Dorst & Mandelkern 2023, Good Guesses

STRUCTURA Metaphysics Reading Group

Helena Fang / Mar. 31st, 2023

TL; DR:

- This paper concerns *guessing*: how people respond to a question when they aren't certain of the answer.
- Dorst & Mandelkern (DM): people aim to optimize a tradeoff between *accuracy* and *informativity* when forming their guess.
- This account yields new theories of belief, assertion and the conjunction fallacy[1], & helps to explain how boundedly rational agents explore the world.

[1] For the time being, I'll skip the section on conjunction fallacy.

1. Take a Guess

Latif is accepted at four law schools: Yale, Harvard, Stanford, and NYU. We don't know his preferences, but here's the data on where applicants who've had the same choice had gone in recent years:

Yale	Harvard	Stanford	NYU
38%	30%	20%	12%

Take a guess: Where do you think he'll go?

Claim: guesses like 'Yale', 'Either Yale or Harvard' sound fine. Meanwhile, it's unnatural to guess 'not Yale', or 'Yale, Stanford, or NYU'. On a first pass, these guesses seem puzzling[2].

[2] For instance, 'Yale' is a fine guess, but its probability is lower than 50%. Moreover, 'Yale or Harvard' sounds fine (it's okay to guess something other than the single most likely school), yet 'Yale, Stanford, or NYU' sounds weird (why leave out 'Harvard'?).

This paper cashes out the underlying pattern for our rational guesses in this sort of questions. The idea is that guessers aim to optimize a tradeoff between *accuracy* and *informativity*, measured by a Jamesian tradeoff, which is treated differently by different guessers in different contexts.

DM further argue that guessing, along with its accuracy-informativity tradeoff, plays a central role in our cognitive lives. Specifically on:

- *Belief*: this account underpins a promising theory of belief due to Holguin (2022), who argues that your beliefs are your best guesses;
- *Assertion*: this theory helps to both explain and generalize the standard pragmatic story about how conversation proceed;
- *Conjunction Fallacy*: this theory helps to explain the *conjunction fallacy*, the psychological finding that people sometimes rank a conjunction as more probable than one of its conjuncts, contra the law of probability.

2. *What We Guess*

What sort of guesses do we tend to make? The answer is both surprising and surprisingly *systematic*. Recall the case of Latif. It seems that there is a range of answers that could reasonably be your guess, given your credences:

(1) a. Yale. ✓
b. Yale or Harvard. ✓
c. Yale or Harvard or Stanford. ✓
d. Yale, Harvard, Stanford, or NYU. ✓

There is also a range of answers that are intuitively unacceptable, for instance:

(2) a. Harvard. ✗
b. Stanford. ✗
c. NYU. ✗
d. Yale or Stanford. ✗
e. Yale or NYU. ✗
f. Harvard or Stanford. ✗
g. Not Yale. ✗
h. Harvard, Stanford, or NYU. ✗
i. Yale, and it's cold in London today. ✗
j. Yale, or he has a birthmark on his left toe. ✗

The claim is *normative*: there is something peculiar - something irrational - about guesses like this.

First off: what a *question* is and what its *answers* are. Drawing from standard theories in semantics and pragmatics, A *question* is a partition of the context set[3] (i.e. a set of mutually exclusive and jointly exhaustive subsets of the context set) (Hamblin 1973, Karttunen 1977, Groenendijk and Stokhof 1984). The cells of the partition are the *complete answers* to the question.

In our current example: [[Where will Latif go to law school?]] = {*Latif will go to Yale*, *Latif will go to Harvard*, *Latif will go to Stanford*, *Latif will go to NYU*}.

A few assumptions: at any given point in a conversation, there's a *question under discussion* (QUD) that guessers aim to address (Roberts, 2012); guessers have credences which can be modeled with a probability function P that is regular over the context set; questions are always finite partitions.

Now, some generalizations from the observations above.
(i) You don't always have to pick an answer that is more likely

[3] A *context set* is a set of possible worlds that comprise all and only the worlds compatible with the assumptions in a given context (Stalnaker 1974, 1978). DM consider cases in which the context set and the guesser's certainties coincide.

than not to obtain.

Improbable Guessing: It's sometimes permissible to answer p even when $P(p) < 0.5$.

(ii) Suppose that your credences are as above, but instead, you're asked: 'Will Latif go to Yale?' i.e. the question $\{$*Yale*, *not Yale*$\}$. Holding your credence as above, when addressing *this* question, 'Yale' is not a very natural guess. Thus:

Question Sensitivity: Whether p is a permissible answer depends not just on the guesser's credence in p but also in what question is being answered.

(iii) One may think that to account for Improbable Guessing and Question Sensitivity, given a question Q, your answer should be the complete answer you have highest credence in. But this overgeneralizes: it predicts that only complete answers are permissible guesses, whereas a range of partial answers (i.e. unions of complete answers) are permissible. Instead we may say:

Optionality: Given any question Q, for any k such that $1 \leq k \leq |Q|$, it's permissible for your guess about Q to be the union of exactly k cells of Q.

(iv) To capture Optionality, we need a further constraint: your guess can't include a complete answer q while excluding a strictly more likely complete answer q'.

Filtering: A guess about Q is permissible only if it is *filtered*: if it includes a complete answer q, it must include all complete answers that are more probable than q. Precisely: p is filtered iff for any $q, q' \in Q$: if $P(q') > P(q)$ and $q \subset p$, then $q' \supset p$.

Optionality and Filtering together predict the admissibility of the answers in (1), together with the inadmissibility of the answers in (2-a)-(2-h).

(v) Regarding answers like (2-i)-(2-j): intuitively, they include *irrelevant* material[4], i.e. they cannot be derived as a union of complete answers to the QUD.

Fit: If a guess crosscuts a complete answer, it's impermissible. Precisely: p is a permissible guess only if there are $q_1, ..., q_k \in Q$ such that $p = q_1 \cup ... \cup q_k$.

These observations – Improbable Guessing, Question Sensitivity, Optionality, Filtering, and Fit – bring out what guesses people tend to make, revealing surprising yet systematic patterns.

[4] Note that some apparent violations of Fit can be felicitous, e.g. 'Latif will go to Yale, and I'm sure he'll love it!' DM: answer of this sort is felicitous, since it is easy to accommodate more fine-grained questions that are in a similar vein to the QUD, e.g. 'Where will Latif go, and will he like it?'. Relative to the finer-grained question, the answer satisfies Fit. In constrast, (2-i)-(2-j) are infelicitous, since thethe finer-grained question which would need to be accommodated to satisfy Fit seem too irrelevant to the original QUD.

3. How We Guess

MD then give a model of *how* we guess, which is intended to be a computational level explanation. The general thought: there's an inevitable tradeoff between two goals. On the one hand, we want *accuracy*: we want our guess to be true. But some guesses, though guaranteed true, say very little[5]. So we also want to have an *informative* guess, one that helpfully narrows down the space of alternatives we're considering.

The question: What problem is a (rational) mind solving when it forms a guess?

The answer: How to optimally trade off accuracy and informativity.

[5] When asked 'Where do you think Latif will go?', 'Somewhere' is sure to be true, but is unhelpful.

3.1 Jamesian guessing

The current approach views guessing as a kind of *epistemic decision problem*: we first say what makes a guess *objectively valuable*, then propose that people aim to maximize this objective value by choosing a guess with the highest *expected* value, given their credences.

Let $V_Q(p)$ be a (real-valued) function which yields the *answer-value* of choosing p as your guess about Q, depending on whether p is true or false. Whenever you're unsure whether p is true, you'll be unsure how much answer-value it has–yet, you can use your credences in the various possibilities to form an *estimate* about how much answer-value it has. p's *expected answer-value* is modeled as such:

$$E_Q(p) := P(p) \cdot V_Q^+(p) + P(\overline{p}) \cdot V_Q^-(p)$$

Guessing as Maximizing: A guess is epistemically permissible given a question iff it has maximal expected answer-value relative to that question, for some permissible measure of answer-value.

The crucial question: Which measures of answer-value are epistemically permissible?

First: true guesses are better than false ones, so any permissible V_Q must be *truth-directed*:

V_Q is **truth-directed** iff any true guess has higher answer-value than any false guess.
Precisely: for all p, r: $V_Q^+(p) > V_Q^-(r)$.

Moreover: informativity matters. Given a question Q and an answer p, let the **informativity** of p relative to Q be the proportion of complete answers to Q that p rules out:

$$Q_p := \frac{|\{q \in Q : p \cap q = \emptyset\}|}{|Q|}.$$

Therefore, second constraint: given the truth-value of p, $V_Q(p)$ should then be fully determined by p's informativity:

> V_Q is **question-based** iff for all p: $V_Q(p)$ is fully determined by p's informativity together with its truth-value.
> Precisely: for all p, r, if $Q_p = Q_r$, then $V_Q^+(p) = V_Q^+(r)$ and $V_Q^-(p) = V_Q^-(r)$.

Therefore: a measure of answer-value is (epistemically) permissible only if it is truth-directed and question-based. This establishes Fit and Filtering.

How about Improbable Guessing and Optionality? DM motivate a particular subclass of truth-directed, question-based measures as the epistemically permissible ones: **Jamesian measures**[6], for which there is some $J \geq 1$ such that, for all p, $V_Q^+(p) = J^{Q_p}$, and $V_Q^+(p) = 0$.

[6] DM note that, the goal of these models is not to *predict* what particular people will guess in particular situation, but rather, to *elucidate* the structural features that the practice of guessing is sensitive to, and thus explain why guesses rationally should–and thus in fact tend to–meet various constraints.

> **Jamesian Expected Answer-Value**:
> $E_Q^J(p) = P(p) \cdot J^{Q_p} + P(\overline{p}) \cdot 0 = P(p) \cdot J^{Q_p}$

The idea: the way to maximize expected answer-value is to pick an informative guess – in the limit, as $J \to \infty$, the way to do so is to pick a *maximally* informative (filtered) guess, regardless of how low its probability is[7].

[7] I'm not sure about the second part.

Precisely: V_Q is **Jamesian** iff, for some $t > 0$ and $J \geq 1$:

$$V_Q = \begin{array}{ll} V_Q^+(p) = J^{Q_p} \cdot t & \text{if } p \text{ is true} \\ V_Q^+(p) = 0 & \text{if } p \text{ is false} \end{array}$$

3.2 *Deriving our constraints*

Fit & Filtering follow from any measure which is truth-directed and question-based. In particular, Filtering is a special case of a more general constraint:

> **Filtered Rankings**: Equally informative answers should be ranked by probability.
> Precisely: if $Q_p = Q_r$, then $E_Q(p) > E_Q(r)$ iff $P(p) > P(r)$.

Improbable Guessing follows from the idea that high informativity can outweigh low probability, especially as J grows large.

Question Sensitivity follows, since Jamesian measures are based on a guess p's informativity Q_p, which in turn is determined by the question Q.

Optionality follows from the basic idea: when J is low, being informative provides little additional value, so the best guess is an uninformative (but definitely true) guess; as J grows, being more informative gradually mattes more you eventually start preferring an answer comprising the union of less cells. That is, we can rationalize guesses of different levels of informativity by ascribing to guessers different J-values, i.e. different weights on informativity.

3.3. Setting $J-$values

A natural question: how are $J-$values set? And how do we know what subjects' $J-$values are?

A natural answer (idea): Your $J-$value is determined by one's mental state in broadly the way credences, utilities, and (on some views) risk profiles are. In general, it's very flexible[8].

[8] Moreover, once we consider the role of guessing, it explains *why* $J-$values would need to be flexible like this.

In particular, a trend that they have observed: as the probabilities of the various complete answers 'cluster' together more tightly, it becomes increasingly strange for your guess to crosscut these clusters– to include some but not all of the cells in a cluster. TO summarize this trend:

> **Clustering**: People tend to avoid making guesses that crosscut clusters of complete answers with similar probabilities.

This indicates that people tend to select a $J-$value that makes their guess *distinction*: one that makes its expected answer-value not only maximal, but distinctively higher than that of alternative guesses.

Precisely: given credences P and a question Q, let the *J-distinctiveness* of a guess p, D^p_J, be the ratio of its expected answer-value to the highest expected answer-value of any other Fit guess (holding fixed J). That is, where F_p is the set of Fit answers to Q other than p, we have:

$$D^p_J := \frac{E^J_Q(p)}{max_{r \in F_p}(E^J_Q(r))}.$$

We then define the *distinctness* of p, D_p, to be the maximal $J-$distinctiveness it can receive, for any value of J: $D^p := \sup\{D^p_J : J \geq 1\}$. We take it to be a natural measure of the *salience* of a guess.

Proposal: there is a tendency (though not obligation) to make guesses with high distinctiveness, and thus to use $J-$value that allow guesses to have high $J-$distinctiveness.

4. When We Guess

4.1 Guess when you believe?

4.2 Guess when you talk?