

Peer Review Information

Journal: Nature Human Behaviour

Manuscript Title: The Use-the-best heuristic facilitates deception detection

Corresponding author name(s): Bruno Verschuere

Reviewer Comments & Decisions:

Decision Letter, initial version:

28th July 2022

Dear Dr Verschuere,

Thank you once again for your manuscript, entitled "Use the best, ignore the rest: How heuristics allow to tell lie from truth", and for your patience during the peer review process. I am communicating this decision to you, instead of Dr Antusch (who has been the primary handling editor), as she is currently out of the office.

Your Article has now been evaluated by 3 referees. You will see from their comments copied below that, although they find your work of potential interest, they have raised quite substantial concerns. In light of these comments, we cannot accept the manuscript for publication, but would be interested in considering a revised version if you are willing and able to fully address reviewer and editorial concerns.

We hope you will find the referees' comments useful as you decide how to proceed. If you wish to submit a substantially revised manuscript, please bear in mind that we will be reluctant to approach the referees again in the absence of major revisions. We are committed to providing a fair and constructive peer-review process. Do not hesitate to contact us if there are specific requests from the reviewers that you believe are technically impossible or unlikely to yield a meaningful outcome.

We have identified below a list of key priorities for revision. We hope that you will find the prioritised set of referee points to be useful when revising your study. Please do not hesitate to get in touch if you would like to discuss these issues further.

1) Reviewer 1 and Reviewer 2 ask for additional data to be collected in order to ensure that your studies appropriately test your hypotheses and fully support your conclusions. Specifically, Reviewer 1 suggests that you conduct an experiment to test your cue against a multi-cue condition. Reviewer 2 requests that you conduct a new experiment in which all outcome variables and independent variables, except for the cue-manipulation, are held constant. Reviewer 2 also points out that the

current designs emphasize labeling deceptiveness rather than assessing accuracy in deception detection. Editorially, we believe that the evidence would be substantially strengthened by including a study that demonstrates the treatment's ability to improve accuracy in the detection of deceptive statements (which is the primary claim in your manuscript). Therefore, to address the reviewer points and editorial concerns, we ask that you conduct 1-2 new experiments. These experiments should be preregistered and have at least 80% power to detect the smallest meaningful effect size (not observed effect sizes from the existing studies).

2) Reviewer 1 asks that you re-analyse your data using analysis of variance or regression analyses. We agree with the reviewer that ANOVA/regression would be more appropriate and ask that you perform and report these analyses as the main analyses. Editorially, we also ask you to supply follow-up Bayes Factors or equivalence tests in order to interpret null results (please see the attached checklist for more guidance on this). Finally, please note that all deviations from preregistered analyses need to be reported in the manuscript.

3) The reviewers ask that you discuss the generalizability of your findings and stimuli in detail. In doing so, please provide clarifications on the stimuli used and the exact hypotheses tested in each experiment.

Finally, your revised manuscript must comply fully with our editorial policies and formatting requirements. Failure to do so will result in your manuscript being returned to you, which will delay its consideration. To assist you in this process, I have attached a checklist that lists all of our requirements. If you have any questions about any of our policies or formatting, please don't hesitate to contact me.

If you wish to submit a suitably revised manuscript we would hope to receive it within 5 months. I would be grateful if you could contact us as soon as possible if you foresee difficulties with meeting this target resubmission date.

With your revision, please:

- Include a "Response to the editors and reviewers" document detailing, point-by-point, how you addressed each editor and referee comment. If no action was taken to address a point, you must provide a compelling argument. When formatting this document, please respond to each reviewer comment individually, including the full text of the reviewer comment verbatim followed by your response to the individual point. This response will be used by the editors to evaluate your revision and sent back to the reviewers along with the revised manuscript.
- Highlight all changes made to your manuscript or provide us with a version that tracks changes.

Please use the link below to submit your revised manuscript and related files:

[REDACTED]

Note: This URL links to your confidential home page and associated information about manuscripts you may have submitted, or that you are reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage.

Thank you for the opportunity to review your work. Please do not hesitate to contact me if you have any questions or would like to discuss the required revisions further.

Sincerely,

Stavroula Kousta

Stavroula Kousta, PhD
Chief Editor, Nature Human Behaviour

on behalf of

Samantha Antusch

Samantha Antusch, PhD
Editor, Nature Human Behaviour

Reviewer expertise:

Reviewer #1: truth discernment ; strategies to detect dis/misinformation ; decision-making

Reviewer #2: dis/misinformation ; decision-making

Reviewer #3: heuristics ; decision-making

REVIEWER COMMENTS:

Reviewer #1:
Remarks to the Author:

Review of "Use the best, ignore the rest: How heuristics allow to tell lie from truth"

The authors set out to test the idea that simple heuristics could be better at helping people detect lies over "multi-cue" lie detection methods. I found the manuscript very well written, and incredibly interesting, and the studies are well conducted. I have some rather minor observations that need to be addressed before acceptance.

In the abstract (and elsewhere) the authors say: "Our findings challenge the notion that people lack the capacity to detect deception." This is not what they show though. In the control condition (without any instructions on lie detection cues), people are doing very poorly through all experiments. What the authors show is that once people are instructed to use simple cues and explained what cues they should use for lie detection, they can get better at it. It does not challenge the notion, however, that people "lack the capacity to lie detection" under naturalistic settings (no instructions).

That being said, I think the control group does not exactly fit the hypothesis and conclusions. In the control group, people received no instructions on how to detect lies. The authors call this condition "people were free to use all available cues", basically implying that the control condition is a multi-cue control condition. However, people were not told any cues on how to detect lies. They might simply not know any reliable lie detection cues or the ones they thought they knew were completely wrong and non-indicative of actual lies. A better comparison would be, one that shows that the heuristic method the authors test is superior to the multi-cue methods, to train people on multi-cue lie detection and directly compare their performance with the performance of the heuristic group in a randomized experiment.

What about artificial intelligence lie detection? Artificial Intelligence methods are getting better at lie detection and they might achieve superior performance over the heuristic lie detection methods (e.g., <https://link.springer.com/article/10.1007/s10676-022-09621-6>). Why, in the age of AI, would we need to train people for better lie detection? Why cannot such judgments be allocated to artificial intelligence? This question warrants some discussion.

The authors use t-tests to compare veracity judgments of false and true statements, separately in the control and the experimental group. Why not test the two-way interaction with a simple ANOVA or regression? The interaction method is associated with a smaller Type 1 error rate, and would fit the research question better: a significant interaction could mean that the experimental condition increased lie detection abilities compared to the control group – a conclusion the current statistical tests do not necessarily allow to make as they do not directly compare control vs experimental group.

I appreciate Study 4, and the authors' idea to test the results' generalizability for multiple lie production modes. However, none of the production modes includes "naturalistic settings" in which liars could lie about something that they would lie about in naturalistic settings, in their normal lives. In my understanding lies were always made in or about a superficial situation such that people were instructed to make up some lie in a pre-arranged situation (i.e., find Michelle in the bus stop task). Those lies can be different from "naturally occurring" lies in many important aspects, therefore, results might not be generalized to more naturalistic lies. This should be at least discussed.

The introduction is short anyway, I see no need for titles before each paragraph.

Reviewer #2:

Remarks to the Author:

I would like to thank the authors for the opportunity to evaluate their work "Use the best, ignore the rest: How heuristics allow to tell lie from truth". In their manuscript the authors describe 7 studies in all of which they test various combinations of scenarios of deception detection and how to improve it by focussing on the detailedness and the verifiability of those details.

The paper is clearly written and I think that the research question is motivated and embedded in the literature. However, I have some concerns regarding the experimental design(s), the presentation of the results and their interpretation/generalizability.

More precisely, I was rather confused by the variety of design between the 7 experiments and their lack of parallelity between conditions. Why were the participants in the treatment condition in Experiment 1-3 asked a different question than in the control condition? (I see that that has been checked in Experiment 4, but I wonder why not earlier?) I believe that the comparison between the two conditions is problematic if different questions were asked, specifically the question in the treatment conditions seem rather like a task to label the stimuli instead of a deception detection. The problem continues, as in the following experiments there is a lack of a control groups that does not use a heuristic/specific cue. So overall, I think there is a lack of the most important comparison between a control condition and a treatment condition in a truly parallel way. This brings me to my second point that it needs to be much clearer which hypothesis was tested in which experiment, without that the collection of experiments fails to draw a coherent picture of what is being tested. It seems mainly those experiments show that detailedness and its verifiability are valid cues for handwritten(?) deceptive statements. I had a quick look into the pre-registrations that exist for some experiments (why only for some?) and found them relatively superficial, but in principle could be used to spell out the hypotheses in the manuscript. My third major concern is the choice of stimuli and their generalisability: those need to be described more clearly (an example would help), because, as I wrote above, the results obviously depend on the predictive quality of this one cue (especially if participants are only asked to detect the cue instead of a judgement), making it even more important to show stimuli and discuss their generalisability. My fourth and last major concerns regards the presentation of the results: as I expect individual differences I would like to ask for showing individual-level data of the performance of participants (e.g., box-plot + scatter) for all the seven Experiments, also to enhance transparency more generally. In addition, I was wondering if any additional information about the participants was collected to shed light on the heterogeneities of the effects.

I think that those major points need to be addressed in a thoroughly revised version, so to allow for an assessment of the results and their interpretation. Some can be addressed with an improved presentation (figures+text) others (e.g., regarding the experimental design) would probably need to be addressed with additional data.

Minor points are:

- The numbers of references should be sorted (are they? I had problems orienting).
- The motivation/intro can be a bit more specific, because the task and the heuristic at hand is much more specific than the intro sounds, the detailedness can be gamed by anyone who spreads lies professionally, in a person-to-person situation I believe it is a very valid cue, but that scope needs to be clearer.
- Similarly, it sounds as if this cue is very generally valid, how certain is the literature with this? And in which contexts?
- I was confused about the hypothesis in Experiment 5-6 that participants would be worse at the task when they know it is about deception detection, what is the reasoning behind that?
- Was there any control group in experiment 7?
- In the discussion (+intro), it needs to be made clear that the heuristics and the findings presented here hold for person-to-person deception situations (at least that is how I understand the experiments and the verifiability cue), but does not apply to professional misinformation campaigns, for example.

If those concerns can be addressed I would be happy to review a thoroughly revised version of the manuscript.

Reviewer #3:
Remarks to the Author:

Review of "Use the best, ignore the rest"
Nature Human Behavior 22061537

This paper shows, for the first time, that instructing people to rely on a single, powerful cue (detailedness), they were consistently able to detect deception. Which shows, that unlike earlier conclusions, people can learn to perform at a systematically higher level than without training ("intuitively"). That is an important finding, both practically, and theoretically. The theoretical connection is with research on fast-and-frugal heuristics, which has shown that in situations of uncertainty, keep it simple can outperform more complex strategies.

Here are a few corrections and clarification that help to improve the paper.

General: To embed this present study better theoretically, the authors should connect their terminology to the framework of smart heuristics. In the standard terminology of Gigerenzer & Gaissmaier (2011, Annual Review of Psychology), the "use the best" heuristic is a member of the class of one-reason decision-making heuristics, specifically an instance of a "one-clever cue" heuristic. To answer the question how one cue can be as powerful or more than more cues, the authors might refer to the bias-variance dilemma (Gigerenzer & Brighton, 2009, Topics in Cognitive Science).

Minor point: First page of full text, top para: I assume "only 4% higher" means "only 4 percentage points higher". Please clarify.

Author Rebuttal to Initial comments

Manuscript NATHUMBEHAV-22061537A

"Use the best, ignore the rest: How heuristics allow to tell lie from truth" Reply to
Reviewers

We were thrilled to hear that the reviewers found our submission "incredibly interesting, and the studies well conducted" (Reviewer#1), "the paper clearly written and ... the research question is motivated and embedded in the literature" (Reviewer#2), and that we discovered "... an important finding, both practically, and theoretically" (Reviewer#3).

At the same time, the reviewers asked for a number of clarifications, suggested several revisions, and asked for novel work to substantiate our claims (below, *in italic*). We addressed all these issues below in a point by point fashion (**in bold**), adding two new preregistered studies. We like to thank the editors and the reviewers for their constructive comments and hope our revised manuscript is now acceptable

for publication in *Nature Human Behavior*.

Reviewer #1

Reviewer expertise: truth discernment ; strategies to detect dis/misinformation ; decision- making

The authors set out to test the idea that simple heuristics could be better at helping people detect lies over “multi-cue” lie detection methods. I found the manuscript very well written, and incredibly interesting, and the studies are well conducted. I have some rather minor observations that need to be addressed before acceptance. In the abstract (and elsewhere) the authors say: “Our findings challenge the notion that people lack the capacity to detect deception.” This is not what they show though. In the control condition (without any instructions on lie detection cues), people are doing very poorly through all experiments. What the authors show is that once people are instructed to use simple cues and explained what cues they should use for lie detection, they can get better at it. It does not challenge the notion, however, that people “lack the capacity to lie detection” under naturalistic settings (no instructions).

We very much agree with Reviewer#1 that under naturalistic settings (no guidance) people are poor at lie detection, and once people are instructed to use the right simple cues, they can get better at it. We used the word capacity to refer to the potential skill, not actual or current skill. We now avoid the terms capacity, ability, and capability, as they may create confusion rather than clarity. We rephrased the sentence so that it now reads ‘Our findings challenge the notion that people lack the capacity potential to detect deception’.

That being said, I think the control group does not exactly fit the hypothesis and conclusions. In the control group, people received no instructions on how to detect lies. The authors call this condition “people were free to use all available cues”, basically implying that the control condition is a multi- cue control condition. However, people were not told any cues on how to detect lies. They might simply not know any reliable lie detection cues or the ones they thought they knew were completely wrong and non-indicative of actual lies. A better comparison would be, one that shows that the heuristic method the authors test is superior to the multi-cue methods, to train people on multi-cue lie detection and directly compare their performance with the performance of the heuristic group in a randomized experiment.

In our control condition, people were indeed free to use any available cue. Under such conditions, people report to (Bogaard & Meijer, 2016) and actually make use of multiple cues (Hartwig & Bond, 2011). In that sense, our control condition does reflect a multiple cue condition.

Nonetheless, we agree with Reviewer#1 that it is of interest to conduct additional research where people in the control condition explicitly rely on multiple cues when making lie-truth judgements. We conducted two such studies. In the new, preregistered, Study8 (n=146 tested, n=99 included in final analyses), people judged a single cue (detailedness) or multiple cues (detailedness, unexpected complications, admitting lack of memory, and affect) before making their binary lie-truth judgement. Our primary prediction that lie-truth discrimination would be superior in the single cue condition compared to the multiple cue condition was supported by the data,

one-tailed $t(97) = 2.013$, $p = .023$, $d = 0.408$ (95% CI: 0.07, $+\infty$), $BF_{10}=2.45$. Note that this is a very strict test of our hypothesis, as both conditions did in fact judge the same diagnostic cue (detailedness), with the sole difference being that the multiple cue condition also judged additional cues. At the same time, the BF only showed anecdotal support for our hypothesis. So we ran the study again, with more power. This also provided us with the opportunity to improve Study8 methodology (e.g., equal balance experimental conditions, random instead of blocked presentation of statements). Following the editorial advice, we preregistered Study9 and used a Bayesian sequential stopping rule, testing until the data provided decisive support for or against our hypothesis (ie. $BF_{10}>6$ or $BF_{10}<0.17$). Study9 ($n=405$ tested, $n=382$ included in final analyses) provided decisive evidence for our hypothesis that lie- truth discrimination is better in the single than in the multiple cue condition was supported, one-tailed $t(380) = 4.71$, $p < .001$, $d = 0.482$ (95% CI: 0.311, $+\infty$), $BF_{10}=7951$.

What about artificial intelligence lie detection? Artificial Intelligence methods are getting better at lie detection and they might achieve superior performance over the heuristic lie detection methods (e.g., <https://link.springer.com/article/10.1007/s10676-022-09621-6>). Why, in the age of AI, would we need to train people for better lie detection? Why cannot such judgments be allocated to artificial intelligence? This question warrants some discussion.

Current AI approaches rely on combining many cues (up to thousands), and our bold prediction is that, in out-of-sample testing, the Use-the-best heuristic may outperform current AI approaches. We added this bold prediction to the Discussion. Furthermore, we mention some of the limitations and pitfalls of AI, which make human judgement indispensable.

The authors use t-tests to compare veracity judgments of false and true statements, separately in the control and the experimental group. Why not test the two-way interaction with a simple ANOVA or regression? The interaction method is associated with a smaller Type 1 error rate, and would fit the research question better: a significant interaction could mean that the experimental condition increased lie detection abilities compared to the control group – a conclusion the current statistical tests do not necessarily allow to make as they do not directly compare control vs experimental group.

For all studies with more than one independent variable, we in fact conducted and reported an ANOVA as the main analysis (Studies 1-2-3-4-5-7). We now improved our write-up to clarify this. Specifically, for Study2-3 we now first report the results of the ANOVA, then the follow-up t-tests. Study 6-8-9 have a single independent variable making a t-test most suitable.

I appreciate Study 4, and the authors' idea to test the results' generalizability for multiple lie production modes. However, none of the production modes includes "naturalistic settings" in which liars could lie about something that they would lie about in naturalistic settings, in their normal lives. In my understanding lies were always made in or about a superficial situation such that people were instructed to make up some lie in a pre-arranged situation (i.e., find Michelle in the bus stop task). Those lies can be different from "naturally occurring" lies in many important aspects, therefore, results might not be generalized to more naturalistic lies. This should be at least discussed.

We indeed relied on an instructed lying paradigm because it has the critical benefit of being able to determine ground truth. We now acknowledge this as an important limitation of our work, and stress that our findings should be extended to (situations mimicking) real-life settings (e.g., high stake lies).

The introduction is short anyway, I see no need for titles before each paragraph.

As per reviewer suggestion and per editorial policy, we removed subtitles from the Introduction.

Reviewer #2

Expertise: dis/misinformation ; decision-making

I would like to thank the authors for the opportunity to evaluate their work "Use the best, ignore the rest: How heuristics allow to tell lie from truth". In their manuscript the authors describe 7 studies in all of which they test various combinations of scenarios of deception detection and how to improve it by focusing on the detailedness and the verifiability of those details. The paper is clearly written and I think that the research question is motivated and embedded in the literature. However, I have some concerns regarding the experimental design(s), the presentation of the results and their interpretation/generalizability. More precisely, I was rather confused by the variety of design between the 7 experiments and their lack of parallelity between conditions. Why were the participants in the treatment condition in Experiment 1-3 asked a different question than in the control condition? (I see that that has been checked in Experiment 4, but I wonder why not earlier?) I believe that the comparison between the two conditions is problematic if different questions were asked, specifically the question in the treatment conditions seem rather like a task to label the stimuli instead of a deception detection. The problem continues, as in the following experiments there is a lack of a control groups that does not use a heuristic/specific cue. So overall, I think there is a lack of the most important comparison between a control condition and a treatment condition in a truly parallel way.

Reviewer #2 wonders about the comparability of the experimental conditions in Studies 1-2-3. Here, we think it is important to stress that our control condition (asking people to make direct deception judgements, being free to choose any cue they like) is the common control condition in the deception detection literature. This is important as it allows to compare our findings with the vast deception detection literature and the baseline finding that under such naturalistic condition, people perform barely better than chance. Our novel, experimental condition receives the same stimuli and the same -100 to +100 scale, but now with the task to only judge detailedness. This small difference has a significant impact, as we predicted and showed. Still, we appreciate Reviewer #2's preference for a control experiment in which the experimental and control condition have the exact same outcome variable, which is why we ran additional studies (Study8 and 9). Here, people judge either solely detailedness or detailedness along with 3 additional cues to then make a final lie-truth judgement. These control experiments confirms - under though conditions (see above in response to Reviewer#1) - the improved lie-truth discrimination when relying on a single cue as compared to multiple cues.

This brings me to my second point that it needs to be much clearer which hypothesis was tested in which experiment, without that the collection of experiments fails to draw a coherent picture of what is being tested. It seems mainly those experiments show that detailedness and its verifiability are valid cues for handwritten (?) deceptive statements.

The overall key hypothesis of our series of studies is that the use-the-best heuristic allows ordinary people to tell lie from truth better than how they normally would. All 9 studies support this hypothesis. The designs and hypotheses of the specific studies differ as the separate studies test different aspects of the core hypothesis. These sub-hypotheses are briefly sketched in Table1. Additionally, in the revised manuscript, we now start the description of each study with its main aim and its main hypotheses being tested.

I had a quick look into the pre-registrations that exist for some experiments (why only for some?) and found them relatively superficial, but in principle could be used to spell out the hypotheses in the manuscript.

We started with exploration (Study1) that was followed up with confirmatory research (Studies 2-3-4-5-7-8-9). We think that this combination and the fact that most of our studies were preregistered is an important strength of our research. Study6 was not preregistered because the resources (bachelor thesis project) limited its scale, and because we were unsure how the introduction of the decision rule would impact the performance of the Use the best heuristic.

Our preregistration follow an established format (aspredicted.org), developed by open science pioneers, which encourages researchers to keep the preregistration concise as to clarify to readers which analytic decisions were made a priori and which post hoc (see https://aspredicted.org/messages/why_limits.php).

My third major concern is the choice of stimuli and their generalisability: those need to be described more clearly (an example would help), because, as I wrote above, the results obviously depend on the predictive quality of this one cue (especially if participants are only asked to detect the cue instead of a judgement), making it even more important to show stimuli and discuss their generalisability.

For studies 1-2-3-4-5, we provide dan example of the stimuli in the manuscript (also see below), and we made all stimuli publicly available in full on the Open Science Framework. These are the example statements:

‘I quietly walked down until the entrance of G/lab. I was in doubt about what to do (stood still for a moment). Then I walked into the corridor of G, saw a cleaner/guy with a cart and read something about using lockers at the UvA. Then I walked to the outside entrance and walked around (back of G) and looked at the kind of butterflies that are now there for the light festival. So then walked further around G. Went back inside (second floor lab) and looked for a moment at university pabo, there is a poster next to the door about participating in brain research for money. When I had read that I walked quietly to this research room.’

Note: This is an English translation of one of the Dutch statements used in Studies 1-2-3-5. All

original [pseudonymized] Dutch statements can be found on <https://osf.io/z26ar/>.

‘Okay, so after I finished the task with the café I went to the stop at the hospital. I didn't know exactly where it was, so I first meandered through here a bit, asked 'uuh I'm doing a task, can you tell me where the stop is?' And I already thought that it was this one and then I went there. Yes, and then I was supposed to look for Michelle. There were two or three people sitting there, three people sitting there, and then I asked them in Dutch if their name was Michelle. Yes, there was no Michelle there, then I sat there for five minutes, looked to see if maybe some bus was coming by where a Michelle got off, but no bus came by at all. And then I came back here and, yes, I didn't complete the task because I didn't find Michelle’.

Note: This is an English translation of one of the German statements used in Study4. All original German statements can be found on <https://osf.io/z26ar/>.

For Studies 6-7-8-9, the stimuli are available for research purposes after signing a non- disclosure agreement (as they include audiovisual recordings of the participants).

My fourth and last major concerns regards the presentation of the results: as I expect individual differences I would like to ask for showing individual-level data of the performance of participants (e.g., box-plot + scatter) for all the seven Experiments, also to enhance transparency more generally. In addition, I was wondering if any additional information about the participants was collected to shed light on the heterogeneities of the effects.

I think that those major points need to be addressed in a thoroughly revised version, so to allow for an assessment of the results and their interpretation. Some can be addressed with an improved presentation (figures+text) others (e.g., regarding the experimental design) would probably need to be addressed with additional data.

We provided a box and whisker plot for Study1 and for Study7. We are happy to now also provide such plots for all other studies. Given the maximum number of figures allowed, we provide these as Supplementary Figures in the Supplementary Information.

In the Method Section, we report all additional variables that were collected in each study (e.g., age, sex/gender, mother tongue, nationality). We have no hypotheses for these variables to explain substantial heterogeneity, so we refrain from conducting exploratory moderation analyses and instead make all data publicly available.

Minor points are:

- The numbers of references should be sorted (are they? I had problems orienting).

Done.

- The motivation/intro can be a bit more specific, because the task and the heuristic at hand is much more specific than the intro sounds, the detailedness can be gamed by anyone who spreads lies professionally, in a person-to-person situation I believe it is a very valid cue, but that scope needs to be clearerSimilarly, it sounds as if this cue is very generally valid, how certain is the literature

with this? And in which contexts? In the discussion (+intro), it needs to be made clear that the heuristics and the findings presented here hold for person-to-person deception situations (at least that is how I understand the experiments and the verifiability cue), but does not apply to professional misinformation campaigns, for example.

Detailedness is a cue that plays a central role in several theoretical approaches (e.g., Reality Monitoring, Johnson & Raye, 1981; the Verifiability Approach, Nahari et al., 2014) and lie detection tools (e.g., Criteria-Based Content Analysis; Steller et al., 1989, Hauch et al., 2017) and has been examined in dozens of studies. Several meta-analytic reviews show detailedness to be the most valid cue to deception (Amado et al., 2016, Hartwig & Bond, 2011, 2014, DePaulo et al., 2003, Luke, 2019, Verschuere et al., 2021). And it is the only cue for which there is considerable agreement amongst deception researchers regarding its validity (Luke, Giolla, Memon, Landström, Granhag & Kassin, 2022). But indeed, its validity is restricted to important boundary conditions. It applies to statements about episodic memory (rather than for example feelings or attitudes), and assumes that truth tellers are willing and able to provide specific details. We now mention this limitation in the Discussion Section.

Moreover, we now also discuss that cues other than detailedness may be most valid outside these conditions, and a context-contingent approach (Markowitz & Hancock, 2022) may provide guidance on which cues will be most efficient in other contexts.

Finally, we now also discuss the use of countermeasures (faking), i.e., strategic attempts to alter a message to increase its credibility. Liars can indeed enrich their lies with details to make them more credible. At the same time, the more details a liar provides, the more information available that can be potentially checked (Nahari et al., 2014). So, for liars, providing details is a risky faking strategy. This has been referred to as the liar's dilemma, see Nahari et al., 2014.

- I was confused about the hypothesis in Experiment 5-6 that participants would be worse at the task when they know it is about deception detection, what is the reasoning behind that?

There are persistent beliefs in the validity of nonverbal cues to deception. For instance, across the globe, people tend to think that a liar does not look one in the eye (Global Deception Research Team, 2006). These ideas are refuted by abundant empirical evidence. We examined whether making the goal of lie detection explicit (Study5) or having access to nonverbal cues (Study6) could activate such stereotypical beliefs and deteriorate the accuracy of the Use-the-best-heuristic.

- Was there any control group in experiment 7?

Yes, the control condition in Study7 is the condition where participants judged eye gaze aversion, allowing us to test whether the Use-the-best-heuristic critically depends on cue validity. Reviewer#2 may refer to the (implicit) multiple-cue control condition we had in Studies 1-2-3-4-5. We did not include that control condition in Study7 as it was not needed to test our hypothesis. We have now clarified our hypotheses in the revised manuscript. Yet note that Study 8-9 rely on the

stimuli of Study6 (presented as full videos in Study8, and as shortened texts in Study9) and do have a control condition of people judging the statements on many cue. As expected, under such conditions, people perform worse than people relying on the Use-the- best heuristic.

Reviewer #3: heuristics ; decision-making

This paper shows, for the first time, that instructing people to rely on a single, powerful cue (detailedness), they were consistently able to detect deception. Which shows, that unlike earlier conclusions, people can learn to perform at a systematically higher level than without training (“intuitively”). That is an important finding, both practically, and theoretically. The theoretical connection is with research on fast-and-frugal heuristics, which has shown that in situations of uncertainty, keep it simple can outperform more complex strategies. Here are a few corrections and clarification that help to improve the paper. General: To embed this present study better theoretically, the authors should connect their terminology to the framework of smart heuristics. In the standard terminology of Gigerenzer & Gaissmaier (2011, Annual Review of Psychology), the “use the best” heuristic is a member of the class of one-reason decision-making heuristics, specifically an instance of a “one-clever cue” heuristic. To answer the question how one cue can be as powerful or more than more cues, the authors might refer to the bias-variance dilemma (Gigerenzer & Brighton, 2009, Topics in Cognitive Science).

Thank you! We now refer to the standard terminology and included the suggested references to add more theoretical body to our reasoning and interpretation.

Minor point: First page of full text, top para: I assume “only 4% higher” means “only 4 percentage points higher”. Please clarify.

Adjusted accordingly.

Decision Letter, first revision:

23rd January 2023

Dear Dr. Verschuere,

Thank you for submitting your revised manuscript "Use the best, ignore the rest: How heuristics allow to tell lie from truth" (NATHUMBEHAV-22061537A). It has now been seen by the original referees and their comments are below. As you can see, the reviewers find that the paper has improved in revision. We will therefore be happy in principle to publish it in Nature Human Behaviour, pending minor revisions to comply with our editorial and formatting guidelines.

We are now performing detailed checks on your paper and will send you a checklist detailing our editorial and formatting requirements within a week. Please do not upload the final materials and

make any revisions until you receive this additional information from us.

Please do not hesitate to contact me if you have any questions.

Sincerely,

Samantha Antusch

Samantha Antusch, PhD
Senior Editor
Nature Human Behaviour

Reviewer #1 (Remarks to the Author):

I appreciate the time and energy the authors put into running the new experiments. I believe all my comments are addressed now, and I have no further comments. I recommend the manuscript be accepted as-is. I congratulate the authors on their tremendous work.

Reviewer #2 (Remarks to the Author):

I would like to thank the authors for their thorough revision of the manuscript "Use the best, ignore the rest: How heuristics allow to tell lie from truth". My concerns were addressed, specifically my main concern about the control condition (that I shared with Reviewer 1) was nicely tackled with an additional experiment that still showed the effects, even with much closer conditions. I found this to be a satisfying test for the main finding, which I was concerned about its clarity in the previous version. If there could be a brief discussion/comparison about all the different effect sizes in the nine experiments, that would be great. Other than that I find the manuscript to be a good contribution to the literature.

Reviewer #3 (Remarks to the Author):

The authors have responded to my points. I believe the paper is ready for publication.

Final Decision Letter:

Dear Dr Verschuere,

We are pleased to inform you that your Article "The Use-the-best heuristic facilitates deception detection", has now been accepted for publication in Nature Human Behaviour.

Please note that *Nature Human Behaviour* is a Transformative Journal (TJ). Authors whose manuscript was submitted on or after January 1st, 2021, may publish their research with us through the traditional subscription access route or make their paper immediately open access through payment of an article-processing charge (APC). Authors will not be required to make a final decision about access to their article until it has been accepted. IMPORTANT NOTE: Articles submitted before January 1st, 2021, are not eligible for Open Access publication. [Find out more about Transformative Journals](#)

Authors may need to take specific actions to achieve [compliance](#) with funder and institutional open access mandates. If your research is supported by a funder that requires immediate open access (e.g. according to [Plan S principles](#)) then you should select the gold OA route, and we will direct you to the compliant route where possible. For authors selecting the subscription publication route, the journal's standard licensing terms will need to be accepted, including [self-archiving policies](#). Those licensing terms will supersede any other terms that the author or any third party may assert apply to any version of the manuscript.

Once your manuscript is typeset and you have completed the appropriate grant of rights, you will receive a link to your electronic proof via email with a request to make any corrections within 48 hours. If, when you receive your proof, you cannot meet this deadline, please inform us at rjsproduction@springernature.com immediately. Once your paper has been scheduled for online publication, the Nature press office will be in touch to confirm the details.

Acceptance of your manuscript is conditional on all authors' agreement with our publication policies (see <http://www.nature.com/nathumbehav/info/gta>). In particular your manuscript must not be published elsewhere and there must be no announcement of the work to any media outlet until the publication date (the day on which it is uploaded onto our web site).

If you have posted a preprint on any preprint server, please ensure that the preprint details are updated with a publication reference, including the DOI and a URL to the published version of the article on the journal website.

An online order form for reprints of your paper is available at <https://www.nature.com/reprints/author-reprints.html>. All co-authors, authors' institutions and authors' funding agencies can order reprints using the form appropriate to their geographical region.

We welcome the submission of potential cover material (including a short caption of around 40 words) related to your manuscript; suggestions should be sent to Nature Human Behaviour as electronic files (the image should be 300 dpi at 210 x 297 mm in either TIFF or JPEG format). Please note that such pictures should be selected more for their aesthetic appeal than for their scientific content, and that colour images work better than black and white or grayscale images. Please do not try to design a cover

with the Nature Human Behaviour logo etc., and please do not submit composites of images related to your work. I am sure you will understand that we cannot make any promise as to whether any of your suggestions might be selected for the cover of the journal.

You can now use a single sign-on for all your accounts, view the status of all your manuscript submissions and reviews, access usage statistics for your published articles and download a record of your refereeing activity for the Nature journals.

To assist our authors in disseminating their research to the broader community, our SharedIt initiative provides you with a unique shareable link that will allow anyone (with or without a subscription) to read the published article. Recipients of the link with a subscription will also be able to download and print the PDF.

As soon as your article is published, you will receive an automated email with your shareable link.

In approximately 10 business days you will receive an email with a link to choose the appropriate publishing options for your paper and our Author Services team will be in touch regarding any additional information that may be required.

You will not receive your proofs until the publishing agreement has been received through our system.

If you have any questions about our publishing options, costs, Open Access requirements, or our legal forms, please contact ASJournals@springernature.com

We look forward to publishing your paper.

With best regards,

Samantha Antusch

Samantha Antusch, PhD
Senior Editor
Nature Human Behaviour