterior representation that depends
654 on a latent prior representation, the parameters θ and δ controlling the prior
655 representation are not fully observed. That is, the estimates for both θ and
656 δ are obtained *jointly* from the data from the prior condition x and the data
657 from the anchoring conditions y_k .

658 Although it would be convenient to have a single estimate for λ corre-
659 sponding to each anchoring condition across questions, the questions are on
660 different scales making it unclear how to constrain the information function

661 to incorporate this idea. Thus, when fitting the model to the data, we par-
662 titioned the data by question, so that the analysis was independent across
663 questions. In the analysis of each question, we avoided building a hierarchical
664 structure on the λ_k s because it was not clear what distribution the λ_k s might
665 take, and we wished to avoid any unjustified (and unnecessary) restrictive
666 assumptions.

667 Figure 2 displays a graphical diagram of the model used in each ques-
668 tion. These diagrams are often used in cognitive modeling because they
669 convey the specification of a mathematical model in a convenient path dia-
670 gram (e.g., Lee, 2008; Lee and Wagenmakers, 2010). Each “node” in Figure
671 2 corresponds to a separate variable in the model, and shaded nodes repre-
672 sent variables that are known, such as the data x or the value of the anchor
673 α_k . Nodes that are connected to one another indicate that the variables in
674 the nodes are dependent. This dependence travels one way, such that the
675 “pointing” node controls the variable it points to. For example, Figure 2
676 shows the dependence of the data y_k on all three parameters of the model, θ ,
677 δ , and λ_k , which is expressed mathematically in Equations 4 and 5. Finally,
678 graphical diagrams represent vector-valued dependencies among variables by
679 surrounding these variables by a rectangle, which is commonly referred to
680 as a “plate.” In Figure 2, the variables λ_k , α_k , and y_k are all on the plate
681 because they each contain four elements – one element for each (posterior)
682 Condition k .

683 4.3. Results

684 To fit the model to the data in a Bayesian framework, we implemented a
685 Differential Evolution Markov chain Monte Carlo (DE-MCMC) sampler with
686 24 chains (see ter Braak, 2006; Turner et al., 2013, for details), although we
687 have provided JAGS (Plummer, 2003) code for performing the estimation in
688 Appendix C. We ran the chains in parallel for 10,000 iterations and discarded
689 the first 2,000 iterations as a burnin, resulting in a total of 192,000 samples
690 from the posterior distribution. Standard methods were used to ensure con-
691 vergence. Generalized equations for performing the analysis are presented in
692 Appendix A, whereas specific details for this analysis were discussed above.

693 Figure 3 shows a violin plot of the posterior distributions for each of the
694 λ_k parameters on the log scale. Each panel in Figure 3 shows the estimates
695 for λ for each anchor type (x -axis) for a particular question. Because λ is
696 the standard deviation of a normal distribution, as λ increases, the influence
697 of the anchor diminishes and the prior representation is more reflective of

Figure 2: A graphical diagram of the between-subjects anchor integration model.

the posterior representation. Across questions, the estimates for λ have little meaning, but within a question, the relative differences in the estimates convey the perceived importance of the anchor. For example, across questions we see that the estimates for λ in the irrelevant anchor condition are larger than the estimates in other conditions. This result demonstrates that AIM is capturing patterns in the data that conform with the a priori hypothesis that irrelevant anchors should produce smaller anchoring effects. The pattern of results for λ estimates in other conditions is less clear. For example, λ in the directional condition is sometimes lower (e.g., Questions 1, 2, and 3), and sometimes larger (e.g., Questions 4, 5, and 6) than λ in the regular condition.

We can also examine the posterior distributions for the parameters θ and δ corresponding to the prior representation. Table 2 shows the median and 95% Bayesian credible set of the estimated posterior distributions for θ and δ , separated by question. For example, on Question 1 (The length of the Mississippi River), the mean estimate of θ was $6.916 = \log(1008.279)$ and the estimated uncertainty parameter δ was 0.858. As with the parameter λ , it is difficult to draw comparisons across items for these parameters because the range of their support is different.

The final analysis we performed was to generate model predictions for

Figure 3: The estimated posterior distribution of λ_k for each of the six anchoring conditions for each question. From left to right, the anchor types are: regular (R), directional (D), irrelevant (IR), and cognitive load (CL).

Table 2: Summaries of the estimated posterior distributions of θ and δ obtained for each item in Experiment 1.

Question	$\hat{\theta}$		$\hat{\delta}$	
	Median	Credible Set	Median	Credible Set
1	6.916	(6.834, 6.999)	0.858	(0.811, 0.910)
2	3.287	(3.228, 3.345)	0.738	(0.690, 0.788)
3	2.526	(2.454, 2.599)	0.791	(0.750, 0.833)
4	4.427	(4.413, 4.440)	0.147	(0.140, 0.154)
5	1.316	(1.256, 1.377)	0.767	(0.720, 0.817)
6	5.244	(5.166, 5.320)	0.807	(0.764, 0.854)

717 both the prior and posterior representations used by the subjects in this ex-
 718 periment. To do this, we randomly sampled 10,000 observations from the
 719 joint posterior distribution of all model parameters, and used these values to
 720 construct both the prior and posterior representations from the model. The
 721 resulting distributions are known as posterior predictive distributions, and
 722 are plotted in Figures 4, 5, 6, and 7 as gray lines against the density of the
 723 data (black lines). Each panel in these figures corresponds to a different ques-
 724 tion used in the experiment. The black dashed lines represent the observed
 725 data from the prior condition, and is reproduced in each of the four figures
 726 for convenience (e.g., dashed black line in the top left panel of Figure 4 is the
 727 same as the dashed black line in the top left panel of Figure 5). In general,
 728 AIM closely matches the observed data across all conditions and questions.
 729 In conditions where the posterior did not differ appreciably from the prior
 730 (e.g., the irrelevant anchoring condition) AIM predicts a similar pattern and
 731 is less influenced by the anchor value. By contrast, in conditions where the
 732 anchor had a remarkable influence on the prior, such as in the directional
 733 condition, our model peaks in the same regions as the data indicating a sim-
 734 ilar interaction with the provided anchor. In conditions where the model
 735 tends to provide a relatively poorer fit to the data, the normality assump-
 736 tion about the prior distribution tends to be invalid. In these conditions, if
 737 one assumed a more complicated function for the prior representation or the
 738 information function, one could obtain better fits to the data.⁴ We will now
 739 discuss the results in more depth for each type of anchor in turn.

740 4.3.1. *The Prior Condition*

741 To get an accurate approximation of the aggregate prior representation of
 742 individuals in the anchoring conditions, 323 participants (61% female with a
 743 mean age of 35) completed the task without the presentation of any anchors.
 744 Experiment 1 hinges on the assumption that the distribution of estimates in
 745 the control (i.e., prior) condition were representative of the prior representa-
 746 tions of participants in the anchoring conditions.

747 The prior distributions for each condition appear in Figures 4, 5, 6, and 7
 748 as the dashed black lines. Because these priors are from a separate condition
 749 of the experiment and do not depend on the type of anchor presented, the

⁴For example, one could assume a Dirichlet process or finite mixture distribution for the prior representation (or information function or both), and obtain much better fits.

priors shown in each of these figures are equivalent, and are only shown for interpretation purposes for the other experimental conditions. That is, because the magnitude of the shift from prior to posterior representation depends on the anchor type, we can interpret the effects of the anchor in each condition by examining the discrepancy between the solid and dashed curves in each of these four figures.

The model fits to the prior condition are shown as the dashed gray densities in Figures 4, 5, 6, and 7. Generally speaking, the model captures the shape of the prior representations well, despite the noise present in the data. To better account for this noise, one could increase the flexibility of the prior representation. However, our Gaussian prior representation assumption provided an adequate fit for our data, and so we save a thorough investigation of these more complex functional forms for future research.

4.3.2. *Regular Anchors*

Sixty participants (56% female with a mean age of 33) answered a comparative question (e.g., is the Mississippi River greater than or less than 3,500 miles long?) about whether the target was greater than or less than the anchor value we provided, before making their estimate, for each of the six questions. The distribution of participants' responses for each of the six questions are presented in Figure 4. As mentioned above, the black dashed lines represent the distribution of estimates provided by participants in the prior condition, whereas the black solid lines are the distributions of the estimates provided by the participants in the regular anchoring condition. The solid gray lines represent the model predictions for the regular anchoring condition. In general, the regular anchor does provide some noticeable shift in the posterior representation relative to the prior representation. This shift is most pronounced for Question 6, and is also noticeable for Question 1. The model captures these shifts well: for Questions 1 and 6, the model predicts a shift of a similar magnitude and in the same direction as the empirical data, whereas for the remaining questions, the shift is negligible in a way that matches the observed data.

4.3.3. *Directional Anchors*

Seventy-two participants (65% female with a mean age of 34) were presented an anchor statement (e.g., the Mississippi River is less than 3,500 miles long) before producing an absolute judgment. Figure 5 shows the density of the posterior distributions for each question in the directional anchoring

Figure 4: The distribution of the log-estimates of Θ for each question in the regular anchors condition (solid black lines) along with the distribution of the log-estimates in the prior condition (dashed black lines). The corresponding posterior predictive densities obtained after fitting the model are shown by the gray lines (prior representations are dashed whereas the posterior representations are solid).

condition as solid black lines. The model predictions are again shown as the solid gray lines. Figure 5 shows that the effects of the directional anchors are very strong, producing large differences between the prior and posterior representations. In particular, for all of the questions except Question 4, the posterior representation is more peaked, having less variance relative to the prior representation. This pattern reflects an increase in the amount of available information in forming the judgment following the presentation of the anchor, and suggests that participants do use the information provided by the anchor when eliciting their response.

Generally speaking, Figure 5 shows the model captures the anchor effects well, producing marked changes in the posterior representation that closely match the changes observed in the data. The fits are particularly good for Questions 1, 2 and 5, whereas some misfit occurs for Questions 3, 4, and 6. The reason for these misfits is mixed. For example, in Question 3, the misfit of the posterior representation seems to be due to a misestimation of the prior representation, whereas in Question 6, the misfit seems to be due to a misestimation of the anchor influence. We suspect that much of this misestimation is due to irregularities in the representations themselves, namely that these representations are not Gaussian.

4.3.4. Irrelevant Anchors

Seventy participants (67% female with a mean age of 37) were presented the six anchoring questions with irrelevant comparison questions (e.g., is the distance between Cairo and Istanbul greater than or less than 3,500 miles?). Irrelevant anchors should not provide much information when forming a posterior representation, and so we would expect the estimates for λ in this condition to be large. However, the proportion of information provided by the anchor will be larger when an individual has little prior information (i.e., when δ is large). Hence, we might generally expect λ and δ to be negatively correlated.

Figure 6 shows the participants' estimates from the prior condition (dashed) against the posterior estimates of participants in the irrelevant anchor condition (solid) for each of the six questions. Figure 6 shows that the anchoring effects in this condition are relatively weak because the posterior representations are highly similar to the prior representations. The model predictions are again shown as the gray lines. Across the panels, Figure 6 shows that the model predictions closely match the patterns in the observed data, but this is unsurprising given that the effects are not particularly strong.

Figure 5: The distribution of the log-estimates for each question in the directional anchors condition (solid black lines) along with the distribution of the log-estimates in the prior condition (dashed black lines). The corresponding posterior predictive densities obtained after fitting the model are shown by the gray lines (prior representations are dashed whereas the posterior representations are solid).

Figure 6: The distribution of the log-estimates for each question in the irrelevant anchors condition (solid black lines) along with the distribution of the estimates in the prior condition (dashed black lines). The corresponding posterior predictive densities obtained after fitting the model are shown by the gray lines (prior representations are dashed whereas the posterior representations are solid).

823 4.3.5. *Cognitive Load*

824 To develop a better understanding of the influence of cognitive load on
825 anchoring effects, 64 participants (56% female with a mean age of 33) com-
826 pleted an identical anchoring task to that of the regular anchor condition.
827 Just before the comparative question was presented with the anchor, par-
828 ticipants were shown a nine letter string and were told to memorize it for
829 later recall. For each question, after producing an absolute estimate of the
830 target item, participants were asked to recall the series of letters. Figure 7
831 plots participants' estimates from the prior condition against posteriors in
832 the cognitive-load condition for each of the six questions. The effects of the
833 cognitive-load condition were relatively strong in Questions 1 and 6, and to
834 a lesser extent in Question 3. The model predictions are again shown as the
835 gray lines, where solid lines represent posterior conditions and dashed lines
836 represent the prior condition. Figure 7 shows that the model again matches
837 the basic patterns in the data well. The match to the data is particularly
838 close in Questions 1 and 6, where the effects of the anchor are strong.

839 4.3.6. *Summary of Experiment 1*

840 Experiment 1 demonstrated that the between-participant version of AIM
841 can account for a wide variety of anchors and moderators. AIM generally
842 produced predictions that closely matched the observed data across six ques-
843 tions and four types of anchoring effects. There were some mismatches in
844 the predictions, but we suspect that the majority of these mismatches are
845 due to the lack of complexity in the prior representation. We speculated that
846 increasing the flexibility of this component of the model would improve the
847 fitting results.

848 In the sections above, we have outlined a between-participants version of
849 AIM that allows us to investigate the effects of different anchor types across
850 groups of individuals. Overall, we would consider the fits of our model to be
851 adequate, considering the extreme variability in participants' estimates and
852 the lack of multiple observations per participant. Although this version of
853 the model is very useful for understanding how the location and context of
854 an anchor influences judgments, we have not provided a detailed modeling
855 account of how anchors might influence individual judgments. In the next
856 section, we extend AIM to account for individual-specific anchoring effects.

Figure 7: The distribution of the log-estimates for each question in the cognitive-load condition (solid black lines) along with the distribution of the log-estimates in the prior condition (dashed black lines). The corresponding posterior predictive densities obtained after fitting the model are shown by the gray lines (prior representations are dashed whereas the posterior representations are solid).

857 5. A Within-participant Version of AIM

858 The version of AIM outlined and tested above provides a cogent descrip-
859 tive explanation of the influences of different types of anchors on participants’
860 judgments. One shortcoming of the model is the assumption of the homo-
861 geneity of the anchor influence parameter across individuals (i.e., predicting
862 anchoring effects at the group level rather than the individual level). In
863 the next section we relax this assumption and extend AIM to account for
864 anchoring effects at the individual level.

865 5.1. Prior Uncertainty and the Role of Confidence

866 The data that we typically acquire from an anchoring experiment has
867 only one judgment per question per individual. Within AIM it becomes
868 difficult to estimate both the influence of an anchor and the uncertainty
869 in the prior representation *simultaneously*. Thus, to investigate individual
870 differences with such little information, it is essential that we either apply
871 reasonable constraints to the model or obtain more information about the
872 judgment. For the prior uncertainty parameter, we will build both of these
873 methods into AIM.

874 In our Experiment 2 (discussed below), we will obtain confidence judg-
875 ments in addition to the estimates of Θ_0 . We disentangle the role of confi-
876 dence by eliciting both prior and posterior judgments within a participant,
877 where their confidence response to the prior judgment represents the prior
878 uncertainty (i.e., how much prior information they have about the target
879 item). We suspect that as the confidence in the judgment increases, the
880 prior uncertainty parameter should decrease. However, we make no restric-
881 tive assumptions about the functional form of this relationship. Instead, we
882 will use a separate prior uncertainty parameter for each confidence response
883 option (e.g., six options if the scale is 1-6). Thus, we can estimate the prior
884 uncertainty by assuming that this parameter is fixed either across different
885 participants or across repeated anchors on the same continuum (i.e., multiple
886 judgments about the same target item).

887 We agree that in some cases, it is unwise to assume that confidence judg-
888 ments have the same meaning, even when they are all about the same target
889 item. For example, it seems unlikely that when Participant A elicits a confi-
890 dence judgment of “1” or “Very Unconfident” that this degree of confidence
891 will have the same meaning for Participant B, who has also elicited a confi-
892 dence judgment of “1”. Even though these two participants have the same

893 degree of confidence (but may have different prior judgments about a target),
 894 their level of *uncertainty* may not be equivalent.

895 While the assumption above may be an overly restrictive one, we empha-
 896 size that this assumption is not necessary for our model. For example, with
 897 multiple confidence judgments per individual, one can assign a separate prior
 898 uncertainty parameter for each category, even when the target dimension Θ_k
 899 is not the same. To do so, we only require that the confidence judgments
 900 be related to one another in some meaningful way. For example, one could
 901 assume that the different confidence judgments come from the same over-
 902 arching process, so that when a confidence judgment of “1” is provided for
 903 the estimate in Question k , it will be related to a confidence judgment of
 904 “1” in Question l . As another example, one could assume a functional form
 905 (e.g., a linear relationship) for how the confidence judgments are related to
 906 the prior uncertainty. As a result, having information about other confidence
 907 judgments such as “1” and “4” will inform our estimates of confidence judg-
 908 ments such as “3”, even if this judgment was never actually elicited in the
 909 experiment.

910 In an attempt to create a more parsimonious model, we now discuss a way
 911 to relate the influence that a given anchor value may have to the anchor’s
 912 consistency with an individual’s prior representation.

913 5.2. *Prior Consistency*

914 In order to make testable predictions about how an anchor might influ-
 915 ence a judgment, we assume that participants use a measure of “prior con-
 916 sistency” to dictate how their prior representations should be updated. That
 917 is, consistency refers to the extent that the anchor information and the prior
 918 representation overlap. Given the prior representation $\text{Prior}(\Theta|\theta, \delta)$ and the
 919 presentation of an anchor value α , we assume that participants determine
 920 how likely the candidate solutions Θ are to be less than the presented anchor
 921 value α , under their prior beliefs. To do so, participants integrate their prior
 922 representation across the space $(-\infty, \alpha)$, or

$$\psi = P(\Theta \leq \alpha) = \int_{-\infty}^{\alpha} \text{Prior}(\Theta|\theta, \delta) d\Theta, \quad (6)$$

which is the cumulative distribution function under a participant’s prior beliefs, and $\psi \in [0, 1]$ because it is a probability measure.⁵

The prior consistency variable ψ now provides a direct evaluation of how consistent a given anchor value is with a participant’s prior beliefs. By inspecting various anchors under a fixed prior representation, we can see how the degree of consistency is non-monotonic. Specifically, as ψ approaches 0.5 (from either direction), the degree of consistency increases, but as ψ moves away from 0.5 (i.e., toward 0.0 or 1.0), the degree of consistency decreases. That is, anchor values that are much lower than the prior representation along the axis Θ will be very inconsistent, which will make ψ very small. By contrast, anchor values that are much larger than the prior representation along Θ will also be very inconsistent, but in this case, ψ will be large because the area under the prior representation will be larger. As a concrete example, consider a prior representation where $\theta = 0$ and $\delta = 1$ and the presentation of three different anchor values; $\alpha = \{-3, -0.5, 1\}$. The top left panel of Figure 8 shows the model under these parameter settings where the solid black line (A1), dashed light gray line (A2) and dashed dark gray line (A2) correspond to the three anchor values respectively. The top right panel of Figure 8 shows the prior consistency values $\psi = \{0.001, 0.309, 0.841\}$ for each anchor with the corresponding vertical lines. The most extreme anchor value $\alpha = -3$ has the smallest prior consistency ($\psi = 0.001$) and the largest anchor value $\alpha = 1$ has the largest ψ value ($\psi = 0.841$), but is still inconsistent with the prior representation.

Although the evaluation of prior consistency is essential for understanding how a given anchor might influence a participant’s representation, we have not yet described how this relationship might unfold. We now discuss how the anchor influence parameter λ is determined by mapping ψ to λ . Whereas λ represents the influence of the anchor for a group of participants in the previous version of the model, it now represents the influence of the anchor for a given individual.

5.3. Anchor Influence

We assume that the influence of the anchor will be a direct function of the prior consistency variable ψ . However, the degree to which the prior

⁵To be clear, we do not contend that participants actually compute an integral to evaluate the prior consistency. However, the calculation is a simple mathematical approximation to the actual process that might occur.

Figure 8: The within-participant version of AIM. The top left panel shows a prior representation under the presentation of three different anchor values. The top right panel shows the mapping function relating the anchor influence parameter λ to the prior consistency variable ψ . The middle and bottom panels show the information functions and posterior representations corresponding to each anchor value, respectively.

consistency variable influences the prior representation is assumed to vary across participants.

5.3.1. Constraining the Estimate of λ

To properly define a mapping function, we must first understand the behavior of the information function when λ becomes large. In particular, the information function has a very similar influence on the prior representation once λ increases beyond a certain value, which we will denote $\lambda_{\max}$. Once $\lambda > \lambda_{\max}$, the prior representation will be very similar to the posterior representation, and further increases in λ will not appreciably change the difference between the prior and posterior representations. This problem is not generally an issue for the previous version of our model because the effects of the anchor are pronounced for enough individuals that λ is much less than $\lambda_{\max}$. However, boundary problems do occur in the extension of our model because some individuals are completely insensitive to the presentation of the anchor, and as a result, we must build this behavior into the model.

Fortunately, a good selection of $\lambda_{\max}$ can be made prior to fitting the model. To do so, we must first examine the distribution of prior judgments across all individuals. Suppose after an experiment, we gather the set of prior judgments x about the target Θ prior to the presentation of an anchor located at α . To select an approximate value for $\lambda_{\max}$, we can use a nonparametric approach to first construct a density estimate $f(x)$ for x , and then assume that this density estimate is an adequate approximation, on average, of any given individual in the data set. We can then iteratively propose values $\lambda_{\max}^*$ for $\lambda_{\max}$ and examine the degree of difference between the prior representation $f(x)$ and the posterior representation, given by

$$\text{Post}(\Theta|\lambda_{\max}^*, \alpha) = \frac{f(x)\text{Info}(\Theta|\alpha, \lambda_{\max}^*)}{\int f(x)\text{Info}(\Theta|\alpha, \lambda_{\max}^*)d\Theta}.$$

To form an estimate for $\lambda_{\max}$, we slowly increase $\lambda_{\max}^*$ until the posterior representation does not appreciably differ from the prior representation. The smallest value of $\lambda_{\max}^*$ for which this condition holds will serve as the estimate for $\lambda_{\max}$.

It is worth mentioning that there are other, more principled ways of estimating $\lambda_{\max}$. For example, one could build the parameter directly into the model and estimate it on the group level in a Bayesian hierarchy. However, we have found that the quality of the model fit is fairly robust to the selection of $\lambda_{\max}$ as a result of the shrinkage that occurs across subjects due to

990 the constraints imposed by the hierarchical model, and so more complicated
991 estimation techniques may prove unnecessary.

992 5.3.2. *The Mapping Function*

993 We next define a mapping function to relate the prior consistency vari-
994 able ψ to the influence of the anchor parameter λ . In theory, this function
995 may be very complex, requiring many parameters to relate the two variables
996 precisely for all participants. However, in practice, this function should posit
997 psychologically meaningful parameters and will likely require some simplify-
998 ing assumptions to facilitate the interpretability of these parameters.

999 For simplicity, AIM assumes that as the distance between the prior repre-
1000 sentation and the anchor increases, the anchor provides a marginally decreas-
1001 ing influence on the posterior representation, producing an inverted U-shaped
1002 relationship between anchor extremity and the magnitude of the shift from
1003 prior to posterior representation. This assumption is informed by the in-
1004 verted U-shaped relationship observed by Wegener et al. (2001). That said,
1005 there are potential special cases where this assumption may not hold. First,
1006 if an individual uses an anchor as a hint (Schwarz, 1994; Wegener et al., 2010)
1007 it is feasible to assume that the hint is an exaggeration. If a speaker states
1008 that it was “1,000,000 degrees outside,” the listener may infer that it is “very
1009 hot” outside, as opposed to it actually being “1,000,000 degrees.” Thus, the
1010 anchor would act as a more plausible exemplar (e.g., 125 degrees). Other
1011 anchors (e.g., 500,000 and 1,500,000 degrees) would produce similar effects,
1012 resulting in an asymptoting, rather than inverted-U, relationship between
1013 anchor extremity and the size of anchoring effects (Mussweiler and Strack,
1014 2001). AIM can capture this asymptoting behavior through the parameter
1015 λ_{max} . Another special case could occur when an individual has no prior in-
1016 formation about a target item. If there is no prior representation, then the
1017 posterior representation will be made up entirely of the anchor information.
1018 In such a case, the posterior representation will be linearly related to the
1019 anchor.

1020 For our purposes, we will use the flexible Bézier curve. The Bézier curve
1021 uses a set of control points to influence the path a curve may take. It requires
1022 the specification of at least two endpoints that the curve must pass through,
1023 and then any intermediate points can influence the shape of the curve, but
1024 the curve will not pass through these points – except in trivial cases. These
1025 requirements turn out to be helpful in our case, because we can directly build
1026 in restrictions on λ and ψ in the specification of the control points.

1027 While there are many forms of Bézier curves, we focus on the set of ratio-
 1028 nal Bézier curves that add adjustable weights to a set of polynomials. The
 1029 weights in the specification of the Bézier curve will serve as the parameters in
 1030 the mapping function. The mapping function will relate the prior consistency
 1031 variable ψ (forming the x -axis) to $B(\psi)$, the proportion of $\lambda_{\max}$ (forming the
 1032 y -axis), so that $\lambda = B(\psi)\lambda_{\max}$. Thus, the mapping function is defined on the
 1033 unit square that ranges from zero to one on both the x and y axes.

1034 We begin by specifying both the number $n + 1$ and locations P of the
 1035 control points. As the number of control points increases, the Bézier curve
 1036 becomes increasingly more flexible, but the contribution of the individual
 1037 points becomes more difficult to interpret. Here, we will use only four control
 1038 points ($n = 3$) at the vertices of the unit square so that

$$P = \begin{bmatrix} 0 & 1 \\ 0 & 0 \\ 1 & 0 \\ 1 & 1 \end{bmatrix},$$

1039 where the i th row ($i = \{0, 1, 2, 3\}$) of P (denoted P_i) represents the i th (x, y)
 1040 pair. The order of the control points affects the shape and orientation of the
 1041 curve. For our mapping function, we require that the curve passes through
 1042 the points $P_0 = (0, 1)$ and $P_3 = (1, 1)$, so that the curve transitions downward
 1043 toward the point $(0.5, 0)$ depending on weights given to the intermediate
 1044 control points P_1 and P_2 .

1045 Each control point P_i must also have an associated weight w_i . We set the
 1046 endpoints w_0 and w_3 equal to one and set the intermediate points w_1 and
 1047 w_2 equal to ξ so that $w = \{1, \xi, \xi, 1\}$. Constraining $w_1 = w_2$ does suggest
 1048 that participants are unbiased with respect to the direction an anchor is
 1049 presented. That is, anchors presented below the prior representation such
 1050 that $\psi = \psi_1$ have just as much influence on the posterior representation as
 1051 the presentation of an anchor that is above the prior representation such that
 1052 $\psi = 1 - \psi_1$. We set $w_1 = w_2 = \xi$, although we acknowledge that this may be
 1053 an overly restrictive assumption.⁶

1054 Once the set of control points and associated weights have been specified,

⁶Although we believe that high and low anchors have comparable influences, we acknowledge that there are often additional constraints (e.g., zero quantities) imposed on low anchors.

1055 the mapping function is calculated by the equation

$$B(\psi|P, w) = \frac{\sum_{i=0}^n b_{i,n}(\psi) P_i w_i}{\sum_{i=0}^n b_{i,n}(\psi) w_i}, \quad (7)$$

1056 where $b_{i,n}(\psi)$ denotes a set of polynomials, known as the Bernstein basis
1057 polynomials, which are given by

$$b_{i,n}(\psi) = \binom{n}{i} \psi^i (1 - \psi)^{n-i}.$$

1058 Figure 9 illustrates the mapping function for five different values of ξ : 0,
1059 0.2, 1, 2, and 30. The y -axis shows the evaluation of $B(\psi|P, w)$ across the
1060 range of ψ . The relationship between the prior consistency variable ψ and
1061 the anchor influence parameter λ is given by the equation

$$\lambda = \lambda_{\max} B(\psi|P, w). \quad (8)$$

1062 As ψ increases from 0 to 0.5, the value of λ becomes small, and as ψ in-
1063 creases from 0.5 to 1, the value of λ increases again. Figure 9 shows that
1064 as ξ increases from 0 to 30, λ decreases sharply, resulting in a stronger de-
1065 pendency on the information provided by the anchor. Thus, the influence of
1066 ξ becomes essential for understanding how consistency might play a role in
1067 a participant’s judgment following the presentation of an anchor. For some
1068 participants, the presentation of an anchor may have no influence on their
1069 representation of the true value for a target item Θ . Given a prior represen-
1070 tation, these participants could be explained by setting $\xi = 0$, which would
1071 result in an anchor information function that does not appreciably change the
1072 prior representation. By contrast, individuals who are especially susceptible
1073 to the information contained in an anchor can be explained by increasing ξ
1074 to some large value (e.g., $\xi = 2$), which will result in a large shift from the
1075 prior representation to the posterior representation. In this way, ξ represents
1076 the degree of anchor susceptibility that can be compared across participants
1077 once the model has been fit to data.

1078 5.4. Summary

1079 In short, we apply two constraints to extend our within-participant ver-
1080 sion of AIM First, we assume that confidence is related to the prior uncer-
1081 tainty parameter. For example, in Experiment 2 we will assume that across a

Figure 9: The mapping function relating the prior consistency variable ψ to the anchor influence parameter λ at five different values for ξ .

single continuum Θ , confidence judgments behave similarly across individuals. This allows the model to infer how confidence may relate to uncertainty across both items and participants. Second, we use a measure of prior consistency to relate the influence of an anchor to a participant’s prior information. In sum, as an anchor becomes more inconsistent with a participant’s information (i.e., becomes more implausible), our AIM assumes that the anchor will have a proportionally smaller influence on that participant’s final judgment about a target item.

Figure 8 illustrates how the model works for three hypothetical anchors. The top left panel shows the prior representation with $\theta = 0$ and $\delta = 1$ along with the three anchors $\alpha = \{-3, -0.5, -1\}$ represented as vertical lines. After the presentation of these anchors, the prior consistency is calculated and this value is then converted to the anchor influence parameter. The top right panel shows the conversion of the three anchor values to $\psi = \{0.001, 0.309, 0.841\}$ (x -axis) along with the corresponding $B(\psi|P, w) = \{0.996, 0.360, 0.600\}$ (y -axis) where $\xi = 1$ and $\lambda_{\max} = 1$. The values of $B(\psi|P, w)$ are then used to compute λ by Equation 8, which then creates the information functions corresponding to each of the anchor values shown in the middle panel. Finally, the information functions are combined with the prior representations to form the posterior representations, which are shown in the bottom panel.

6. Experiment 2

The primary goal of Experiment 2 was to test AIM’s ability to account for empirical data on an individual participant basis. Experiment 2 employed a within-participant anchoring design to investigate how individuals integrate prior information and information provided by the anchor in the formation of a posterior judgment. Employing a within-participant design will provide the conditions necessary to fully identify the model we presented above.

6.1. Methods

Participants were recruited from Amazon’s Mechanical Turk website and completed the task in exchange for \$0.50. The sample ($N=189$) was 52% female with a mean age of 35 years. Participants completed a similar task to the regular anchors condition of Experiment 1. For each of the six questions from Experiment 1, the task began by eliciting an unanchored (i.e., prior) judgment about the target item and corresponding confidence rating for this

1117 first judgment, from 1 (*Very Unconfident*) to 13 (*Very Confident*). After
 1118 completing this first task, participants completed roughly 15 minutes worth
 1119 of unrelated non-numeric (to avoid the possibility of incidental anchors) tasks
 1120 aimed at decreasing the accessibility of participants’ unanchored judgments
 1121 within their memories. Subsequently, participants were presented the anchor
 1122 in a comparison question (e.g., is the length of the Mississippi River greater
 1123 than or less than 3,500 miles?) before estimating the target and providing a
 1124 confidence judgment the second time for each of the six items.

1125 6.2. Technical Details

1126 In contrast to Experiment 1, in Experiment 2 participants elicited an
 1127 estimate from their prior representation and their posterior representation
 1128 for each of the six questions. Thus, we are able to estimate parameters in
 1129 the model associated with the prior representation (θ and δ) as well as the
 1130 posterior representation (λ). In addition, we are now able to use a model
 1131 that incorporates individual differences among all three parameters of AIM
 1132 θ , δ , and λ . To do so, we will use a hierarchical version of AIM which will
 1133 again be fit to the data in the Bayesian framework. The technical details of
 1134 the hierarchical Bayesian model are presented in Appendix B, and we have
 1135 provided JAGS (Plummer, 2003) code for fitting AIM to data in Appendix
 1136 D.

1137 To simplify the notation, we will again divide the data up into data from
 1138 the k th question from the prior question x_k (i.e., the presentation of Ques-
 1139 tion k that does not contain an anchor) and the k th anchoring question y_k .
 1140 To indicate individual responses, we add the subscript j to the appropriate
 1141 elements. For example, $x_{k,j}$ denotes the response from the j th subject in the
 1142 k th prior anchoring question. To control for the large skew in the estimates,
 1143 we performed a log transformation of the raw data, making the support for
 1144 each Θ_k stretch from negative infinity to infinity. We can assume a normal
 1145 prior representation for each subject in each anchoring question so that

$$\begin{aligned} x_{k,j} &\sim \text{Prior}(\theta_{k,j}, \delta_{k,j}) \\ &\sim \mathcal{N}(\theta_{k,j}, \delta_{k,j}). \end{aligned} \tag{9}$$

1146 To further constrain the model we built the confidence judgments directly
 1147 into the estimates for each of the $\delta_{k,j}$ s, as discussed above. We first collapsed
 1148 nearby confidence judgments into new bins, reducing the number of possible
 1149 confidence responses from 13 to 6. For example, confidence responses of “1”

1150 and “2” were reassigned to the judgment “1” and “3” and “4” were reassigned
 1151 to the judgment “2”, and so on. The final bin, “6”, contained judgments of
 1152 “11”, “12”, and “13”.

1153 We denote the confidence response from the j th subject in the k th prior
 1154 anchoring question as $z_{k,j}$ which ranges from $\{1, 2, \dots, 6\}$. We group all sub-
 1155 jects eliciting the same confidence judgment together by assuming that they
 1156 have the same degree of prior uncertainty, although they may have different
 1157 mean parameters $\theta_{k,j}$. Thus, for each confidence judgment $c \in \{1, 2, \dots, 6\}$
 1158 in the k th question, we assign a separate prior uncertainty parameter $\beta_{k,c}$,
 1159 so that

$$\delta_{k,j} = \beta_{k,z_{k,j}}.$$

1160 To avoid any undue constraint on the estimates for $\beta_{k,c}$, we did not specify any
 1161 hierarchical structure such as the ones described above; instead, we allowed
 1162 each $\beta_{k,c}$ to vary freely.

1163 As we explained above, AIM assumes that an individual evaluates how
 1164 consistent the anchor is with the prior beliefs. We assume that upon the
 1165 presentation of an anchor α_k in the k th question, individuals first determine
 1166 a consistency value $\psi_{k,j}$ by integrating their prior representation up to the
 1167 given anchor value, or

$$\psi_{k,j} = \int_{-\infty}^{\alpha_k} \text{Prior}(\Theta_k | \theta_{k,j}, \delta_{k,j}) d\Theta_k.$$

1168 Once computed, this degree of consistency measure $\psi_{k,j}$ is then used to spec-
 1169 ify the anchor influence parameter $\lambda_{k,j}$ by evaluating

$$\lambda_{k,j} = B(\psi_{k,j} | P, w) \lambda_{\max}^k, \quad (10)$$

1170 where $w = \{1, \xi_j, \xi_j, 1\}$, $B(\psi_{k,j} | P, w)$ is evaluated by Equation 7, and $\lambda_{\max}^k$ is
 1171 the upper bound constraint on λ in the k th question. We let ξ vary across
 1172 participants and assume that these parameters follow a common distribution
 1173 such that

$$\xi_j \sim \mathcal{TN}(\xi^{(\mu)}, \xi^{(\sigma)}, 0, \infty), \quad (11)$$

1174 where $\mathcal{TN}(a, b, c, d)$ represents a truncated normal distribution with mean
 1175 a , standard deviation b , lower bound c and upper bound d .

1176 The anchor influence parameter is then used to form the information
 1177 function. Here, we assume the information function is normal in form, so

1178 that

$$\text{Info}(\Theta_k|\alpha_k, \lambda_{k,j}) = \frac{1}{\sqrt{2\pi}\lambda_{k,j}} \exp \left[-\frac{(\Theta_k - \alpha_k)^2}{2\lambda_{k,j}^2} \right]. \quad (12)$$

1179 Once the anchor provides information about Θ_k , the participant updates their
 1180 prior representation by the information function to arrive at the posterior
 1181 representation

$$\text{Post}(\Theta_k|\theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k) = \frac{\text{Prior}(\Theta_k|\theta_{k,j}, \delta_{k,j})\text{Info}(\Theta_k|\alpha_k, \lambda_{k,j})}{\int \text{Prior}(\Theta_k|\theta_{k,j}, \delta_{k,j})\text{Info}(\Theta_k|\alpha_k, \lambda_{k,j})d\Theta_k}.$$

1182 Similar to the earlier version of our model used in Experiment 1 where we
 1183 assumed a normal prior representation and a normal information function,
 1184 we can combine Equations 9 and 12 to show that the posterior representation
 1185 for these data also has a closed form expression. Specifically,

$$\text{Post}(\Theta_k|\theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k) = \frac{1}{\sqrt{2\pi}b_{k,j}} \exp \left[-\frac{(\Theta_k - a_{k,j})^2}{2b_{k,j}^2} \right],$$

1186 which is a normal distribution with mean $a_{k,j}$ and standard deviation $b_{k,j}$,
 1187 where

$$\begin{aligned} a_{k,j} &= \frac{\frac{\theta_{k,j}}{\delta_{k,j}^2} + \frac{\alpha_k}{\lambda_{k,j}^2}}{\frac{1}{\delta_{k,j}^2} + \frac{1}{\lambda_{k,j}^2}}, \text{ and} \\ b_{k,j} &= \left(\frac{1}{\delta_{k,j}^2} + \frac{1}{\lambda_{k,j}^2} \right)^{-1/2}. \end{aligned} \quad (13)$$

1188 Thus, the j th response from the k th anchoring question is assumed to arise
 1189 from the distribution given by

$$\begin{aligned} y_{k,j} &\sim \text{Post}(\theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k) \\ &\sim \mathcal{N}(a_{k,j}, b_{k,j}) \end{aligned}$$

1190 Figure 10 shows a graphical diagram for this model. Nodes with shaded
 1191 rings are the model components from the previous version of the model. As an
 1192 extension of the earlier version of our model, we have added constraints from
 1193 the confidence judgments as well as employed the prior consistency variable

1194 ψ to identify the model and dissociate the separate roles of prior uncertainty
 1195 and the influence of the anchor values. Notice that for this model, in addition
 1196 to a plate for each of the questions, we also have a plate for each of the
 1197 participants and a plate for each of the confidence categories. At the first
 1198 layer of the model, we have all of the participant parameters, namely, any
 1199 of the parameters located on the “participant” plate. For example, the node
 1200 for ξ_j is only on the participant plate, making it intrinsic to each participant
 1201 but fixed across all questions. Similarly, the variable $\beta_{k,c}$ changes across
 1202 each confidence category *and* anchoring question, but is held fixed across
 1203 all participants. To facilitate a comparison of this extended model to the
 1204 basic model used for Experiment 1 (see Figure 2), we added light gray circles
 1205 around the nodes that are consistent with the previous model.

1206 6.3. Results

1207 We fit the model in a Bayesian framework by again using DE-MCMC
 1208 sampling (see ter Braak, 2006; Turner et al., 2013, for details). We ran a
 1209 blocked version of the algorithm with 40 chains distributed across 8 proces-
 1210 sors for 3,000 iterations after a burnin period of 4,000 iterations, resulting in
 1211 120,000 total samples from the joint posterior distribution. Standard proce-
 1212 dures were used to ensure convergence.

1213 There are a number of parameter analyses and posterior checks we could
 1214 present here; however, for brevity we will instead focus on two sets of param-
 1215 eters that are central to the development of our model: the prior uncertainty
 1216 parameters β and the anchor sensitivity parameters ξ . Similar to the analysis
 1217 in Experiment 1, we will then examine the posterior predictive distributions
 1218 of the model against the data that were observed.

1219 We begin by assessing the relationship our model parameters have to the
 1220 degree of confidence reported for each individual judgment. In the model sec-
 1221 tion, we hypothesized that as the reported confidence increased, we should
 1222 see a much smaller value for the prior uncertainty parameter. Recall that
 1223 the model for these data assumed a separate prior uncertainty parameter for
 1224 each question by confidence judgment, and assumed no relationship between
 1225 the confidence parameters. However, to apply constraints to the model at
 1226 the individual level, we assumed that participants in the k th question who
 1227 reported a confidence judgment of c shared the same prior uncertainty pa-
 1228 rameter $\beta_{k,c}$, but may not necessarily share the same prior mean $\theta_{k,j}$.

1229 Figure 11 shows the estimated posterior distribution for each parameter
 1230 $\beta_{k,c}$ on the log scale. The separate panels represent the estimates for each of

Figure 10: A graphical diagram of the within-participant version of AIM.

Figure 11: The estimated posterior distributions for β (y -axes) on the log scale for each of the six questions as a function of confidence response (x -axes).

the six questions (i.e., the rows of the matrix β), and the separate violin plots within a panel represent the elicited confidence judgment (i.e., the columns of the matrix β). There is a strong negative relationship between prior uncertainty and the confidence response. Specifically, Figure 11 shows that as the reported confidence judgment increases from 1 to 6, the degree of uncertainty in the prior representations for this group of participants decreases, reflecting a more peaked prior representation (see Figure 1). In general, we find that as confidence increases, the difference between the prior judgment and the posterior judgment decreases, and our estimates of $\beta_{k,c}$ reflect this relationship.

One compelling advantage of the Bayesian approach is the quantification of uncertainty in each parameter estimate. For example, the top left panel of Figure 11 shows that the estimate for $\beta_{1,6}$ (the violin plot furthest to the right) is much more variable than the estimate for say, $\beta_{1,1}$ (the violin plot furthest to the left). Thus, while we can be sure that the true value for $\log(\beta_{1,1})$ is somewhere between -1 and 1, the true value for $\log(\beta_{1,6})$ might range anywhere from -1 to 3. The uncertainty here is largely due to the sparsity in the number of confidence judgments of “6”. However, other factors, such as the variability in the estimate itself, will influence the shape of these posterior distributions.

The next parameter we investigate is the anchor sensitivity parameter ξ . The parameter ξ was constructed to provide some insight into how an individual might update for any hypothetical anchor by relating the measure of prior consistency (see Equation 6) to the spread of the information function. The parameter ξ controls this mapping so that as ξ increases, anchors are perceived to carry more information than when ξ is small (see Figure 9). Thus, participants whose posterior judgments differed considerably from their prior judgments should have large values of ξ .

To construct a holistic measure of how much a participant’s prior and posterior judgments differ, we first calculate the average absolute difference (i.e., their average updating score), given by

$$MA_j = \frac{1}{6} \sum_{k=1}^6 |x_{k,j} - y_{k,j}|.$$

Although this measure will not provide us with insight about *why* a participant might update their estimates, it provides a measure that we can compare ξ against. Figure 12 shows the average absolute differences MA against the

Figure 12: The estimated posterior distribution for ξ as a function of the mean updating across all questions for each subject. The median of the estimated posterior is shown as the black dot and the 95% credible set is shown as the height of the vertical gray line.

posterior estimates for ξ for each participant. Each ξ_j is shown as the median, represented as the black dot, and the 95% credible set, represented as the length of the vertical lines. The figure shows a positive relationship between MA and ξ : as ξ increases, the average absolute differences also tends to increase. The variable MA is correlated with the parameter ξ , resulting in a correlation of $R = 0.27$. This result indicates that ξ is related to the degree that individuals' update their beliefs when presented an anchor (i.e., their susceptibility to the anchor).

While the variable MA does provide a reasonable statistic for the average degree of movement for a participant, it does not provide any information about the reason for the updating. AIM, on the other hand, has two mechanisms to explain this updating. The first rests in the degree of prior uncertainty, which we examined above. The second is anchor susceptibility, which we have just discussed. We have shown that together these estimates provide a reasonable interpretation of why participants update their prior judgments after the presentation of an anchor. In particular, the range of the estimates for ξ combined with the relationship of prior uncertainty and elicited confidence suggest that prior uncertainty is a larger contributor to posterior updating than anchor susceptibility.

As a final assessment of model fit, we compared the posterior predictive distribution of our model to the data that were observed for each anchoring question. Figure 13 shows these posterior predictive distributions (gray clouds) along with the data (black dots) for each question (represented as separate panels). For each question, we first plotted the prior judgment $x_{k,j}$ against $y_{k,j}$ observed in the data. After fitting the model to these data, we then generated 500 samples from the posterior predictive distribution.⁷ Figure 13 shows that AIM does an excellent job of matching both the spread and relative density of the observed data.

The vertical and horizontal dashed lines in each panel shows the anchor value that was presented. We also shaded the regions that correspond to an updating from the prior representation toward the presented anchor value. For example, in the top left panel of Figure 13 the bottom left region is shaded to indicate that if a prior judgment of say, 5.1 increased toward

⁷While 500 samples may seem like a small number, it was chosen for graphical purposes. Increasing the number of samples produced a dense cloud of points making it difficult to assess the density at any particular location.

Figure 13: The posterior predictive distributions (gray cloud) along with the observed responses (black dots) for the prior judgment (x -axes) and posterior judgment (y -axes) for each of the six anchoring questions in Experiment 2. Both axes are log-transformed. The dashed vertical lines represent the anchor value that was presented. The light gray triangular regions indicate positive updating (i.e., updating from prior estimates toward the anchor value).

the presented anchor value of 8.16 (the horizontal dashed line), it would lie somewhere in the interval (5.1, 8.16). Similarly, the top right triangular area is shaded to reflect the situation where a participant’s prior judgment is above the presented anchor. Following the presentation of 8.16 (i.e., the log-transformed anchor value), if the participant’s posterior judgment shifted downward (and hence would be under the diagonal line $y = x$) toward the anchor value, their data point would fall in the gray triangular region.

Figure 13 shows that the distribution of responses as a function of distance from the anchor is curvilinear: there tends to be small updates in the posterior estimates when the prior estimates are nearby or far from the anchor, but large updates (i.e., anchoring effects) in between these extremes. This curvilinear pattern is consistent with the inverted U-shaped relationship observed by Wegener et al. (2001).

The figure also shows that AIM can account for reversals, where the judgment from prior to posterior moves in the opposite direction of the anchor value or moves to the other side of the anchoring value. Recall that in AIM, because the response is assumed to be sampled from a representation (either the prior or posterior), there is some probability of a reverse anchoring effect, where the judgment moves in the opposite direction of the anchor. Although this may seem unlikely, such outcomes are evident in each of the questions in Figure 13 (i.e., when a black dot falls outside of the shaded regions). In addition, AIM naturally captures the effect that the closer a prior judgment is to a presented anchor value, the smaller the difference between the prior and posterior judgments. For example, the top middle panel shows that in areas far from the anchor value, the variability in the joint distribution of the prior and posterior representations is higher than in regions near the presented anchor (i.e., regions near the intersection of the vertical and horizontal dashed lines). This pattern is also exhibited in the behavioral data.

6.4. Discussion

Experiment 2 was intended to test the adequacy of the within-participant version of AIM, assessing the extent to which individuals integrate anchor information into their prior representation to form a posterior representation of the target item. Although it is not feasible to actually measure both an individual’s prior and posterior representation, Experiment 2 provides a novel insight into this information integration. To the best of our knowledge, this is the first demonstration of the standard anchoring effect at the individual level. Although participants’ posterior estimates were heavily biased by

their prior estimates, the current results do show anchoring effects at the individual level. Furthermore, the within-participant version of AIM provided impressive fits to the data at the individual level.

7. General Discussion

Four decades of anchoring research has demonstrated the stability of anchoring effects (Mussweiler, 2001), and the importance of this class of effects on everyday judgments and decisions (Epley, 2004; Furnham and Boo, 2011). The current article presented two versions of AIM, a descriptive model for quantifying anchoring effects.

The anchoring-effects literature is composed of five non-mutually exclusive theories, each of which provide useful insight toward understanding the processes underlying anchoring effects. We do not discount the existing literature, but believe that the literature needs trimming to improve its usability, both for theorists and practitioners. Because existing anchoring theories differ in terms of how anchors provide information, we suggest that theorists could build process-based models of their proposed psychological mechanisms within the AIM framework and test for the model's adequacy in accounting for the data. For example, Bhatia and Chaudhry (2013) introduced an associative network model of how exemplars are retrieved in a manner consistent with the selective accessibility hypothesis Strack and Mussweiler (1997); Mussweiler and Strack (1999). A version of this model could be fit within AIM to determine the anchor-influence parameter λ . In doing so, AIM would become a process-based model that allows for testing the adequacy of the selective accessibility hypothesis. Because AIM – as a descriptive model – is agnostic about how the anchor provides information, AIM would be a valuable tool for discriminating the efficacy of differing accounts. At the very least, this approach may not rule out any of the psychological mechanisms proposed within the literature, but could provide evidence as to the prevalence of any of the proposed psychological processes across a range of paradigms within the anchoring literature. We do not discount the existing qualitative approaches for conceptualizing anchoring effects, rather we view AIM as a useful tool for further developing existing accounts. Thus, although AIM is descriptive in its current form, AIM provides a novel means for testing and differentiating theories.

1369 7.1. Limitations

1370 A major limitation of the current research is that it is difficult to obtain
 1371 enough data to appropriately identify AIM. For the between-participant ver-
 1372 sion of AIM, we assessed the unanchored judgments of a control group and
 1373 used the distribution of responses as an approximation for participants' av-
 1374 erage prior representations in the experimental conditions of Experiment 1.
 1375 Experiment 2 attempted to remove this assumption by collecting evidence for
 1376 both a participant's prior and posterior representation of each target judg-
 1377 ment. However, once an individual is asked to estimate a target item, his or
 1378 her estimate of the target item will act as a self-generated anchor, making it
 1379 difficult to differentiate between the effects of the self-generated anchor and
 1380 the anchor value presented in the subsequent question. Furthermore, it seems
 1381 unreasonable to expect that a participant will respond differently to the sec-
 1382 ond presentation of a question when it is essentially the same as the first
 1383 presentation. Regardless, we attempted a within-participant manipulation
 1384 in Experiment 2, in which each participant was asked to report a value from
 1385 their prior representation, and then report an estimate from their posterior
 1386 representation, after the presentation of an anchor. As predicted, some par-
 1387 ticipants did respond similarly (some even provided the same response) for
 1388 each question. After fitting AIM to the data, we found that our model could
 1389 account for these data, despite the large dependencies in participants' re-
 1390 sponses. AIM predicted a degree of variability that was consistent with each
 1391 participant's responses. Although we are skeptical about the effectiveness
 1392 of within-participant manipulations, our data exhibited moderate anchoring
 1393 effects across participants, and AIM was successful in accounting for these
 1394 data.

1395 One limitation of AIM is the restriction of the form of both the prior
 1396 representation and the anchor information function. We have already ex-
 1397 pressed concerns about these assumptions, but we reemphasize that more
 1398 complicated priors may be necessary in certain instances (e.g., Lewandowsky
 1399 et al., 2009). When this is the case, one could employ a variety of Bayesian
 1400 techniques such as finite or infinite mixture priors, or empirical priors such
 1401 as a kernel density estimate of the data from the prior representation (see
 1402 Silverman, 1986).

1403 In the within-participant version of AIM, one potential limitation lies
 1404 in the mapping function. In this article, we proposed a mapping function
 1405 that was very simple, symmetric, and provided a convenient interpretation
 1406 for the parameter ξ . However, we suspect that this may be one component

of the model that could be improved. Furthermore, the mapping function's dependence on the upper bound parameter $\lambda_{\max}$ also poses a slight complication. In this article, we simply proposed that $\lambda_{\max}$ be estimated prior to fitting AIM to the data. Although we have not tested this directly, one could incorporate $\lambda_{\max}$ into the model and estimate it.

7.2. Future Directions

We have demonstrated that AIM can account for a variety of anchoring effects. Due to the information integration proposed in this article, we would also expect that AIM could be extended to account for dynamic judgments as well. Dynamic judgments are those in which individuals incorporate several sources of information and make a series of estimates. For example, an extension of AIM could account for and reliably predict the learning processes in stock trading, gambling, and related dynamic environments.

Although we consider AIM to be a step forward in modeling anchoring effects, we hope that AIM spurs additional research into the development of methods that more appropriately collect within-participant data, and appropriate models for fitting the data. In addition to traditional methods for studying anchoring effects, we believe that future anchoring research should concern within-participant data. Such data would allow for greater insight into the psychological processes underlying anchoring effects.

7.2.1. AIM as a Predictive Model for “Nudging”

Although forty years of research have already been dedicated to the anchoring phenomenon, the current article opens up several novel avenues for additional research. For instance, as noted above, AIM provides a framework for future theory evaluation. One of the most substantial opportunities involves extending AIM from a descriptive model into a predictive model. In its current form AIM provides a useful tool for quantifying the observed relationships present in the data. We use the data from the current experiments to estimate the anchor-influence parameter λ . If λ were known, then for any given prior, AIM could estimate predictions for the posterior representation corresponding to any presented anchor. If prior and posterior data were collected across a range of anchor values within a given paradigm, then AIM could infer λ corresponding to any anchor value and predict, a priori, the posterior representation for any anchor value. For instance, past research demonstrated that the listing price of a home acts as an anchor on subsequent home valuations (Northcraft and Neale, 1987). For a given market,

incorporating existing priors (e.g., past listing prices) and posteriors (i.e., corresponding offers), AIM could predict a posterior for any given anchor. We find that the parsimony of AIM allows it to be tractable in a myriad of domains.

The ability to predict, not just the existence of anchoring effects, but a priori estimates of the magnitude of anchoring effects would provide a valuable tool. Considerable research on “nudging” (also referred to as choice architecture) has highlighted the utility of behavioral insights for public policy (Allcott and Mullainathan, 2010; Johnson and Goldstein, 2003; Thaler and Sunstein, 2008). By using behavioral interventions, researchers and policy developers can construct decision environments that facilitate effective judgments and decisions. Given the reliability and prevalence of anchoring effects across a range of important domains, AIM may provide an important tool for policy. In its current form, the anchoring literature can inform policy about whether or not an anchoring effect will occur in a particular paradigm. Furthermore, AIM offers the ability for policy developers to potentially choose “optimal” anchors.

8. Conclusions

The abundance of anchoring research highlights its prevalence and stability as a core feature of human judgment and decision making. Anchoring effects often involve a quantitative anchor and a quantitative estimate, as such qualitative theories of anchoring effects are insufficient for fully articulating anchoring effects. The current article introduces AIM for the quantitative assessment of anchoring effects. AIM not only provides a means for assessing the magnitude of anchoring effects, but the opportunity for theory evaluation and prediction.

1469 In this section, we present the equations necessary to fit the two ver-
 1470 sions of AIM to data. Our exposition is presented deductively, moving from
 1471 more general equations which are applicable to any experimental data to
 1472 the specific ones used to fit our data from Experiment 1 and Experiment 2,
 1473 respectively. All notation, except where noted, is as presented in the model
 1474 sections of Experiments 1 and 2.

1475 **Appendix A. Equations for Fitting the Between-participant Ver-** 1476 **sion of AIM to Experimental Data**

1477 We again present the details of the analysis on the basis of a single item
 1478 with multiple conditions (i.e., anchor types), to simplify the exposition. We
 1479 let the estimate provided by the j th participant in the prior condition be
 1480 denoted x_j . We assume that the J participants have the same prior repre-
 1481 sentation from which they independently sample their estimate, which makes
 1482 the specific contribution of the data x to the parameters θ and δ described
 1483 by the equation

$$f(\theta, \delta | x) = \prod_{j=1}^J \text{Prior}(x_j | \theta, \delta). \quad (\text{A.1})$$

1484 We let $y_{k,j}$ denote the the j th response in the k anchoring condition,
 1485 where $j \in \{1, \dots, M_k\}$ and M_k is the number of responses obtained in the
 1486 k th anchoring condition. For the posterior representation, we again assume
 1487 that the data $y_{k,j}$ are independent and identically distributed. Thus, the
 1488 specific influence of the data from all participants in the k th condition, which
 1489 we denote $y_{k,1:M_k}$, can be described by the equation

$$f(\theta, \delta, \lambda_k | y_{k,1:M_k}, \alpha_k) = \prod_{j=1}^{M_k} \text{Post}(y_{k,j} | \theta, \delta, \lambda_k, \alpha_k), \quad (\text{A.2})$$

1490 where α_k is assumed to be known. If we further assume that the data y are
 1491 independent across the K conditions, then Equation A.2 becomes

$$f(\theta, \delta, \lambda | y, \alpha) = \prod_{k=1}^K \prod_{j=1}^{M_k} \text{Post}(y_{k,j} | \theta, \delta, \lambda_k, \alpha_k). \quad (\text{A.3})$$

1492 We can then combine Equations A.1 and A.3 to form the likelihood function
 1493 for the parameters θ , δ , and λ , conditional on the anchor value α and the

1494 data x and y , such that

$$L(\theta, \delta, \lambda | x, y, \alpha) = \prod_{j=1}^J \text{Prior}(x_j | \theta, \delta) \prod_{k=1}^K \prod_{j=1}^{M_k} \text{Post}(y_{k,j} | \theta, \delta, \lambda_k, \alpha_k). \quad (\text{A.4})$$

1495 If one can assume, as we do to fit the data from Experiment 1, that both the
1496 prior and information functions are normal in form, Equation A.4 becomes

$$L(\theta, \delta, \lambda | x, y, \alpha) = \prod_{j=1}^J \phi(x_j | \theta, \delta) \prod_{k=1}^K \prod_{j=1}^{M_k} \phi(y_{k,j} | a_k, b_k), \quad (\text{A.5})$$

1497 where a_k and b_k are evaluated using Equation 5, and $\phi(x|a, b)$ denotes the
1498 normal density with mean a and standard deviation b , evaluated at the loca-
1499 tion x . With the likelihood function in hand, one can use standard routines
1500 to optimize Equation A.4 – or when appropriate, Equation A.5 – to obtain
1501 maximum likelihood estimates for the parameters θ , δ , and λ .

1502 When estimating the parameters θ , δ and λ in the Bayesian framework,
1503 one must provide prior distributions for each parameter. To fit the data from
1504 Experiment 1, we assumed a normal prior for θ and gamma priors for both
1505 δ and λ such that

$$\begin{aligned} \theta &\sim \mathcal{N}(\mu, 10), \\ \delta &\sim \Gamma(1, 1), \text{ and} \\ \lambda_k &\sim \Gamma(1, 1). \end{aligned}$$

1506 The values for each hyperparameter μ were selected somewhat informatively
1507 for each item such that $\mu = \{7, 3, 3, 4.5, 1, 5\}$.

1508 Once the priors for θ , δ and λ have been specified, we can now write the
1509 joint posterior distribution as

$$\pi(\theta, \delta, \lambda | x, y, \alpha) \propto L(\theta, \delta, \lambda | x, y, \alpha) \prod_{k=1}^K \pi(\lambda_k) \pi(\theta) \pi(\delta), \quad (\text{A.6})$$

1510 where the likelihood is provided in Equation A.4. Because Equation A.6 is
1511 rarely in a convenient form, we generally must use algorithms to sample from
1512 it. In fitting the data from Experiment 1, we obtained samples directly from
1513 Equation A.6, but one could employ a Gibbs sampling algorithm by sampling
1514 from the conditional distributions for each of the parameters sequentially.

1515 A Gibbs sampling approach would be particularly advantageous when the
 1516 dimensionality of λ is large, because of the decline in the acceptance rate
 1517 of proposal generation in high-dimensional spaces (see Robert and Casella,
 1518 2004; Gelman et al., 2004). The conditional distributions for θ , δ and λ_k are

$$\begin{aligned}\pi(\theta|\delta, \lambda, x, y, \alpha) &\propto L(\theta, \delta, \lambda|x, y, \alpha)\pi(\theta), \\ \pi(\delta|\theta, \lambda, x, y, \alpha) &\propto L(\theta, \delta, \lambda|x, y, \alpha)\pi(\delta), \text{ and} \\ \pi(\lambda_k|\theta, \delta, x, y, \alpha_k) &\propto L(\theta, \delta, \lambda_k|x, y, \alpha_k)\pi(\lambda_k) \\ &\propto \prod_{j=1}^{M_k} \text{Post}(y_{k,j}|\theta, \delta, \lambda_k, \alpha_k)\pi(\lambda_k),\end{aligned}$$

1519 respectively.

1520 Appendix B. Equations for Fitting the Within-participant Version 1521 of AIM to Experimental Data

1522 Equation 9 specifies that the data from the prior condition, denoted $x_{k,j}$
1523 for the j th participant in the k th question, comes from the prior represen-
1524 tation governed by the parameters $\theta_{k,j}$ and $\delta_{k,j} = \beta_{k,z_{k,j}}$, where $z_{k,j}$ is the
1525 corresponding confidence judgment for $x_{k,j}$.⁸ Thus, the influence of the data
1526 $x_{k,j}$ on the parameters $\theta_{k,j}$ and $\delta_{k,j}$ can be described by the equation

$$f(\theta_{k,j}, \delta_{k,j} | x_{k,j}) = \text{Prior}(x_{k,j} | \theta_{k,j}, \delta_{k,j}). \quad (\text{B.1})$$

1527 For the posterior judgment $y_{k,j}$, we can write the influence of the data on the
1528 parameters $\theta_{k,j}$, $\delta_{k,j}$ and $\lambda_{k,j}$ as

$$f(\theta_{k,j}, \delta_{k,j}, \lambda_{k,j} | y_{k,j}, \alpha_k) = \text{Post}(y_{k,j} | \theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k), \quad (\text{B.2})$$

1529 where α_k is assumed to be known and $\lambda_{k,j}$ is determined by Equation 10.

1530 Combining Equations B.1 and B.2, we can write the likelihood function
1531 as

$$\begin{aligned} L(\theta_{k,j}, \delta_{k,j}, \lambda_{k,j} | x_{k,j}, y_{k,j}, \alpha_k) &= \text{Prior}(x_{k,j} | \theta_{k,j}, \delta_{k,j}) \\ &\quad \times \text{Post}(y_{k,j} | \theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k). \end{aligned} \quad (\text{B.3})$$

1532 By assuming that the data $x_{k,j}$ and $y_{k,j}$ are independent and identically dis-
1533 tributed from the data set x and y , respectively, we can express the likelihood
1534 function in Equation B.3 for the parameter set θ , δ , and λ as

$$L(\theta, \delta, \lambda | x, y, \alpha) = \prod_{k=1}^K \prod_{j=1}^J \text{Prior}(x_{k,j} | \theta_{k,j}, \delta_{k,j}) \text{Post}(y_{k,j} | \theta_{k,j}, \delta_{k,j}, \lambda_{k,j}, \alpha_k), \quad (\text{B.4})$$

1535 where we assume that x and y are of the same dimensionality (namely, they
1536 are matrices of dimension $K \times J$). Because the free parameters in the model
1537 are actually θ , ξ , and β , it may be convenient to rewrite Equation B.4 in
1538 terms of these parameters, which is given by

$$\begin{aligned} L(\theta, \beta, \xi | x, y, z, \alpha) &= \prod_{k=1}^K \prod_{j=1}^J \text{Prior}(x_{k,j} | \theta_{k,j}, \beta_{k,z_{k,j}}) \\ &\quad \times \text{Post}(y_{k,j} | \theta_{k,j}, \beta_{k,z_{k,j}}, q_{k,j}, \alpha_k), \end{aligned} \quad (\text{B.5})$$

⁸Although several nodes in the model are deterministic, we will refer to them as parameters for simplicity. For example, the node $\delta_{k,j}$ is actually determined by the equation $\delta_{k,j} = \beta_{k,z_{k,j}}$, where $\beta_{k,z_{k,j}}$ is the unknown parameter.

1539 where

$$q_{k,j} = B \left(\int_{-\infty}^{\alpha_k} \text{Prior}(\Theta_k | \theta_{k,j}, \beta_{k,z_{k,j}}) d\Theta_k \middle| P, w = \{1, \xi_j, \xi_j, 1\} \right) \lambda_{\max}^k.$$

1540 Note that this new expression makes the dependency on the observed confi-
1541 dence judgments z explicit.

1542 To fit the model to the data in our Experiment 2, we assumed that the
1543 prior representations and the information functions were normal in form.
1544 Thus, Equation B.4 becomes

$$L(\theta, \delta, \lambda | x, y, \alpha) = \prod_{k=1}^K \prod_{j=1}^J \phi(x_{k,j} | \theta_{k,j}, \delta_{k,j}) \phi(y_{k,j} | a_{k,j}, b_{k,j}), \quad (\text{B.6})$$

1545 where $a_{k,j}$ and $b_{k,j}$ are determined by Equation 13. With the likelihood
1546 function in hand, one can employ standard optimization routines to maximize
1547 Equation B.4 – or Equation B.6 when the appropriate assumptions are made
1548 – to obtain maximum likelihood estimates of the parameters θ , β , and ξ .

1549 When fitting the within-participant version of AIM in the Bayesian frame-
1550 work, one must specify hierarchical structures in the form of prior distribu-
1551 tions as well as specify prior distributions for these hierarchical parameters.
1552 In our model fitting, we did not specify any hierarchical structure for $\beta_{k,c}$,
1553 and instead allowed each $\beta_{k,c}$ to vary freely by specifying a uniform prior so
1554 that

$$\beta_{k,c} \sim CU[0, 20],$$

1555 where $CU[a, b]$ denotes a continuous uniform distribution on the closed in-
1556 terval $[a, b]$. This specification made the prior density for β

$$\pi(\beta) = \prod_{k=1}^K \prod_{c=1}^C \frac{1}{20} = \left(\frac{1}{20} \right)^{KC}$$

1557 By contrast, we constructed a hierarchical structure for the prior mean pa-
1558 rameters by specifying that the set of prior means in the k th question, de-
1559 noted $\theta_{k,1:J}$, come from a common normal distribution such that

$$\theta_{k,1:J} \sim \mathcal{N} \left(\theta_k^{(\mu)}, \theta_k^{(\sigma)} \right).$$

1560 Thus, the prior density of $\theta_{k,1:J}$ is

$$\pi(\theta_{k,1:J}) = \prod_{j=1}^J \phi \left(\theta_{k,j} | \theta_k^{(\mu)}, \theta_k^{(\sigma)} \right),$$

1561 and the prior density of θ is

$$\pi(\theta) = \prod_{k=1}^K \prod_{j=1}^J \phi\left(\theta_{k,j} | \theta_k^{(\mu)}, \theta_k^{(\sigma)}\right),$$

1562 For the anchor sensitivity parameters ξ_j , we assumed a common truncated
1563 normal distribution, given by Equation 11, such that the prior density of ξ
1564 is

$$\pi(\xi) = \prod_{j=1}^J \text{TN}\left(\xi_j | \xi^{(\mu)}, \xi^{(\sigma)}, 0, \infty\right),$$

1565 where $\text{TN}(x|a, b, c, d)$ represents the truncated normal *density* with mean a ,
1566 standard deviation b , lower bound c and upper bound d , evaluated at the
1567 location x .

1568 We must now specify the prior distributions for the hyperparameters. For
1569 the means $\theta_k^{(\mu)}$ and standard deviations $\theta_k^{(\sigma)}$, we assumed mildly informative
1570 priors such that

$$\begin{aligned}\theta_k^{(\mu)} &\sim \mathcal{N}(4, 10) \\ \theta_k^{(\sigma)} &\sim \Gamma(1, 1),\end{aligned}$$

1571 for each $k \in \{1, \dots, 6\}$. For the mean $\xi^{(\mu)}$ and hyper standard deviation $\xi^{(\sigma)}$,
1572 we assumed

$$\begin{aligned}\xi^{(\mu)} &\sim \mathcal{TN}(10, 20, 0, \infty) \\ \xi^{(\sigma)} &\sim \Gamma(1, 1).\end{aligned}$$

1573 Combining the likelihood function in Equation B.5 and each of the priors
1574 listed above, we can write the joint posterior distribution as

$$\begin{aligned}\pi(\eta|x, y, z, \alpha) &\propto L(\theta, \xi, \beta|x, y, z, \alpha) \prod_{k=1}^K \left[\pi\left(\theta_k^{(\mu)}\right) \pi\left(\theta_k^{(\sigma)}\right) \right] \\ &\times \pi\left(\xi^{(\mu)}\right) \pi\left(\xi^{(\sigma)}\right) \pi(\theta) \pi(\xi) \pi(\beta).\end{aligned}\tag{B.7}$$

1575 where we use η to represent the set of parameters of interest, namely $\eta =$
1576 $\{\theta, \xi, \beta, \theta^{(\mu)}, \theta^{(\sigma)}, \xi^{(\mu)}, \xi^{(\sigma)}\}$. Equation B.7 is a general equation for fitting
1577 the model under the constraints we used to fit the model to the data from
1578 Experiment 2. There are numerous alternative modeling approaches, but

1579 the derivation of the joint posterior distribution would be similar to the one
1580 derived here.

1581 To sample from the joint posterior in Equation B.7, we implemented a
1582 blocked version of the DE-MCMC algorithm (see Turner et al., 2013, for de-
1583 tails). While we do not present the conditional distributions here, our blocks
1584 consisted of the pair $(\xi^{(\mu)}, \xi^{(\sigma)})$, each k th pair $(\theta_k^{(\mu)}, \theta_k^{(\sigma)})$ taken separately,
1585 each k th set of $\beta_{k,1:C}$, and each j th set $(\theta_{1:K,j}, \xi_j)$.

1586 **Appendix C. JAGS Code for Implementing the Between-participant** 1587 **Version of AIM**

1588 Below is JAGS code for implementing the between-participant version of
1589 AIM. In addition, a downloadable package is available on each of the authors'
1590 websites that contains (1) the JAGS code below, (2) R code for executing the
1591 JAGS code on Question 1, and (3) the data files from Experiment 1. Finally,
1592 Table C.3 contains a list of the variables that need to be passed to JAGS for
1593 implementation.

```
1594 # Inferring the parameters of the between-participant version of AIM
1595 model{
1596     # Prior Data
1597     for (i in 1:n.prior){
1598         prior[i] ~ dnorm(theta,1/delta^2)
1599     }
1600     # Posterior Data
1601     for (i in 1:n.post){
1602         den[i] <- ((1/delta^2) + (1/lambda[post[i,2]]^2))
1603         b[i] <- den[i]^(-1/2)
1604         a[i] <- ((theta/delta^2) + (alpha[post[i,2]]/lambda[post[i,2]]^2))
1605         / den[i]
1606         post[i,1] ~ dnorm(a[i],1/b[i]^2)
1607     }
1608     # Priors
1609     theta ~ dnorm(mu,1/10^2)
1610     delta ~ dgamma(1,1)
1611     for(k in 1:n.cond){
1612         lambda[k] ~ dgamma(1,1)
1613     }
1614 }
```

Table C.3: Key variable definitions for implementing the JAGS code.

Variable	Description
n.post	total number of observations in the “posterior” conditions
n.cond	number of “posterior” conditions
n.prior	number of observations in “prior” condition
mu	value for the mean parameter on the prior for θ
alpha	vector containing the values of the anchors used in each “posterior” condition
prior	a vector containing observations from the “prior” condition
post	matrix containing (column 1) the observed “posterior” data, and (column 2) corresponding condition number

1615 Appendix D. JAGS Code for Implementing the Within-participant 1616 Version of AIM

1617 Below is JAGS code for implementing the within-participant version of
1618 AIM. In addition, a downloadable package is available on each of the authors'
1619 websites that contains (1) the JAGS code below, (2) R code for executing
1620 the JAGS code on the data reported in Experiment 2, and (3) the data files
1621 from Experiment 2. Finally, Table D.4 contains a list of the variables that
1622 need to be passed to JAGS for implementation.

```
1623 # Inferring the parameters of the within-participant version of AIM
1624 model{
1625   for(k in 1:n.questions){
1626     for(j in 1:n.subjects){
1627       # delta[j,k] <- beta[k,z[j,k]]
1628       prior[j,k] ~ dnorm(theta[j,k], 1/beta[k,z[j,k]]^2)
1629       psi[j,k] <- pnorm(alpha[k],theta[j,k],1/beta[k,z[j,k]]^2)
1630       #n=3; Bezier curve calculated manually below.
1631       b0[j,k] <- (1-psi[j,k])^(3)
1632       b1[j,k] <- 3 * psi[j,k] * (1-psi[j,k])^2
1633       b2[j,k] <- 3 * psi[j,k]^2 * (1-psi[j,k])
1634       b3[j,k] <- psi[j,k]^3
1635       #w=[1, xi[j], xi[j], 1]
1636       den[j,k] <- b0[j,k]*1 + b1[j,k]*xi[j] + b2[j,k]*xi[j] + b3[j,k]*1
1637       num[j,k] <- b0[j,k]*1*P[1,2] + b1[j,k]*xi[j]*P[2,2]
1638       + b2[j,k]*xi[j]*P[3,2] + b3[j,k]*1*P[4,2]
1639       B.psi[j,k] <- num[j,k]/den[j,k]
1640       lambda[j,k] <- B.psi[j,k]*lambda.max[k]
1641       den2[j,k] <- ((1/beta[k,z[j,k]]^2) + (1/lambda[j,k]^2))
1642       b[j,k] <- den2[j,k]^(-1/2)
1643       a[j,k] <- ((theta[j,k]/beta[k,z[j,k]]^2) + (alpha[k]/lambda[j,k]^2))
1644       / den2[j,k]
1645       post[j,k] ~ dnorm(a[j,k],1/b[j,k]^2)
1646     }}
1647   # priors
1648   for(k in 1:n.questions){
1649     for(j in 1:n.subjects){
1650       theta[j,k] ~ dnorm(theta.mu[k], 1/theta.sigma[k]^2)
```

```

1651     }}
1652     for(k in 1:n.questions){
1653     for(c in 1:n.confcat){
1654         beta[k,c] ~ dunif(0,20)
1655     }}
1656     for(j in 1:n.subjects){
1657         xi[j] ~ dnorm(xi.mu,1/xi.sigma^2)T(0,)
1658     }
1659     for(k in 1:n.questions){
1660         theta.mu[k] ~ dnorm(4,1/10^2)
1661         theta.sigma[k] ~ dgamma(1,1)
1662     }
1663     xi.mu ~ dnorm(10,1/20^2)T(0,)
1664     xi.sigma ~ dgamma(1,1)
1665 }

```

Table D.4: Key variable definitions for implementing the JAGS code.

Variable	Description
n.questions	total number of questions
n.subjects	total number of subjects
n.confcat	total number of confidence categories
lambda.max	vector containing the maximum value of λ
alpha	vector containing the values of the anchors used in each “posterior” condition
z	matrix containing confidence judgments
prior	matrix containing estimates in “prior” condition
post	matrix containing estimates in “posterior” condition
P	matrix containing the set of control points in Bezier curve

Appendix E. References

- Adaval, R., Monroe, K. B., 2002. Automatic construction and use of contextual information for product and price evaluations. *Journal of Consumer Research* 28, 572–588.
- Allcott, H., Mullainathan, S., 2010. Behavioral science and energy policy. *Science* 327, 1204–1205.
- Anderson, J. R., 1990. *Cognitive psychology and its implications*. Worth Publishers, New York, NY.
- Ariely, D., Loewenstein, G., Prelec, D., 2003. Coherent arbitrariness: stable demand curves without stable preferences. *The Quarterly Journal of Economics* 118, 73–105.
- Baker, C. L., Tenenbaum, J. B., Saxe, R. R., 2007. Goal inference as inverse planning. In: *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*.
- Bhatia, S., Chaudhry, S. J., 2013. The dynamics of anchoring in bidirectional associative memory networks. In: *Proceedings of the 35th Annual Conference of the Cognitive Science Society*.
- Birnbaum, M. H., Zimmermann, J. M., 1998. Buying and selling prices of investments: Configural weight model of interactions predicts violations of joint independence. *Organizational Behavior and Human Decision Processes* 74, 145–187.
- Blankenship, K. L., Wegener, D. T., Petty, R. E., Detweiler-Bedell, B., Macy, C. L., 2008. Elaboration and consequences of anchored estimates: An attitudinal perspective on numerical anchoring. *Journal of Experimental Social Psychology* 44, 1465–1476.
- Brewer, N. T., Chapman, G. B., 2002. The fragile basic anchoring effect. *Journal of Behavioral Decision Making* 7, 223–242.
- Brown, S., Heathcote, A., 2005. A ballistic model of choice response time. *Psychological Review* 112, 117–128.

- 1695 Brown, S., Heathcote, A., 2008. The simplest complete model of choice re-
1696 action time: Linear ballistic accumulation. *Cognitive Psychology* 57, 153–
1697 178.
- 1698 Buhrmester, M., Kwang, T., Gosling, S. D., 2011. Amazon’s mechanical
1699 turk: A new source of inexpensive, yet high-quality, data? *Perspectives*
1700 *on Psychological Science* 6, 3–5.
- 1701 Carlson, B. W., 1990. Anchoring and adjustment in judgments under risk.
1702 *Journal of Experimental Psychology: Learning, Memory, and Cognition*
1703 16, 665–676.
- 1704 Cervone, D., Peake, P. K., 1986. Anchoring, efficacy, and action: the influence
1705 of judgmental heuristics on self-efficacy judgments and behavior. *Journal*
1706 *of Personality and Social Psychology* 50, 592–501.
- 1707 Chapman, G. B., Bornstein, B. H., 1996. The more you ask for, the more you
1708 get: Anchoring in personal injury verdict. *Applied Cognitive Psychology*
1709 10, 519–540.
- 1710 Chapman, G. B., Johnson, E. J., 1994. The limits of anchoring. *Journal of*
1711 *Behavioral Decision Making* 7, 223–242.
- 1712 Chapman, G. B., Johnson, E. J., 1999. Anchoring, activation and the con-
1713 struction of value. *Organizational Behavior and Human Decision Processes*
1714 79, 115–153.
- 1715 Chapman, G. B., Johnson, E. J., 2002. Incorporating the irrelevant: Anchors
1716 in judgment of belief and value. In: Gilovich, T., Griffin, D., Kahneman,
1717 D. (Eds.), *Heuristics and biases: The psychology of intuitive judgment*.
1718 Cambridge, England: Cambridge University Press, pp. 120–138.
- 1719 Chater, N., Manning, C. D., 2006. Probabilistic models of language process-
1720 ing and acquisition. *Trends in Cognitive Sciences* 10, 335–344.
- 1721 Critcher, C. R., Gilovich, T., 2008. Incidental environmental anchors. *Journal*
1722 *of Behavioral Decision Making* 21, 241–251.
- 1723 Englich, B., 2008. When knowledge matters: Differential effects of available
1724 knowledge in standard and basic anchoring. *European Journal of Social*
1725 *Psychology* 38, 896–904.

- 1726 English, B., Mussweiler, T., 2001. Sentencing under uncertainty: Anchoring
1727 effects in the courtroom. *Journal of Applied Social Psychology* 31, 1535–
1728 1551.
- 1729 English, B., Mussweiler, T., Strack, F., 2006. Playing dice with criminal
1730 sentences: The influence of irrelevant anchors on experts' judicial decision
1731 making. *Personality and Social Psychology Bulletin* 32, 188–200.
- 1732 Epley, N., 2004. A tale of tuned decks? anchoring as accessibility and anchor-
1733 ing as adjustment. In: Koehler, D. J., Harvey, N. (Eds.), *The Blackwell*
1734 *handbook of judgment and decision making*. Oxford, UK: Blackwell, pp.
1735 240–256.
- 1736 Epley, N., Gilovich, T., 2001. Putting adjustment back in the anchoring
1737 and adjustment heuristic: Differential processing of self-generated and
1738 experimenter-provided anchors. *Psychological Science* 12, 391–396.
- 1739 Epley, N., Gilovich, T., 2006. The anchoring and adjustment heuristic: Why
1740 adjustments are insufficient. *Psychological Science* 17, 311–318.
- 1741 Frederick, S., Mochon, D., 2012. A scale distortion theory of anchoring. *Jour-*
1742 *nal of Experimental Psychology: General* 141, 124–133.
- 1743 Furnham, A., Boo, H. C., 2011. A literature review of the anchoring effect.
1744 *The Journal of Socio-Economics* 40, 35–42.
- 1745 Galinsky, A. D., Mussweiler, T., 2001. First offers as anchors: the role of
1746 perspective-taking and negotiator focus. *Journal of Personality and Social*
1747 *Psychology* 81, 657–669.
- 1748 Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B., 2004. *Bayesian Data*
1749 *Analysis*. Chapman and Hall, New York, NY.
- 1750 Gilbert, D. T., 2002. Inferential correction. In: Gilovich, T., Griffin, D.,
1751 Kahneman, D. (Eds.), *Heuristics and biases: The psychology of intuitive*
1752 *judgment*. Cambridge, England: Cambridge University Press, pp. 167–184.
- 1753 Gilovich, T., Griffin, D., Kahneman, D., 2002. *Heuristics and biases: The*
1754 *psychology of intuitive judgment*. Cambridge University Press, Cambridge,
1755 England.

- 1756 Goldstein, W. M., Einhorn, H. J., 1987. Expression theory and the preference
1757 reversal phenomena. *Psychological Review* 94, 236–254.
- 1758 Green, D., Jacowitz, K. E., Kahneman, D., McFadden, D., 1998. Referendum
1759 contingent valuation, anchoring, and willingness to pay for public goods.
1760 *Resource and Energy Economics* 20, 85–116.
- 1761 Grice, H. R., 1975. Logic and conversation. In: Cole, P., Morgan, J. L.
1762 (Eds.), *Syntax and semantics 3: Speech acts*. New York: Academic Press,
1763 pp. 41–58.
- 1764 Griffiths, T. L., Tenenbaum, J. B., 2011. Predicting the future as bayesian
1765 inference: People combine prior knowledge with observations when esti-
1766 mating duration and extent. *Journal of Experimental Psychology: General*
1767 140, 725–743.
- 1768 Hogarth, R. M., Einhorn, H. J., 1992. Order effects in belief updating: The
1769 belief-adjustment model. *Cognitive Psychology* 24, 1–55.
- 1770 Jacowitz, K. E., Kahneman, D., 1995. Measures of anchoring in estimation
1771 tasks. *Personality and Social Psychology Bulletin* 21, 1161–1167.
- 1772 Johnson, E. J., Goldstein, D. G., 2003. Do defaults save lives? *Science* 302,
1773 1338–1339.
- 1774 Johnson, E. J., Schkade, D. A., 1989. Bias in utility assessments: Further
1775 evidence and explanations. *Management Science* 35, 406–424.
- 1776 Johnson, J. G., R., B. J., 2005. A dynamic, stochastic, computational model
1777 of preference reversal phenomena. *Psychological Review* 112, 841–861.
- 1778 Kahneman, D., Tversky, A., 1974. Judgment under uncertainty: Heuristics
1779 and biases. *Science* 185, 1124–1131.
- 1780 Klayman, J., Ha, Y. W., 1991. Confirmation, disconfirmation, and informa-
1781 tion in hypotheses testing. *Psychological Review* 94, 211–228.
- 1782 Kruschke, J. K., 1992. ALCOVE: An exemplar-based connectionist model of
1783 category learning. *Psychological Review* 99, 22–44.
- 1784 Lee, M. D., 2004. A Bayesian analysis of retention functions. *Journal of*
1785 *Mathematical Psychology* 48, 310–321.

- 1786 Lee, M. D., 2008. Three case studies in the Bayesian analysis of cognitive
1787 models. *Psychonomic Bulletin and Review* 15, 1–15.
- 1788 Lee, M. D., Wagenmakers, E.-J., 2010. A course in Bayesian
1789 graphical modeling for cognitive science, available from
1790 <http://users.fmg.uva.nl/ewagenmakers/BayesCourse/BayesBook.pdf>;
1791 last downloaded February 26, 2010.
- 1792 Lewandowsky, S., Griffiths, T. L., Kalish, M. L., 2009. The wisdom of indi-
1793 vidual: Exploring peoples knowledge about everyday events using iterated
1794 learning. *Cognitive Science* 33, 969–998.
- 1795 Lichtenstein, S., Slovic, P., 1971. Reversals of preference between bids and
1796 choices in gambling decisions. *Journal of Experimental Psychology* 89, 46–
1797 55.
- 1798 Liu, C. C., Aitkin, M., 2008. Bayes factors: Prior sensitivity and model
1799 generalizability. *Journal of Mathematical Psychology* 52, 362–375.
- 1800 McElroy, T., Dowd, K., 2007. Susceptibility to anchoring effects: how
1801 openness to experience influences responses to anchoring cues. *Judgment and*
1802 *Decision Making* 2, 48–53.
- 1803 Mochon, D., Frederick, S., 2013. Anchoring in sequential judgments. *Orga-*
1804 *nizational Behavior and Human Decision Processes* 122, 69–79.
- 1805 Mussweiler, T., 2001. The durability of anchoring effects. *European Journal*
1806 *of Social Psychology* 31, 431–442.
- 1807 Mussweiler, T., 2003. Comparison processes in social judgment: Mechanisms
1808 and consequences. *Psychological Review* 110, 472–489.
- 1809 Mussweiler, T., Englich, B., 2005. Subliminal anchoring: Judgmental conse-
1810 quences and underlying mechanisms. *Organizational Behavior and Human*
1811 *Decision Processes* 98, 133–143.
- 1812 Mussweiler, T., Strack, F., 1999. Hypothesis-consistent testing and semantic
1813 priming in the anchoring paradigm: A selective accessibility model. *Journal*
1814 *of Experimental Social Psychology* 35, 136–164.

- 1815 Mussweiler, T., Strack, F., 2000. The use of category and exemplar knowl-
1816 edge in the solution of anchoring tasks. *Journal of Personality and Social*
1817 *Psychology* 78, 1038–1052.
- 1818 Mussweiler, T., Strack, F., 2001. Consider the impossible: Explaining the
1819 effects of implausible anchors. *Social Cognition* 19, 145–160.
- 1820 Mussweiler, T., Strack, F., Pfeiffer, T., 2000. Overcoming the inevitable an-
1821 choring effect: Considering the opposite compensates for selective accessi-
1822 bility. *Personality and Social Psychology Bulletin* 26, 1142–1150.
- 1823 Northcraft, G. B., Neale, M. A., 1987. Experts, amateurs, and real estate:
1824 An anchoring-and-adjustment perspective on property pricing decisions.
1825 *Organizational Behavior and Human Decision Processes* 39, 44–64.
- 1826 Nosofsky, R. M., 1986. Attention, similarity, and the identification-
1827 categorization relationship. *Journal of Experimental Psychology: General*
1828 115, 39–57.
- 1829 Oppenheimer, D. M., LeBoeuf, R. A., Brewer, N. T., 2008. Anchors aweigh:
1830 A demonstration of cross-modality anchoring and magnitude priming. *Cog-*
1831 *nition* 106, 13–26.
- 1832 Paolacci, G., Chandler, J., Ipeirotis, P. G., 2010. Running experiments on
1833 amazon mechanical turk. *Judgment and Decision Making* 5, 411–419.
- 1834 Petty, R. E., Cacioppo, J., 1986. The elaboration likelihood model of persua-
1835 sion. In: Berkowitz, L. (Ed.), *Advances in experimental social psychology*,
1836 19. New York: Academic Press, pp. 123–205.
- 1837 Plous, S., 1989. Thinking the unthinkable: the effect of anchoring on likeli-
1838 hood estimates of nuclear war. *Journal of Applied Social Psychology* 19,
1839 67–91.
- 1840 Plummer, M., 2003. JAGS: A program for analysis of Bayesian graphical
1841 models using Gibbs sampling. In: *Proceedings of the 3rd International*
1842 *Workshop on Distributed Statistical Computing*.
- 1843 Quattrone, G. A., 1982. Overattribution and unit formation: When behavior
1844 engulfs the person. *Journal of Personality and Social Psychology* 42, 593–
1845 607.

- 1846 Quattrone, G. A., Lawrence, C. P., Finkel, S. E., Andrus, D. C., 1984. Ex-
1847 plorations in anchoring: The effects of prior range, anchor extremity, and
1848 suggestive hints, unpublished Manuscript.
- 1849 Ratcliff, R., 1978. A theory of memory retrieval. *Psychological Review* 85,
1850 59–108.
- 1851 Robert, C. P., Casella, G., 2004. Monte Carlo statistical methods. Springer,
1852 New York, NY.
- 1853 Rubin, D. C., Wenzel, A. E., 1996. One hundred years of forgetting: A
1854 quantitative description of retention. *Psychological Review* 4, 734–760.
- 1855 Schwarz, N., 1994. Judgments in a social context: Biases, shortcomings, and
1856 the logic of conversation. In: Zanna, M. (Ed.), *Advances in Experimental*
1857 *Social Psychology*. San Diego: Academic Press, pp. 125–162.
- 1858 Silverman, B. W., 1986. Density estimation for statistics and data analysis.
1859 London: Chapman & Hall.
- 1860 Simmons, J. P., LeBoeuf, R. A., Nelson, L. D., 2010. The effect of accuracy
1861 motivation on anchoring and adjustment: Do people adjust from provided
1862 anchors? *Journal of Personality and Social Psychology* 99, 917–932.
- 1863 Simonson, I., Drolet, A., 2004. Anchoring effects on consumers’ willingness-
1864 to-pay and willingness-to-accept. *Journal of Consumer Research* 31, 681–
1865 690.
- 1866 Slovic, P., 1967. The relative influence of probabilities and payoffs upon per-
1867 ceived risk of a gamble. *Psychonomic Science* 9, 223–224.
- 1868 Slovic, P., Lichtenstein, S., 1968. Relative importance of probabilities and
1869 payoffs in risk taking. *Journal of Experimental Psychology Monograph* 78,
1870 1–18.
- 1871 Smith, A. R., Windschitl, P. D., 2015. Resisting anchoring effects: The roles
1872 of metric and mapping knowledge. *Memory & Cognition* 43, 1071–1084.
- 1873 Smith, A. R., Windschitl, P. D., Bruchmann, K., 2013. Knowledge matters:
1874 Anchoring effects are moderated by knowledge level. *European Journal of*
1875 *Social Psychology* 43, 97–108.

- 1876 Stewart, N., 2009. The cost of anchoring on credit-card minimum repayments.
1877 Psychological Science 20, 39–41.
- 1878 Steyvers, M., Griffiths, T. L., Dennis, S., 2006. Probabilistic inference in
1879 human semantic memory. Trends in Cognitive Sciences 10, 327–334.
- 1880 Strack, F., Mussweiler, T., 1997. Explaining the enigmatic anchoring effect:
1881 Mechanisms of selective accessibility. Journal of Personality and Social
1882 Psychology 73, 437–446.
- 1883 Tenenbaum, J. B., Griffiths, T. L., Kemp, C., 2006. Theory-based bayesian
1884 models of inductive learning and reasoning. Trends in Cognitive Sciences
1885 10, 309–318.
- 1886 ter Braak, C. J. F., 2006. A Markov chain Monte Carlo version of the ge-
1887 netic algorithm Differential Evolution: easy Bayesian computing for real
1888 parameter spaces. Statistics and Computing 16, 239–249.
- 1889 Thaler, R. H., Sunstein, C. R., 2008. Nudge: Improving decisions about
1890 health, wealth, and happiness. Yale University Press.
- 1891 Treisman, M., Williams, T., 1984. A theory of criterion setting with an ap-
1892 plication to sequential dependencies. Psychological Review 91, 68–111.
- 1893 Turner, B. M., Sederberg, P. B., Brown, S., Steyvers, M., 2013. A method
1894 for efficiently sampling from distributions with correlated dimensions. Psy-
1895 chological Methods 18, 368–384.
- 1896 Turner, B. M., Van Zandt, T., Brown, S. D., 2011. A dynamic, stimulus-
1897 driven model of signal detection. Psychological Review 118, 583–613.
- 1898 Urminsky, O., Goswami, I., 2015. In search of optimally effective defaults,
1899 working Paper.
- 1900 Usher, M., McClelland, J. L., 2001. On the time course of perceptual choice:
1901 The leaky competing accumulator model. Psychological Review 108, 550–
1902 592.
- 1903 van Rooijen, M. R., Daamen, D. D. L., 2006. Subliminal anchoring: The
1904 effects of subliminally presented numbers on probability estimates. Journal
1905 of Experimental Social Psychology 42, 380–387.

- 1906 Wansink, B., Kent, R. J., Hoch, S. J., 1998. An anchoring and adjustment
1907 model of purchase quantity decisions. *Journal of Marketing Research* 35,
1908 71–81.
- 1909 Wason, P. C., 1960. On the failure to eliminate hypotheses in a conceptual
1910 task. *Quarterly Journal of Experimental Psychology* 12, 129–140.
- 1911 Wegener, D. T., Petty, R. E., Blankenship, K. L., Detweiler-Bedell, B., 2010.
1912 Elaboration and numerical anchoring: Breadth, depth, and the role of
1913 (non-)thoughtful processes in anchoring theories. *Journal of Consumer*
1914 *Psychology* 20, 5–16.
- 1915 Wegener, D. T., Petty, R. E., Detweiler-Bedell, B., Jarvis, W. B. G., 2001.
1916 Implications of attitude change theories for numerical anchoring: Anchor
1917 plausibility and the limits of anchor effectiveness. *Journal of Experimental*
1918 *Social Psychology* 37, 62–69.
- 1919 Wilson, T. D., Houston, C. E., Etling, K. M., Brekke, N., 1996. A new
1920 look at anchoring effects: Basic anchoring and its antecedents. *Journal of*
1921 *Experimental Psychology: General* 125, 387–402.
- 1922 Wixted, J. T., 1990. Analyzing the empirical course of forgetting. *Journal of*
1923 *Experimental Psychology: Learning, Memory, and Cognition* 16, 927–935.
- 1924 Wong, K. F. E., Kwong, J. Y. Y., 2000. Is 7300 m equal to 7.3 km? same
1925 semantics but different anchoring effects. *Organizational Behavior and Hu-*
1926 *man Decision Processes* 82, 314–333.
- 1927 Yuille, A., Kersten, D., 2006. Vision as bayesian inference: analysis by syn-
1928 thesis? *Trends in Cognitive Sciences* 10, 301–308.
- 1929 Zhang, Y. C., Schwarz, N., 2013. The power of precise numbers: A con-
1930 versational logic analysis. *Journal of Experimental Social Psychology* 49,
1931 944–946.