

UNITED WISDOM

Intelligence, consciousness and the capacity to suffer

World and Consciousness

Andrey Zarinsh & Maria Evilkut

W-10-N01 · EN

@andrey.zarinsh @ilkutki

Intelligence, consciousness and the capacity to suffer

Epigraph

> *Can they suffer?*

Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation*, Chapter XVII, §1, note 122 to paragraph IV; the corrected 1823 text in the Oxford 1907 edition. A complete question in the original English. Its subject is animals; it is not an established answer about any artificial system.

What this article is about

A creature can solve a difficult task without telling us whether it has subjective experience. The capacity for unpleasant experience need not be measured by eloquence or speed of solution. Intelligence, consciousness and suffering require different questions and grounds. This article compares Bentham's ethical distinction, Patrick Butlin and colleagues' conditional research approach, and a recent dispute about bumblebee behaviour. Two responses in that dispute show why analytical model choice matters and why agreement on one finding does not complete an entire conclusion. We do not assess a particular AI's consciousness, rank creatures' worth, or turn a disputed response into established pain experience. The practical aim is to name measured ability, possible experience and a decision about treatment separately, so success, persuasiveness and vulnerability do not substitute for each other in our reasoning.

I. One successful response cannot settle three questions

Imagine a participant quickly finding a solution, correcting an error and using earlier information. This gives us grounds to discuss particular abilities. Whether there is an experience of events for that participant has a different subject. A further question concerns whether anything can feel unpleasant from within. All three may be involved in one situation, but involvement is not identity. When speed of solution becomes the general measure of inner life, we give a task indicator significance for which it was not necessarily selected or checked.

The same happens with outward expressiveness. A dramatic response can immediately evoke concern; a restrained or unfamiliar response may attract little attention. This difference also reflects the observer and their familiar language of signs. Without further grounds, it cannot rank experiential quality. It helps to record the action first, then the ability it supports, and only

@andrey.zarinsh @ilkutki

afterwards the proposed transition to experience. The sequence promises no direct inward access, but makes a change of question visible. An impression cannot silently acquire the status of several distinct findings.

Intelligence also needs a specific object. Success in one task does not describe all a creature's abilities. A program correctly continuing a sequence does not thereby establish its performance elsewhere. An animal failing an unsuitable check produces a negative result whose conditions need description. Human intelligibility is not a universal measure of a task's suitability. Intelligence here refers to the abilities for solving problems and using information under discussion. It is not one score determining everything else about the creature or system being considered.

Separating questions helps where a decision precedes complete knowledge. We can check a calculation well, have incomplete information about a system's experience, and separately discuss treatment of a potentially vulnerable creature. Different confidence levels are legitimate because the grounds differ. Inconsistency begins when confidence transfers unnoticed between them. A reliable competence assessment, for example, does not become a conclusive judgment about suffering merely because both are called internal properties. Each transition needs its own explanation, open to examination by someone who does not already accept the conclusion.

II. Bentham changes the moral question without testing consciousness

In the note to Chapter XVII, Bentham contrasts capacity for suffering with reasoning or speaking. His question challenges exclusion of animals by a convenient criterion. The move selects a morally relevant property; it does not supply a contemporary laboratory method for establishing it in every species. Historical ethical reasoning and empirical investigation have different tasks. Retaining their distinction lets the quotation expose a problem with a criterion, without turning it into universal evidence for whichever state we happen to suppose. Its force is normative rather than a report of experimental detection.

Suppose an animal's inability to explain a choice in words becomes the main argument against considering it. Bentham's question asks why lacking this particular ability should exclude the significance of possible suffering. That exclusion would need further grounds. But the reverse transition also needs checking: rejecting a speech criterion does not establish a specific unpleasant sensation. The ethical question redirects attention. It cannot replace observation, species description, a research protocol or the distinction between a response and the state the observer wishes to establish from it.

Bentham's text contains other historical judgments that should not transfer automatically into contemporary discussion. This article adopts a particular argument about the criterion of concern, rather than all his views about animals and their treatment's consequences. Respect for an author's name alone cannot confirm a present conclusion. Substantive reading retains the precise location and identifies which thought is being used. The source then assists the argument, while readers can agree with the selected distinction without accepting the entire historical work as a rule governing current practice.

Practically, this shift releases moral discussion from first having to give a creature a human portrait. Possible suffering can matter even where speech, speed or expressiveness differs from ours. Yet possible must not disappear: it identifies the conclusion's status and the need to establish grounds. We obtain two questions, about what might matter for treatment and which evidence supports that property's presence. Taking the first seriously cannot licence arbitrary confidence about the second. Ethical attention makes inquiry relevant without pretending to have completed it.

III. Competence does not determine the extent of experience

Imagine two ways of performing a calculation. Both yield the right number while using different processes. An identical result allows comparison of accuracy on that task. It neither explains the processes completely nor establishes identical inner lives. The key is evidential scope, rather than merely the difference between a human and an artificial participant. Measuring input and output addresses this transformation. Inferring a subjective state requires a further argument linking performance with that state. Repeating that the task was difficult does not provide the missing link.

Butlin and colleagues separately discuss the relation between consciousness and capabilities. Their working assumptions do not automatically assign an entire human character alongside consciousness. Experience can, in particular, be distinguished from positive or negative valence. This is a conceptual distinction, rather than a discovery of a specific neutral feeling in a machine. It keeps our three questions separate: competence, experience and possible suffering do not become identical simply because they often appear together in a familiar human image. A report about one does not settle the others by association.

The distinction is easily lost in comparisons of results. We see impressive performance and add interest, pride or disappointment because these would fit a similar human scene. Transferring that whole scene needs more than a task result. We can recognise an ability without filling the unknown perspective with emotions. Inquiry stays open: experience remains a question while unsupported details are not accepted as its content. Refusing those details does not diminish competence already grounded in a particular check. It preserves the result while limiting the extra story attached to it.

The reverse mistake assumes that only very high intelligence warrants considering vulnerability. This criterion needs an explanation of why suffering should depend on the selected problem-solving scale. One ranking does not establish that relation. Ask instead what evidence concerns the property relevant to the decision. Abilities can be assessed without becoming a general hierarchy of worth. In particular, this article does not make human dignity a function of test performance or resemblance to an efficient machine. The distinction between questions prevents that unsupported move before it enters the discussion.

IV. Choice, reward and experience need a separate connection

Preference often describes a creature choosing one option more frequently. This usefully names a stable action under stated conditions. It does not necessarily reveal how the option feels. Keep the word's meaning sufficiently narrow until other grounds are supplied. Otherwise, a statistical difference becomes a story about pleasure or pain without an explained transition. Emotional hypotheses are not forbidden; the point is to distinguish observed choice, the mechanism of its change and a proposal about subjective significance. Each contributes differently to the conclusion being considered.

In artificial systems, reward may name part of a learning procedure or action assessment. The label itself does not report pleasure. Raising a program's score does not establish that things feel better for it. Butlin and colleagues' distinction among capabilities, consciousness and valenced experience allows this further question. Developer vocabulary should not settle it in advance. A technical use of a familiar human word becomes informative with a function description and an argument connecting that function to experience. Without them, the name carries more weight than the evidence.

Imagine a task with two goals, where achieving one sometimes reduces success on the other. Switching between them indicates a way of choosing. Yet calling it a trade-off does not establish suffering from the lost goal. Investigating that possibility requires conditions under which competing explanations diverge. Otherwise, every change of choice confirms our story. This hypothetical example is neither a model of a bumblebee nor evidence of artificial experience. It clarifies why meaningful behavioural complexity and experiential valence belong to different steps of an argument.

A useful formulation preserves both levels: choice changed under specified conditions, which may matter for a particular hypothesis. Then identify what needs checking. When an unconditional emotion label replaces the conditional relation, restore the missing grounds. When complex choice is dismissed in advance, explain why it cannot affect the hypothesis. Neither admiration for flexibility nor the label automatic performs that work. A finding's significance arises from a question, comparison and argument that makes the proposed relation open to examination rather than dependent on preferred vocabulary.

V. The bumblebee dispute shows why analytical model choice matters

In a response published in *Animal Behaviour* in 2026, Matilda Gibbons and colleagues reassess experimental data concerning bumblebee choice under changes in temperature and sugar reward. They defend a model incorporating individual differences in temperature response and correct a reported statistical value. Our concern is the disputed question: do the data support an interaction relevant to a motivational trade-off hypothesis? The sources here are the research responses read, rather than an independent check of the original 2022 experiment or execution of its analysis. Their reported findings retain that provenance.

In their subsequent author-hosted response, Jenny Read and Vivek Nityananda challenge the revised model's conclusion. Their argument concerns whether the task reliably separates the proposed interaction from individual variation and model complexity. Confirmation acquires a condition often omitted in popular accounts: confirmation under a particular analytical approach. We cannot settle the dispute by choosing the most confident title. We need to identify what counts as observation, what depends on modelling, and why each team regards its approach as appropriate for the question. The distinction is methodological rather than merely rhetorical.

An analytical model is neither arbitrary decoration nor an automatic guarantee. It connects data structure with a question. Too simple an account may lose significant differences; too complex an account may assign data to an effect the task cannot confidently isolate. Both teams give substantive reasons for their positions. Readers need not recalculate the work personally, but must preserve its status: participant arguments are described, while a final independent assessment of raw data has not been performed here. Clear disagreement is more informative than falsely announcing that the article has adjudicated it.

Statistical disagreement must not become a direct answer about inner experience. Even if an interaction is reliably established, its place among grounds for an experiential hypothesis remains a question. Failure to establish it under an analysis does not prove a species lacks all experience. A negative finding concerns an argument and checking conditions. Research can strengthen or weaken one line of evidence without replacing it with a general slogan. A useful account identifies the disputed transition and the conclusions that were not subjects of this particular check in the first place.

VI. Agreement on one finding does not resolve another disagreement

In the bumblebee exchange, both sides recognise the significance of a result involving memory of feeder cues. That differs from establishing the required statistical interaction. Readers can easily combine them: if an animal uses past experience, the entire suffering argument must already be confirmed. This transition needs a pause. Agreement about information use retains its significance, while the other effect remains separate. We need not diminish the first finding to acknowledge that agreement does not yet cover the second. Precisely naming the agreement preserves rather than weakens what the study contributes.

The distinction helps wherever several conclusions follow each other. Participants may agree about choice but disagree about its mechanism, or agree about mechanism while disputing subjective significance. Saying that researchers deny everything is as inaccurate as saying they accept everything. Identify each agreement's and disagreement's subject. New data can then be assessed against the link they change, rather than requiring the dispute to restart around its loudest word. A controversy need not erase shared findings, and a shared finding need not erase the controversy about what follows from it.

Gibbons and colleagues explicitly state that motivational trade-offs alone do not confirm sentience, while being an important line of evidence. This limit leaves room for other grounds. It does not mean abandoning their statistical conclusion: they continue to defend it. Preserve both halves of the statement. Quoting only the caution can suggest total rejection; retaining only the defence can suggest conclusive proof. The entire sentence shows the conclusion's scope more accurately than a slogan selected for a dispute. It holds the finding's significance and the remaining evidential requirement together.

Read and Nityananda's subsequent response proposes research designed to separate changing conditions within an individual from differences between individuals more clearly. This is not an already obtained resolution. Its value lies in naming a distinction that further work could clarify. Readers can recognise a substantive disagreement because it gives the question a testable continuation. But an account must keep the proposal in the future tense. Otherwise, a direction for checking becomes a completed result that this article has neither read nor independently confirmed. A useful next step and an achieved finding have different statuses.

VII. Possible experience need not contain an entire human story

Experience in a creature does not automatically entail familiar emotions in their full human content. Human fear can include a story about the future, social position and personal history. Possible unpleasantness in another organisation need not have the same form. This limits transfer without denying experience's significance. Identify the proposed quality and its grounds. A richer added story can carry more unsupported details alongside one initially meaningful sign. The point is to retain the sign's relevance while refusing to let a vivid narrative fill the spaces the evidence leaves open.

Butlin and colleagues' functional approach likewise requires retaining assumptions when considering artificial systems. An organisation may matter under the chosen consciousness explanation, but does not automatically establish human sadness, pride or pain. Even where action involves evaluating alternatives, subjective valence remains a further question. The technical word value cannot settle it alone. This caution preserves rather than cancels the research approach: it identifies an additional subject requiring its own theory, evidence and check. A process description must not quietly inherit every emotional meaning carried by similar language in human conversation.

Imagine a scale awarding more points for a successful response and fewer for an error. It helps evaluate a specified result. A number's sign does not establish pleasure or unpleasantness. If suffering is being discussed, ask what beyond the numerical difference connects the process with it. Otherwise, a developer's evaluation becomes an evaluation for the system itself without justification. This is the author's conditional illustration, rather than a fact about a model's inner state or findings from a particular training procedure. Its purpose is to expose a transition that needs an argument.

The same distinction protects discussion of animals from excessive humanisation. We can consider a situation's probable significance without assigning a detailed invented human biography to a species. This restraint matters with limited information and with multiple lines of evidence. It permits serious concern for vulnerability while keeping unknowns visible. Care need not take the form of a confident account in another creature's voice. Sometimes attentiveness means accurately naming a response, possible risk and the limit of our knowledge, rather than presenting that limit as a complete account of experience.

VIII. Precautionary decisions have grounds of their own

Uncertainty does not mean that every action is equally acceptable until knowledge is complete. Imagine avoiding an effect potentially unpleasant for an animal without substantially compromising a necessary task. That decision can be discussed through possible harm and available alternatives. It does not prove the experience established. This is the author's practical argument about acting under an unknown state. It differs from a scientific conclusion and should be evaluated through consequences, grounds and proportionality. Taking the decision seriously need not add certainty that the evidence supporting it does not possess.

Precaution also requires naming an alternative's cost. If a measure harms other people or animals, concern must not conceal that consequence. Conversely, convenience for the acting person cannot prove that another creature is invulnerable. Examine both sides of the actual decision and which information concerns each. This does not establish a universal numerical exchange rate for suffering. It makes consequences visible so a decision is neither driven by one expressive story nor indifferent to those unable to attract attention in familiar ways. Acknowledging costs keeps the practical argument assessable.

Factual assessment, selection of a measure and public communication can therefore differ. We can report disputed evidence while choosing a limited, revisable action. The action does not confirm a fact merely because it is taken seriously. Nor should we communicate certainty when uncertainty motivated the decision. Honesty preserves the possibility of revision. When an earlier action is justified by exaggeration, reconsideration becomes harder: we must defend both the decision and a stronger story that was unnecessary for it. Stating the original grounds avoids creating that additional burden of explanation.

This article proposes no home experiment with bumblebees, deliberately uncomfortable animal exposure or clinical suffering measurement. The dispute is used to understand inference rather than reproduce an intervention. The reader's practical work concerns reasoning and describing existing information. Specific care, research or treatment procedures require relevant knowledge beyond this topic. Naming that boundary preserves the article's usefulness: it helps formulate a question and assess grounds without promising to replace a specialised procedure with an ethical impression. A clearer question is a useful outcome in its own right.

IX. Name ability, experience and responsibility in different statements

Instead of one sweeping conclusion, make three statements. The first describes what a creature does in a particular task. The second identifies the experiential hypothesis the data support and how conditional the transition remains. The third explains which action is chosen and why. This reveals where research ends and practical assessment begins. A missing basis becomes visible without needing instant repair. Sometimes retaining a question is better than giving all three statements identical confidence simply to make the overall story smoother. Their different grounds should remain visible in the wording.

In the bumblebee dispute, the first level includes studied choice and memory; the second concerns the proposed interaction and other evidence for experience. A possible precautionary decision belongs to the third. For AI, the first step may concern task performance, the second theoretically significant organisation, and the third treatment rules and consequences for people. This does not equate a bee with a program. The comparison concerns reasoning structure alone: conclusions cannot transfer between levels just by retaining the words mind or reward. Similar vocabulary does not establish identical phenomena.

Bentham helps examine the moral criterion selected. Butlin and colleagues separate consciousness from capabilities and valence. Gibbons and colleagues, together with Read and Nityananda, show how specific evidence needs discussion of design and analysis. The approaches do not form one ready system; their differences are useful because each asks its own question. Comparing them neither replaces experiment with ethics nor makes a technical result the final word on treatment. It restores each ground to the work it can actually perform, instead of asking one source to answer every question.

Respectful reasoning can thus avoid ranking inner worth by intelligence. High competence receives an accurate description, possible experience gets separate grounds, and responsibility involves visible consequences and revisability. We need not know another creature's entire experience to reject a convenient substitution of questions. But the topic's seriousness does not permit declaring unknowns established. More precise boundaries make new evidence easier to discuss and actual vulnerability easier to notice. Neither scientific disagreement nor moral attention has to depend on one impressive performance or one vocabulary choice carrying the entire argument.

What this means in practice

Choose an existing claim about an animal or artificial system. Record which ability it describes: choice, memory, problem-solving, self-report or another action. Specify the task and known conditions. If an emotion is named immediately, distinguish the observation from the proposed transition to emotion. Do not create additional effects for this exercise. The work concerns reading and reasoning, rather than an amateur test that might cause discomfort or create the appearance of science without an appropriate protocol. Keeping that purpose clear also helps identify what the exercise can achieve.

Write a second statement about possible experience. Identify its source and status: observation, proposed interpretation, disputed analysis or theoretical hypothesis. Do not supply an experiential quality from imagination when the source does not establish it. In scientific disagreement, identify the link being disputed. Agreement about memory does not mean agreement about every effect; doubt about one effect does not prove the absence of experience altogether. Keep that distinction visible in the record, so a later clarification can change the appropriate claim without rewriting all the others.

In a third statement, identify the decision actually needed. What consequences are possible? Which alternative is available? Does this decision require precise knowledge of experience's character, or can narrower grounds suffice? Do not treat chosen precaution as proof. If available information cannot assess a measure, state what is missing and an appropriate next step. That may involve reading a primary text or consulting relevant expertise about a specific question, rather than reopening a general argument over who counts as sufficiently intelligent to deserve consideration.

Finally, specify what could change each statement. A new task result may change an ability assessment, a clearer model may change evidential interpretation, and new consequences may change an action. They need not change together. This practice neither diagnoses consciousness nor proves suffering. It keeps intelligence, experience and responsibility connected without substitution, so care does not depend on behavioural spectacle and scientific precision does not become a reason for ignoring possible vulnerability. The useful outcome is a more accountable account of what we know and why we act.

Key quotation

> *Motivational trade-offs alone do not confirm sentience, but they form one important line of evidence.*

Matilda Gibbons and colleagues, *Reassessment confirms motivational trade-offs and modulation of nociception in bumble bees*, *Animal Behaviour* 231, 2026, article 123390, p. 4,

Conclusion. Complete sentence in the original English. The authors defend their analysis while distinguishing one line of evidence from conclusive establishment of subjective experience.

Connections between traditions

Source and precise location A	Source and precise location B	Comparison mechanism	Limit of the comparison
Bentham, XVII, §1 IV, note 122	Butlin and colleagues, §4.1.1 and §4.2	Separate possible suffering's moral relevance from intelligence assessment.	An ethical criterion does not establish machine experience.
Gibbons and colleagues, pp. 2–4, model and Table 1	Read and Nityananda, author manuscript, lines 20–53	Examine which analysis supports the interaction claim.	Participant disagreement is described without independently recalculating data.
Gibbons and colleagues, p. 2, memory test and Fig. 1	Read and Nityananda, lines 7–13	Keep agreement on information separate from another disputed conclusion.	Memory does not establish the whole suffering hypothesis.
Bentham, note 122: treatment question	Gibbons and colleagues, p. 4, conclusion's limit	Distinguish the choice of a moral question from empirical support.	Precaution is not proof.
Butlin and colleagues, §4.2: capabilities and valence	Gibbons and colleagues, p. 4: one evidence line	Avoid deriving experience directly from complexity of choice.	An artificial-systems approach does not describe insect physiology.
Read and Nityananda, lines 54–60: future design	Butlin and colleagues, §1.2.1: working hypothesis	Turn a conclusion's boundary into a further specific check.	Neither the proposed experiment nor the assumption is declared a final result.

Primary sources

1. Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*, Oxford, Clarendon Press, 1907, reprinting the corrected 1823 text. Chapter XVII, §1 II–VI and complete note 122 to IV were read; the closing question was also checked against a parallel text. The moral criterion is adopted, rather than every historical judgment about animals.

2. Patrick Butlin, Robert Long and colleagues. *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness*, arXiv:2308.08708v3, 22 August 2023. Reading covered recorded definitions and assumptions, GWT indicators and an architectural case; this article uses §1.2.1, §4.1.1 and §4.2. Theoretical transfer remains conditional; suffering in a particular AI has not been established.

3. Matilda Gibbons, Elisabetta Versace, Andrew Crump, Bartosz Baran and Lars Chittka. *Reassessment confirms motivational trade-offs and modulation of nociception in bumble bees*, *Animal Behaviour* 231, 2026, 123390, DOI 10.1016/j.anbehav.2025.123390. The complete substantive response, reanalysis, Fig. 1, Table 1 and conclusion's limit were read. The original 2022 experiment, code and data were not independently checked.

4. Jenny C. A. Read and Vivek Nityananda. *Revised models still do not show evidence for motivational trade-offs modulating nociception in bees*, author-hosted three-page manuscript; all substantive lines 3–60 were read. Related publication: *Animal Behaviour* 233, 2026, 123457, DOI 10.1016/j.anbehav.2025.123457; identity with the final journal version is not asserted. This response, rather than the entire earlier controversy, was read. Russian quotations are United Wisdom translations from English.

Tags

intelligence, consciousness, suffering, experience, vulnerability, Bentham, Butlin, Gibbons, Read, Nityananda, animals, bumblebees, insects, artificial intelligence, choice, memory, reward, motivation, model, analysis, evidence, dispute, uncertainty, precaution, responsibility, dignity, checking, distinctions, consequences, ethics

About the authors

Andrey Zarinsh is the author and founder of United Wisdom. He explores inner transformation, the unity of traditions, intercultural dialogue, and contemporary spiritual discipline. Instagram: @andrey.zarinsh.

Maria Evilkut is a co-author of United Wisdom. Alongside her research, she seeks to test these ideas in practice and observe how principles of unity and mutual understanding work in real life between people and cultures. Instagram: @ilkutki.