

Normative Detachment Theory in Context

Literature Review 0.15

Bree Beal · September 2026 · Reader companion

The account of NDT follows Formal Specification 1.6.5 and Scope and Commitments 1.18. Appendix A identifies all 48 records in Comparison Dataset 1.20, with their retained classifications and relevant source qualifications. Full publication entries are collected in Master Bibliography 1.10.

1. The answer and the thought that produces it

A parent is preparing breakfast. Before pouring milk over a child's cereal, the parent notices a sour smell, believes the milk has spoiled, and decides to discard it. The child will receive something else. The decision is sensible, but knowing the decision does not yet tell us how the parent reached it.

Suppose the child's well-being helps make serving the milk unacceptable. The child matters to the parent in the child's own right; their shared life connects the child's needs to the parent's response; and the parent can prevent the child from drinking the milk. Those features of the represented relationship help organize the judgment. The parent need not silently recite them. An immediate decision can have a structured explanation.

Now imagine the same parent, breakfast, spoilage belief, options, and conclusion, but a different contributing basis. The parent applies a food-handling rule learned at work: discard milk showing signs of spoilage. Applying the rule at home produces the discard judgment. Rule use alone would not distinguish the episodes: a rule can express care. For the contrast, therefore, suppose the rule supplies the entire basis of the second judgment, without the relationship to the child or another moral relationship helping form it. The parent may still love the child. A relationship can remain important in someone's life without contributing to every thought.

NDT proposes that the difference between those two ways of forming a judgment matters to its cognitive type. The question is not which parent is better, whether breakfast is safe, or whether following a rule deserves less moral credit. The question concerns a particular occurrence of judgment: does a represented moral relationship help organize the evaluation used in forming it? NDT calls the contributing source and organization of that occurrence its *provenance*.

The imagined comparison makes a conceptual distinction available. It is not a behavioral experiment, and the identical act of discarding milk could not establish which imagined explanation describes a real parent. The distinction between an answer and its contributing basis must first be understood before asking what evidence could identify either basis in an actual episode.

Many intellectual traditions address neighboring questions. Some explain how means follow from ends, some distinguish morally important concerns, some explain judgments of harm, and some investigate the reasons for which an attitude is held. A useful comparison asks what each account contributes to the breakfast case and where the explanatory task changes. Matching the conclusion is one achievement. Explaining the contributing organization is another. Identifying a judgment as moral is another again.


One breakfast, two contributing bases

Kept the same: parent, child, spoilage belief, options, and discard conclusion.

Relationship-supported judgment



Procedure-only judgment, as stipulated



The parent may love the child in both episodes.

Routine does not establish the absent contribution; the comparison stipulates it.

Figure 1. Two stipulated ways of reaching the discard judgment. In one, the represented relationship helps form the evaluation; in the other, the rule supplies the basis without a moral relationship contributing. The diagram explains the proposed distinction, not a measured contrast between actual parents.

2. From a represented situation to practical guidance

Before asking what makes the breakfast judgment moral, consider what makes it practical. A sour smell by itself is not an instruction to discard anything. The parent understands the smell as evidence of spoilage, connects spoilage with possible illness, and treats serving the milk as unacceptable or discarding it as preferable. The judgment also depends on alternatives: the milk has not yet been served, and another breakfast is possible. Remove the relevant evaluation or change the available response and the inference changes.

Practical-inference traditions make part of that structure explicit. Aristotle, Anscombe, and Davidson examine relations among an end, beliefs about means, and the practical conclusion. Their accounts do not all identify the conclusion in the same way: action and intention have roles that NDT's narrower focus on an occurrence of judgment does not replace. The bounded correspondence is guidance from represented ends and means, not an identity between judgment, intention, and successful action. A parent can judge that the milk should be discarded and still fail to discard it.

Hume introduces a different boundary. The famous is–ought passage asks that a transition to an ought relation be noticed and explained. NDT does not derive a prescription from inputs stipulated to lack every evaluative ordering, goal, requirement, or commitment. Its distinctive question is partly about the inputs: must evaluation always be supplied by a separately represented desire or verdict, or can evaluation belong to what makes a representation function as the representation it is?

The smell of spoilage provides an entrance into the question without answering it. A person might first represent a neutral chemical condition and then apply a separate concern about health. Alternatively, a representation might owe part of its functional identity to its place in an organism's stake-sensitive control: what the organism selects, maintains, avoids losing, or restores. NDT permits the second possibility as encoded evaluation. The claim is stronger than saying that information sometimes affects behavior, but narrower than saying that every representation is evaluative or that every evaluation determines a specific action.

How can anything matter to an organism rather than merely affect it? Di Paolo's enactive account connects significance to an organism's self-maintaining organization and its capacity to regulate conditions that threaten continued viability. Mere self-production is insufficient: the organism must also distinguish and respond to tendencies toward or away from its viability boundary. Di Paolo also sketches how values could arise beyond immediate metabolic needs. NDT's encoded evaluation addresses a related question at the level of a representation's functional role. The comparison concerns how significance belongs to cognitive organization; it does not identify Di Paolo's adaptivity with NDT's normative detachment or its criteria for moral judgment. [Di Paolo \(2005\), pp. 436–441, 445–448.](#)

Three relations must therefore remain distinct. A representation can succeed in representing its object. Outcomes can be ranked as better or worse within the person's practical organization. Available actions can be evaluated through their expected consequences and applicable requirements. Accurate temperature information, for example, does not alone settle whether to warm a patient, preserve a sample, or recalibrate an instrument. Moving between those questions requires represented connections. NDT supplies those connections as part of a decision model: the person's organized representation of the practical problem, not a second thinker inside the person.

Whether encoded control counts as passion in a sufficiently broad Humean psychology remains a separate interpretive issue. Hume's discussion includes calm passions; the familiar contrast between reason and passion does not independently settle every modern functional decomposition. The present comparison therefore does not announce a refutation of Hume or classify all Humean accounts together. The unresolved object is the proposed separation between representation and the source of its evaluative role.

Evaluation can also enter through rules and institutions. A promise may be understood as creating an obligation because the speaker and listener recognize an institution and its operative conditions. Searle's promise case makes that institutional bridge salient. Input/output logic supplies a different resource: a way to connect conditions with normative outputs without reducing the governing norm to an ordinary material conditional. Neither resource alone decides whether a particular use of the promise or rule is moral in NDT's sense. The parent's remembered food-handling rule remains the useful reminder: a rule's content, institutional legitimacy, or historical moral purpose does not establish the contributing basis of every current application.

Source notes: practical inference—Aristotle, *De Motu Animalium*, part 7; Anscombe, *Intention*, §33; Davidson, “Intending” (1978, reprinted 2006), p. 132 of *The Essential Davidson*. Hume—*Treatise* 3.1.1.27, SBN 469–470; motivation and calm passions—2.3.3.3 and 2.3.3.8. Institutional and rule transformations—[Searle \(1964\)](#); [Makinson and van der Torre \(2007\)](#). NDT—Formal Specification §§0.2–0.3, 2–3.

3. What the represented relationship adds

The two breakfast explanations differ within a common account of normativity. Both can produce guidance. NDT's additional proposal is that moral cognition involves a represented subject in relation to an object, with the relationship helping shape the operative evaluation. “Object” names a role in the explanation, not a denial of personhood: the child is the object of concern, while the parent is the represented subject who can respond.

The relationship contains three distinguishable questions. Does the child matter in the child's own right, beyond instrumental usefulness? How is the child situated in relation to the parent? What response does the parent understand as possible? In the breakfast case, positive inherent value, intimacy, and an available protective response make the questions easy to answer. The theoretical decomposition becomes useful when their answers come apart.

A person can regard distant people as valuable without representing those particular people as connected to the person's own response. A professional role can connect a physician to a patient without intimacy. A stranger's testimony can bring a particular need directly into someone's practical understanding without a pre-existing shared identity. NDT names the joining role *Relational Positioning*. Its current families are identity-mediated positioning, intimacy, and direct encounter. Identity-mediated positioning can involve a shared identity, a contrastive identity, or complementary roles. Bare membership in a category or possession of a job title is insufficient: the represented connection must do the relevant work.

Nor is the value question exhausted by benevolence. NDT keeps positive and negative inherent value separate. Someone can represent an object as good in itself, bad in itself, both, or neither. A political opponent need not be represented as inherently bad; contrastive identity and negative value answer different questions. Even represented badness does not alone yield a duty to punish, destroy, or avoid. Those responses require additional represented connections. The theory describes how a judgment is formed without endorsing the values or consequences the person represents.

Agency supplies another boundary. It concerns a represented possibility of responding to the object or relationship, not merely awareness of the object. The relevant response may be immediate, future, reparative, testimonial, or expressive within the declared horizon. Attributing moral subjecthood to someone else adds a different question: does the observer represent that other agent as capable of recognizing and integrating the relevant value, relations, and possibilities of response? A human label does not automatically answer yes, and a nonhuman label does not automatically answer no. Blame then requires further answerability and responsibility conditions.

Finally, all those represented features must belong to the judgment being explained. It would not suffice to find the parent–child relationship somewhere in memory while an unrelated rule alone forms the discard judgment. The relationship, the represented responder, and the particular object must contribute

together to the evaluation. These requirements make NDT's moral relationship an organizational proposal, not a label for every social tie or a synonym for concern.

The account invites comparisons with content theories, relational theories, and measures of moral concern. Those comparisons should now ask which of the newly separated questions the external construct answers. A familiar moral theme can be relevant without supplying every relation that NDT distinguishes.

Source note: the NDT account is specified in Formal Specification §§4a–4b and Scope §§7–8. The preceding cases explain the current decomposition; they are not independent demonstrations that the proposed variables have been measured or that rival decompositions fail.

4. Moral Foundations Theory: plurality, learning, and the organization of judgment

4.1 What the foundations are meant to explain

Why do moral concerns exhibit recurrent patterns while differing so markedly across people and cultures? Moral Foundations Theory approaches the problem through evolved preparations for learning and culturally developed moralities. Haidt and Joseph's early account connects prepared intuitions with culturally variable virtues; it does not describe adult moral life as a finished collection of rules supplied at birth. Graham and colleagues' 2013 synthesis integrates nativism, cultural learning, intuitionism, and pluralism. The number of foundations remains open, and substantial innate organization is defended alongside experiential development. The permissive discussion of modularity does not erase the nativist commitment. [Haidt and Joseph \(2004\)](#); [Graham et al. \(2013\)](#), pp. 58, 61–65, 99–104.

The importance of the developmental account is easiest to see in the breakfast example. Recognizing a threat to a child's health can involve learned circumstances, readily elicited concern, and a rapid response. NDT need not exclude such processing. Its prepared-module comparison concerns a possible source of content and response organization within a judgment-producing process. The comparison does not establish that a prepared module is necessary for every NDT judgment or that its activation suffices to make the judgment moral.

The familiar five-foundation questionnaire has a different job. Graham and colleagues' 2011 article studies five correlated dimensions, compares hierarchical alternatives, and leaves revision of the inventory open. The article also adopts a functional account of moral systems concerned with regulating selfishness and enabling social life, rather than simply defining morality by a content list. An early questionnaire's five dimensions should therefore not become a claim that mature MFT admits exactly five kinds of moral concern. [Graham et al. \(2011\)](#), pp. 368, 372–373, 382.

Later development changes the comparison further. Atari and colleagues distinguish equality from proportionality in a six-foundation instrument and examine culturally varying networks of association. Their six dimensions are care, equality, proportionality, loyalty, authority, and purity; the sixth is not simply the addition of liberty to the earlier five. The network edges are statistical relations, not demonstrated causal priorities within an individual judgment. Additional foundations remain possible. [Atari et al. \(2023\)](#), “Gaps in Theory and Measurement,” *Online First* pp. 3–5, 9, 15–16, 26.

4.2 The structural comparison after the count is left open

An open inventory does not answer every question about organization. Researchers can revise a taxonomy over decades while a psychological theory makes a distinct proposal about what one person can construct during an episode. Conversely, a theory can have fixed structural commitments while allowing changing contents. NDT itself specifies three builder roles and a governed set of positioning families. Calling one account fixed and the other dynamic without identifying the variable would compare different levels.

The precise NDT question is how a concern enters the current judgment. In the first breakfast episode, the child's represented standing and relation to the parent help make serving the milk unacceptable. In the second stipulated episode, the remembered rule produces the same guidance without such contribution. A foundation classification can identify related content in the two episodes while NDT distinguishes their contributing organization. Whether a foundation-based account can recover that contrast depends on more than assigning both cases to care. It needs a corresponding account of the different bases and their role in cognitive classification.

Hierarchy also needs a named relation. A distinction between general and specific concerns is not the same as a priority between two duties. A statistical hierarchy among questionnaire factors is not the same as a dependency by which one represented relationship constructs a requirement. NDT connects the moral subset to broader normativity and connects a particular requirement to the relationship that helps produce it. Those are the relations to compare. Neither the presence nor the absence of a branching factor diagram settles their equivalence.

The remaining challenge is therefore not "MFT has no structure." It is whether foundation-linked classification explains the token-level boundary NDT proposes, and whether the developed account supplies the dependencies needed for that explanation. A five-foundation partition and a provenance partition become empirical competitors only when their units and explanatory tasks are aligned. The fact that one groups content while another groups judgment occurrences is not itself a result favoring either. Comparable evidence would require separately specified classifications and a common basis of comparison, not a declaration that entire theories contradict one another.

4.3 Moral identification and evolved function

The same precision matters to the charge of circularity. Selecting familiar moral judgments for research does not, by itself, explain what makes those judgments moral. But selecting examples of an explanandum is not automatically a circular argument either. The question is whether an account uses the moral classification as a premise in the very explanation that is supposed to establish that classification.

Graham and colleagues assess candidate foundations using several criteria, including third-party normative judgment, affect, cultural distribution, preparedness, and an adaptive history. Their discussion also acknowledges relational context. Those commitments must be represented before evaluating a charge that the theory treats isolated content as intrinsically moral. [Graham et al. \(2013\), pp. 106–114.](#)

NDT's criticism can then be stated without a caricature. A recurring content family, an evolutionary explanation, or a social-system function does not yet supply the proposed distinction between the parent's two occurrences of judgment. A bridge from those claims to the contributing organization of an occurrence remains to be specified. The omission is relevant to NDT's explanatory problem; it does not

establish that every foundation-based extension is impossible, or that an evolved function requires every present act to benefit cooperation.

The mature encounter leaves two conclusions together. Foundation-linked learning and processing are possible resources for NDT, while a foundation taxonomy and a provenance criterion are different proposed ways of organizing the subject matter. The source development sharpens the remaining challenge rather than dissolving it into either whole-theory agreement or a contest over the correct number of foundations.

Compare the same explanatory job

Development and learning

What is prepared, and what is learned?

MFT: evolved preparations and culturally developed moralities.

NDT: possible sources of represented content and processing.

Measurement

What patterns appear in reported responses?

MFT instruments: dimensions and statistical associations.

A factor or network is not, by itself, a judgment-producing dependency.

Judgment formation

What helps form one particular judgment?

NDT: an exact represented relationship contributes to evaluation.

The comparison asks whether another account explains that distinction.

Different levels—not stages of one process or a claim of containment.

Figure 2. A learning architecture, a measurement structure, and a judgment-producing dependency answer different questions. The panels identify comparison levels, not a historical pipeline or a finding that one account contains another.

5. Dyadic morality: representing harm and explaining its moral force

5.1 The developed alternative

The Theory of Dyadic Morality asks how acts come to seem morally wrong and gives perceived harm a central role. The relevant harm is not simply objective injury. Schein and Gray's 2018 synthesis connects a graded intentional-agent/vulnerable-patient template with norms, negative affect, learning, and culturally informed understanding. The earlier agency-and-experience study is a precursor, not a substitute for that architecture. The synthesis describes harm as fundamental while qualifying its central prediction statistically: exceptions are expected to be rare, unstable, and effortfully maintained, not logically impossible. [Schein and Gray \(2018\)](#), pp. 32, 35, 38–43, 48–51; [Gray, Gray, and Wegner \(2007\)](#).

Two comparison jobs follow. One asks whether NDT can represent the content of a harmful dyad. Another asks whether perceived harm supplies the fundamental basis of moral judgment. Showing that NDT can use an agent–patient representation answers the first question; it does not settle the second.

Imagine an observer who learns that a competent adult deliberately served spoiled milk to a child despite an available alternative. The observer represents an actor, a vulnerable patient, and a causal connection to possible illness. A harmful-dyad description helps articulate what the observer sees. NDT asks additional questions: how does the observer represent the adult's relation to the child, which capacities are attributed to the adult, and how does that relationship help construct the requirement or appraisal? The observer's own concern for the child cannot simply substitute for the represented adult–child relationship when the observer judges what the adult owed the child.

The added NDT structure does not show that the dyadic theory is unaware of context or caring. It identifies the proposed organizational difference. Dyadic content can describe the subject matter of an appraisal; NDT distinguishes that content from the exact relationship's contribution to the appraisal's moral type. The issue is whether the source's account of moralization supplies the same discrimination, another explanation of the contrast, or a different target.

5.2 What would a counterexample establish?

Return to the original breakfast comparison. Both stipulated episodes can include the belief that serving spoiled milk could harm the child. In the rule-only episode, perceived harm can remain part of the represented situation without a qualifying moral relationship helping form the judgment. NDT consequently does not treat harm content as sufficient for moral type. Calling both episodes harm-related does not erase the conceptual contrast.

The converse NDT claim is that a moral judgment need not represent an intentional actor damaging a vulnerable patient. Consider a person deciding how to honor a valued practice or care for a valued object without representing injury to a patient as the basis of the requirement. A complete NDT application would still need the represented subject, object, value, positioning, agency, and contributing structure. The hypothetical does not earn moral type merely because the narrator calls the practice sacred. It identifies the kind of non-harm organization the theory permits once its conditions are satisfied.

A dyadic theorist can reasonably ask whether harm has really been excluded. Absence of physical injury is not absence of perceived harm. A person's statement that nobody is harmed is also not conclusive evidence about every relevant representation. Conversely, an analyst cannot add an unobserved harmed patient solely because a judgment has already been called moral. The comparison requires independently supported interpretations on both sides.

Modal strength matters equally. An intelligible non-harm case can challenge an exceptionless definition. It cannot by itself refute a claim that perceived harm is overwhelmingly characteristic of actual moralization. Assessing the statistical claim would require observations in a specified population and domain, with an adequate account of perceived harm and of the judgment being classified. NDT's internal examples establish what its architecture permits; they do not report the frequency of real non-harm judgments.

NDT rejects harm as a necessary basis of moral classification. An exceptionless harm-based boundary would therefore conflict with NDT's proposal. Schein and Gray's statistical qualification prevents treating their developed account as simply that exceptionless boundary. The substantive question is how the claim that harm is fundamental constrains moral classification once the qualification is retained. A conceptual distinction can clarify that question without supplying an empirical refutation.

5.3 Affect changes the explanation without replacing the comparison

Later affective work sharpens rather than introduces the role of feeling. Gray and colleagues' Affective Harm Account distinguishes visceral arousal from stimulus-directed appraisal; its authors also limit causal conclusions drawn from cross-sectional harm and appraisal reports. The development does not warrant describing the earlier dyadic account as emotion-free. [Gray et al. \(2022\), pp. 1204–1206, 1217.](#)

For NDT, a feeling's presence still leaves a functional question. Does the feeling alter what the person represents as valuable, focus attention on a frame, affect a consequence estimate, or help form the appraisal? Several contributions can coexist, but one observation does not establish every contribution. An affective rating can be informative without identifying the moral relationship or its exact role. The comparison accordingly concerns identifiable operations, not whether one theory contains emotion and another contains cognition.

Perceived harm: two separate comparisons

Dyadic-morality claim

Content of a harmful dyad

An intentional actor, perceived harm, and a vulnerable patient.

Question for NDT

Can the content be represented?

NDT can use those representations within its account of a judgment.

Harm as a fundamental basis

The developed theory qualifies its central claim statistically.

What makes the judgment moral?

NDT proposes a separate criterion: a contributing moral relationship.

Shared content does not settle the second question.

A conceptual counterexample is not an empirical test of frequency.

Figure 3. Reusable content does not settle a claim about the basis of moral judgment. The second comparison retains the dyadic theory's statistical qualification; no empirical refutation is depicted.

6. Relationships, social domains, and the moral boundary

Other programs approach the boundary through relationships, rules, or social function. Rai and Fiske's Relationship Regulation Theory explains moral judgment and motivation through unity, hierarchy,

equality, and proportionality, associated with communal sharing, authority ranking, equality matching, and market pricing. These are not merely labels for visible social arrangements. The account includes tacit motives, combinations of relational models, and obligations arising from relations among relationships. It distinguishes the motives actually underlying judgment from the rationales people state and predicts moral indifference when relationship motivation is absent. The paper develops a theoretical synthesis and predictions, with some questions about identifying the operative motives still open. [Rai and Fiske \(2011\), pp. 57–60, 64, 66–69.](#)

The comparison with NDT therefore concerns the proposed organization of a relationship and its contribution, not whether relationships or cognitive motives matter. A physician–patient case needs more than an authority label in either account: which relationship and motive does the physician actually take up? NDT further asks how the patient matters to the physician, how the patient is positioned relative to the physician, what response is represented as possible, and whether that organized relationship helps form the judgment. A relationship-regulating motive and NDT's complete moral-relationship condition are different proposed explanations of the case. Their adequacy depends on the distinctions they preserve, not on counting motives against builders.

Two questions remain separate: how relationship regulation classifies particular motives, including their combinations, and whether its four basic motives exhaust the basis of moral judgment. The universal-basis claim is the point of disagreement, not the existence of the four distinctions or the recognition of actual motivation. A claim that all moral motives fall within four types must not silently become the stronger claim that any occurrence of a corresponding social arrangement suffices for moral cognition. The stronger claim requires additional source support. NDT's permitted cases outside the taxonomy and its cases lacking contribution motivate the comparison; their empirical prevalence and classification are not established merely by displaying them in a formal model.

Social Domain Theory brings the problem into developmental research. Turiel's criterion-judgment tasks ask, among other things, whether an act would remain wrong without a rule or authority, and whether the judgment generalizes. Such response profiles help distinguish familiar moral and conventional cases. A pattern in responses, however, and a constitution of cognitive type have different explanatory roles. [Turiel \(1983\), chapter 4, pp. 61, 64.](#)

The parent's food-handling rule makes the distinction concrete. A person could apply a stored instruction to condemn an act whether or not an authority permits it. Authority-independence alone would not reveal whether a relationship helped construct the evaluation. Conversely, a relationship-grounded judgment need not reproduce every expected answer in a particular task. NDT therefore treats criterion profiles as possible evidence rather than sufficient classification by themselves. A constitutive reading of “definitional” conflicts with that proposal; a merely diagnostic reading does not establish the same conflict. The interpretation needs to remain visible rather than allowing the word to decide the matter silently.

Nichols offers a different explanation of the moral/conventional distinction, combining prohibited-action information with an affective mechanism. NDT can ask what prohibition and affect contribute to the breakfast judgment while still withholding an automatic moral classification from their conjunction. The source supplies an empirical rival organization of task performance, not a source-stated necessary-and-

sufficient classifier for NDT's judgment occurrences. A fair comparison keeps the measured task and the proposed cognitive boundary separate. [Nichols \(2002\)](#).

Morality-as-Cooperation introduces yet another level: biological and cultural solutions to recurrent cooperation problems. NDT permits cooperative normative judgments without moral provenance and moral judgments not organized as cooperative solutions. But an evolutionary or domain-level account must be compared with an appropriately aligned NDT account, rather than refuted by pointing out that it does not use individual judgment occurrences as its original unit. The disagreement concerns whether cooperation exhausts the moral domain under that alignment. It does not attribute an unstated token classifier to Curry and colleagues. [Curry, Mullins, and Whitehouse \(2019\)](#).

These encounters share a task without sharing one conclusion. Relational taxonomies organize motives; developmental profiles distinguish responses; affective accounts propose mechanisms; cooperative accounts organize functions. NDT asks how each bears on the moral/nonmoral boundary in judgment formation. The answer must preserve the external account's own target before identifying what remains unexplained.

7. Changing what follows, and hearing more than one voice

The breakfast example has so far held the represented danger fixed. Suppose the parent instead receives a credible medical explanation that a particular product is fermented rather than spoiled. The changed information can alter the predicted consequence of serving it. Suppose an emergency instead makes a normally prohibited action the only available way to avert a greater danger. The requirement may remain applicable while another consideration supplies a reason to override it. Those changes do different work.

NDT uses *bender* for a represented influence that changes what follows from an underlying relationship while preserving the exact relationship core and its continuing contribution across the specified comparison. Changing the child's represented value, the parent–child connection, or the relevant agency would instead change the core on that operation. One representation can perform both kinds of work on separately supported connections; a topic such as vulnerability or ownership does not carry an automatic classification.

Familiar philosophical approaches illuminate different forms of practical organization. Bentham's utility principle supplies an outcome-maximizing family. A represented Kantian humanity requirement can exclude an outcome-better act. Ross's prima-facie duties make the persistence of an overridden requirement intelligible. Held's care framework and Scanlon's reasonable rejectability organize still other guidance. NDT represents those restricted families without deciding which first-order standard is correct or declaring that every use of any family is moral in cognitive type.

The distinctions prevent an artificial opposition between calculation and relationship. A relationship can help construct an outcome ordering, a prohibition, or a complaint-sensitive standard. Equally, a person can apply a sophisticated rule without a moral relationship contributing on that occasion. The question is not whether the person sounds Kantian or consequentialist but what organizes the particular evaluation.

Conflict can also occur within one decision problem or between different interpretations. If every feasible act violates a requirement, a person may still judge one act the least unacceptable while retaining regret or residue. NDT represents a violation-sensitive fallback only when that fallback is itself represented; it does

not automatically drop the requirements and maximize outcomes. Such tragedy differs from two interpretations supporting incompatible judgments.

The term *polyphony* comes from Bakhtin's account of independent, unmerged voices in Dostoevsky's novels. Beal's use moves from literary organization to the possibility of several coherent normative interpretations within a person. NDT adds conditions not attributed to Bakhtin: distinct frame activations, a common case, appropriate comparison bridges, and conflicting outputs. The literary source supplies a lineage, not an experimental validation or a formal precursor for every added condition. [Bakhtin \(1984\), chapter 1, p. 6.](#)

Consider a person who has promised confidentiality and then learns that disclosure could protect someone. Under one interpretation, disclosure is betrayal of the confidant; under another, disclosure is protection of a threatened person. The example needs more than two emotionally opposed words. Are betrayal and protection descriptions of the same possible disclosure? Do the judgments answer the same practical question? If one appraisal concerns loyalty and another concerns the act's effectiveness, opposite evaluations need not conflict at a common dimension. NDT requires the represented bridge appropriate to the comparison before calling the pair polyphonic.

A further interpretation may arbitrate between the two judgments, but the theory does not require a final voice that silences every other voice. Nor does several-frame organization make each judgment justified. Polyphony concerns the structure of the person's represented conflict. Truth, defensibility, and eventual action remain additional questions.

Source notes: [Bentham, chapter I, §II](#); Kant, *Groundwork*, Formula of Humanity, Academy 4:429; [Ross, chapter 2, p. 28](#); [Held \(1995\)](#); [Scanlon \(1998\), chapter 5, p. 189](#). NDT—Formal Specification §§3.2, 5–6. The confidentiality case is an illustration of comparison requirements, not a new empirical finding.

8. Producing support, reported reasons, and available explanations

The review can now return to the observer at breakfast. The observer sees the parent discard the milk and asks why. The parent says, “I followed the rule,” or “I wanted to protect my child.” Either answer supplies evidence, but neither automatically reconstructs the thought that helped form the judgment. The first answer might omit a contributing relationship; the second might offer a fitting explanation after the decision.

Haidt's Social Intuitionist Model gives rapid intuition a central producing role and treats much later reasoning as socially situated justification, while retaining other routes. Its relevance is not that NDT must be intuition-first. It makes the difference between producing support and subsequent explanation impossible to ignore. NDT permits intuition-first, reasoning-first, and mixed processes; timing or reported sequence alone does not confer moral type. [Haidt \(2001\), p. 817.](#)

Before asking what made the parent's practical judgment moral, consider the simpler belief that the milk has spoiled. Suppose the parent remembers that the carton was left unrefrigerated overnight but forms the spoilage belief from its sour smell alone. The refrigeration information is a reason available to the parent, yet it is not the reason on which the present belief rests. In a further imagined variation, remembering the overnight lapse prompts the parent to inspect the carton, while the smell supplies the basis of the resulting belief. The memory helps explain why the parent acquired the belief without necessarily being its

evidential basis. Possession, causal influence, and basing can therefore mark different relations even when the belief's content stays fixed.

Work on *epistemic basing* examines such relations for belief. Neta distinguishes reasons someone has, reasons why an attitude occurs, and reasons for which it is held. McCain asks how an interventionist account of causation can distinguish basing from merely being causally upstream. The NDT question changes the target: which represented relationship helps form the evaluation that makes the judgment moral? The epistemological distinction prepares that question without answering it.

The change of target becomes particularly important when justification enters the comparison. A person might possess a good reason while believing for an unrelated bad reason. Does putting those two facts together establish belief for a good reason? Sylvan and Lord reject a merely composite answer: the appropriate connection must itself obtain. Silva, addressing doxastic justification, challenges a general demand for the proposed kind of basing. Those positions ask what makes a belief properly based or justified. NDT's requirement concerns descriptive moral type, so neither conclusion transfers merely because the word "basis" appears in both discussions.

The boundary is productive. The parent can possess a reason without using it, and a reason can be causally upstream without doing the particular evaluative work claimed. NDT also refuses a split explanation in which one inactive consideration supplies moral-looking content while another alone produces the judgment. But a shared concern about basing is not evidence that the epistemological accounts are components of NDT or that a conclusion about justification transfers to cognitive classification.

Redundant support creates another difficulty. Suppose, in an explicit hypothetical, the parent's concern for the child and the remembered rule both help produce the discard judgment. Either would have been enough to preserve the spoken conclusion. Removing the rule and hearing the same answer would not establish that the rule had done nothing. Compare a different hypothetical in which the rule remains stored but unused: the rule would have supplied an answer if needed, yet contributes nothing to the actual judgment. An available backup and two operating sources are different structures. Looking only for a changed answer after removal can miss that difference.

Formal causal tools help state the narrower questions. A sufficient set contains enough operating support for the specified result; a member can be necessary within that set even when another sufficient set also operates. Wright's NESS condition captures that form of credit. NDT uses a bounded profile-level counterpart within its broader contribution architecture. The whole claim is not that a relationship occurs somewhere in a sufficient explanation: the exact profile must do the specified work, and the remaining contribution and binding conditions must still hold.

Halpern and Pearl's structural approach offers another resource, examining dependence under selected contingencies. For the breakfast hypothetical, a contingency asks what would happen under a specified change to support while other relevant conditions are fixed. Choosing those conditions is part of the substantive account, not an automatic consequence of drawing arrows. Neither a structural-model verdict nor the NESS operation alone establishes NDT's full contribution conditions or moral type.

Challenges introduce another distinction. If the parent learns that a warning was fabricated, the information challenges the credibility of a premise. If the warning is genuine but irrelevant to the

particular product, the challenge concerns a connection. By contrast, a reason to accept a risk changes the practical evaluation or priority.

Now imagine the parent receives two reports: one disputes the warning's credibility, while the other disputes the reliability of the first report. Listing both challenges does not settle which support survives. A resolution rule is needed, especially if two challenges undermine each other and neither has independent support that decides the dispute. NDT's cautious reference rule does not simply choose whichever conclusion restores the desired judgment. It begins from the admitted challenge structure and settles only what that structure warrants under the selected rule; unresolved challenges are not counted as defeated.

Dung's grounded semantics supplies the bounded least-fixed-point calculation for that reference rule. The calculation can be specified precisely once the admitted attack graph is given. Deciding which challenges the person represents, what they target, and whether they contribute to the episode remains additional work. Other challenge resolvers remain possible implementations; the exact computational match does not identify a universal psychology of doubt.

Empirical research supplies further resources without closing the explanatory gap. Van Bavel and colleagues' changing evaluative instructions provide a signature of flexible construal. Cushman and Young show the relevance of domain-general causal and intentional representations. Kleiman-Weiner and colleagues provide hierarchical Bayesian machinery for learning others' value weights. A global-workspace account offers one possible realization of functional availability. None of those findings or mechanisms by itself identifies the complete contributing organization of the parent's judgment.

That limitation should not be turned into a promise that sufficiently many behavioral probes will eventually reveal an exact mental route. The conceptual distinctions do not depend on a completed measurement procedure. Empirical applications do depend on evidence adequate to their claims, and several rival functional explanations may remain unresolved. Latency, self-report, transfer, intervention, and other observations can constrain interpretations; their combination is not a guaranteed window into every constitutive relation.

Source notes: [Neta \(2019\)](#); [McCain \(2012\)](#); [Lord and Sylvan \(2020\)](#), §2.2, indexed as Sylvan and Lord (2019) in the passage record; [Silva \(2015\)](#); [Wright \(1985\)](#), p. 1774; [Halpern and Pearl \(2005\)](#); [Dung \(1995\)](#), definition 20, p. 329. Flexible construal, domain-general representation, Bayesian learning, and workspace broadcasting are separately identified in Appendix A.

9. What the comparisons warrant

NDT's proposed contribution is an integrated distinction among normative formation, moral relationship, contributing provenance, and comparison-ready conflict. The literature provides useful operations, empirical signatures, philosophical contrasts, rival partitions, and substantive disagreements. Those relations are not interchangeable versions of support.

Kant's shopkeeper, for example, makes outward conformity insufficient for moral worth. NDT likewise refuses to classify solely by surface conformity, but it classifies a judgment's cognitive type rather than the worth of the action or the purity of a motive. A person-centered account of character inference can explain what an act suggests about its agent without itself supplying a moral/nonmoral token boundary. A

word-association graph can predict moral relevance without reconstructing a decision's provenance. Each achievement deserves its own description.

The possibility of a unified moral kind also remains a substantive question. Sinnott-Armstrong and Wheatley challenge familiar candidates for unity and recommend finer-grained inquiry; Kumar defends homeostatic unity. NDT proposes relationship provenance as a different unifier. Its formal definition does not establish that provenance-based groups are empirically superior, any more than ordinary agreement about the word “moral” establishes a single psychological mechanism. The relative explanatory value of the proposed groups must be assessed without defining success as agreement with NDT.

Across the 48 governed records, no exact same-level match has been established for NDT's complete architecture. The result is bounded by the specified constructs and sources. More importantly, the review identifies where a judgment-content taxonomy, a learned mechanism, a philosophical standard, or an observable response pattern leaves a distinct organizational question. Partial overlap does not answer that question by itself; a different vocabulary does not prove that no other account can answer it.

The breakfast decision provides the final test of the review's purpose. Readers can agree that the milk should be discarded while disagreeing about what explains the judgment and what makes it moral. NDT asks for those questions to be separated and then reconnected through an explicit account of contributing thought. The literature supplies both resources and challenges for that task. Neither an internal model demonstration nor a familiar label completes it.

Source notes: Kant, *Groundwork*, Academy 4:397; [Uhlmann, Pizarro, and Diermeier \(2015\)](#); [Ramezani and Xu \(2025\)](#); [Sinnott-Armstrong and Wheatley \(2014\)](#); [Kumar \(2016\)](#). The 48-record reference follows.

Appendix A. Construct-level comparison reference

The 48 entries retain Dataset 1.20's record identities, comparison objects, and exact relation assignments. Forty-six entries have the record class “comparison”; one is a scope-boundary record and one is a lineage record. The latter two have no assigned relation field, as indicated below. Names identify particular claims or operations, not whole-theory scores.

The prose identifies the claim being compared and distinguishes a retained classification from its source-based limits. In particular, the MFT inventory is open, the dyadic basis claim carries a statistical qualification, and the constitutive readings of social-domain and relational-motive claims require the qualifications printed with their entries. Retaining an identifier and relation does not convert those qualifications into a resolved theoretical adjudication.

An “exact submodule” is a match to a bounded formal operation. An “implementation family” supplies possible content or machinery for an NDT operation. A “restricted decision-model family” occupies a limited region of NDT's permitted practical structures. An “empirical signature” supplies observations that may inform an application without constituting the target. An “empirical rival partition” proposes a competing grouping for aligned observations. “Tokenwise cross-classification” preserves different criteria applied to individual occurrences. An “output-only approximation” can match a result without recovering its derivation. “Contradicts NDT” identifies a source-bounded conflicting claim, subject to the qualifications stated. “Scope boundary,” “conceptual context,” and “lineage” do not establish computational equivalence. “Source claim underdetermined” withholds a stable same-level classification.

CTX001. Goal-and-means reasoning

Retained exact relation: **restricted decision-model family**.

Objects compared: Explicit-goal detachment / Practical-inference guide structure.

NDT: When a person represents an end and a feasible means, NDT can use that structure to form prospective guidance. NDT studies the judgment produced and makes no claim that the judgment is itself an action or intention.

External claim: Practical inference organizes ends or goods and available means into guidance about what to do; its terminus varies across Aristotle, Anscombe, and Davidson.

Sources: [Aristotle, n.d.-a \(S01\)](#); [Anscombe, 1957 \(S46\)](#); [Davidson, 1978 \(S47\)](#).

CTX002. What supplies motivation

Retained exact relation: **source claim underdetermined**.

Objects compared: Explicit conative evaluation / Passional source requirement for motivation.

NDT: NDT allows motivation through a separately represented desire and through evaluation built into control. It does not yet settle whether the second kind should count as a Humean calm passion.

External claim: Hume holds that reason alone cannot motivate: an impulse must come from passion, including calm passions known mainly through their effects.

Sources: [Hume, n.d.-a \(S02\)](#); [Hume, n.d.-c \(S49\)](#).

CTX003. Rule-based guidance

Retained exact relation: **implementation family**.

Objects compared: Rule-conditioned detachment / Input/output logic.

NDT: NDT can form a possible normative conclusion from a requirement represented as available for appraisal or guidance in the decision's perspective. The perspective, the requirement's source, and its addressee remain separate, and the requirement need not be treated as an ordinary if-then proposition.

External claim: Input/output logic maps stated conditions to normative conclusions without treating the governing rule as an ordinary if-then proposition.

Sources: [Makinson & van der Torre, 2007 \(S03\)](#).

CTX004. Institution-based requirements

Retained exact relation: **restricted decision-model family**.

Objects compared: Institutional norm formation / Institutional promise obligation.

NDT: When an institution, its rule, and the conditions for that rule are represented as active, the institution can supply or shape a requirement.

External claim: A promise made within an active institution can create an obligation because of the institution's rules.

Sources: [Searle, 1964 \(S04\)](#).

CTX005. Cautious challenge resolution

Retained exact relation: **exact submodule**.

Objects compared: Grounded challenge semantics / Grounded argumentation semantics.

NDT: NDT allows an admitted challenge to block a reason, a step in the reasoning, or the proposed conclusion. Its cautious default accepts only what survives the resulting attacks.

External claim: Grounded semantics resolves a network of attacks into one cautious set of accepted claims.

Sources: [Dung, 1995-a \(S05\)](#).

CTX006. Outcome-maximizing guidance

Retained exact relation: **restricted decision-model family**.

Objects compared: Outcome-maximizing guidance / Principle of utility.

NDT: Some NDT models guide action by choosing the best represented outcome. This is neither the only kind of guidance NDT permits nor enough to make a judgment moral.

External claim: The principle of utility judges actions by whether they increase or reduce the happiness of those affected.

Sources: [Bentham, n.d. \(S06\)](#).

CTX007. Requirements that outcomes cannot outweigh

Retained exact relation: **restricted decision-model family**.

Objects compared: Hard requirements / Formula of Humanity.

NDT: An applicable requirement that has not been overridden can rule out an action even when that action would produce a better outcome.

External claim: The Formula of Humanity requires treating rational nature as an end in itself.

Sources: [Kant, n.d. \(S07\)](#).

CTX008. Overridden duties that still matter

Retained exact relation: **restricted decision-model family**.

Objects compared: Override and residue / Prima-facie duty and moral residue.

NDT: A higher-priority requirement can license choosing an action that violates a lower-priority requirement without erasing the lower requirement or its residue. This license is not a primitive permission output.

External claim: A prima-facie duty can be overridden while the unfulfilled duty still matters morally.

Sources: [Ross, 1930 \(S08\)](#).

CTX009. Relationship-grounded guidance

Retained exact relation: **restricted decision-model family**.

Objects compared: Relationship-grounded guidance / Care as a moral framework.

NDT: A represented relationship can create the requirement that guides action, rather than merely motivating someone to follow a rule fixed independently.

External claim: Care can be treated as the wider moral framework within which justice should operate.

Sources: [Held, 1995 \(S10\)](#).

CTX010. Complaint-sensitive guidance

Retained exact relation: **restricted decision-model family**.

Objects compared: Complaint-sensitive guidance / Reasonable rejectability.

NDT: NDT can represent cases in which complaints or reasonable rejection guide a decision, but those cases do not define all normative thinking.

External claim: Reasonable rejectability asks whether conduct can be justified to others on terms they could not reasonably reject.

Sources: [Scanlon, 1998 \(S11\)](#).

CTX011. Content does not settle moral type

Retained exact relation: **implementation family**.

Objects compared: Content–provenance separation / Prepared-module moral-intuition architecture.

NDT: NDT denies that activation of a prepared moral-content module is enough to determine whether a particular judgment is moral; qualifying provenance remains necessary.

External claim: Haidt and Joseph present prepared modules as one possible source of moral intuitions; Graham and colleagues describe related modules working together to guide and constrain responses.

Qualification: The implementation comparison does not classify the whole theory. The 2013 synthesis's learning and nativist commitments remain relevant; the foundation-domain partition is a separate comparison.

Sources: [Haidt & Joseph, 2004 \(S50\)](#); [Graham et al., 2013 \(S51\)](#).

CTX012. Representing a harmful dyad

Retained exact relation: **implementation family**.

Objects compared: Dyadic harm appraisal / Dyadic harm template.

NDT: NDT can represent an observer's judgment that an intentional actor harmed a vulnerable person. That pattern neither determines whether the judgment is moral nor covers every moral structure NDT permits.

External claim: The dyadic harm template represents harm as an intentional actor damaging a vulnerable patient.

Sources: [Schein & Gray, 2018-a \(S17\)](#).

CTX013. Response patterns do not define a domain

Retained exact relation: **contradicts NDT**.

Objects compared: Domain-type provenance / Criterion-judgment domain identification.

NDT: NDT allows judgments classified as moral, conventional, or personal to show the same rule-independence, generalizability, and related response patterns. These patterns may provide evidence about a judgment's domain, but they do not by themselves determine its cognitive type; judgments with the same pattern can arise from different evaluative organization and provenance.

External claim: The registered comparison reads Turiel's criterion judgments as definitional features distinguishing domains. Whether the relevant wording constitutes cognitive type or identifies it diagnostically remains an interpretive qualification.

Qualification: The retained contradiction is conditional on the constitutive reading. Criterion responses as fallible indicators are treated separately in the empirical-signature record; a diagnostic interpretation alone does not establish contradiction.

Sources: [Turiel, 1983-b \(S53\)](#).

CTX014. Cooperation does not constitute moral type

Retained exact relation: **contradicts NDT**.

Objects compared: Cooperation–moral-type independence / Cooperation constitution of morality.

NDT: Cooperation can shape judgments, but NDT permits cooperative judgments that are nonmoral and moral judgments that are not cooperative solutions.

External claim: Morality-as-Cooperation says morality consists of biological and cultural solutions to recurrent cooperation problems.

Qualification: The comparison requires declared domain-level alignment. The source is not credited with an unstated necessary-and-sufficient classifier of individual judgment occurrences.

Sources: [Curry, Mullins & Whitehouse, 2019 \(S19\)](#).

CTX015. What makes a represented relationship morally qualifying

Retained exact relation: **tokenwise cross-classification**.

Objects compared: Gate 2 represented moral relationship qualification / Relational-model individuation of moral motives.

NDT: NDT qualifies a represented relationship through three independently variable roles: a nonempty inherent-valence profile with separate positive and negative components, Relational Positioning of the exact object relative to the exact subject, and perceived agency. The subject gate, functional access, and same-chain profile and object contribution remain separate. Positioning can be identity-mediated, intimate, or encounter-based; identity-mediated forms can be shared, contrastive, or complementary-role. No relational-model label, positioning route, identity form, or strict bender decides qualification by itself.

External claim: This construct distinguishes moral motives according to which of four basic kinds of social relationship they regulate.

Qualification: Relationship qualification, subjecthood, access, and contribution are separate NDT conditions. A relationship's social-model label alone does not establish their satisfaction. The two-way comparison is a formal proposal, not an observed distribution of cases.

Sources: [Rai & Fiske, 2011-a \(S20\)](#).

CTX016. Learning other people's moral weights

Retained exact relation: **implementation family**.

Objects compared: Learned social-value weighting / Hierarchical Bayesian moral-weight inference.

NDT: Learned estimates of how much people value abstract moral principles can shape NDT's comparison of outcomes, but they do not determine whether a particular judgment is moral.

External claim: A hierarchical Bayesian model uses observed behavior to estimate how strongly other people weight abstract moral principles.

Sources: [Kleiman-Weiner, Saxe & Tenenbaum, 2017 \(S21\)](#).

CTX020. NESS credit is one part of the contribution test

Retained exact relation: **exact submodule**.

Objects compared: Profile-level NESS credit within actual contribution / NESS minimal-sufficient-set membership.

NDT: Within NDT's broader actual-contribution architecture, a relationship profile can receive NESS credit when it is necessary within an actually operating support profile that is sufficient for the judgment.

External claim: The NESS condition identifies a cause as a necessary member of an actually instantiated set sufficient for the outcome.

Sources: [Wright, 1985 \(S38\)](#).

CTX022. Bargaining-based decision models

Retained exact relation: **restricted decision-model family**.

Objects compared: Bargaining-derived decision models / Resource-rational agreement approximation.

NDT: NDT can represent bargaining-based standards and decision routes, but using them does not determine whether a particular judgment is moral.

External claim: Resource-rational contractualism treats many action- and welfare-based moral judgments as efficient approximations of agreement-based judgment.

Sources: [Levine et al., 2026 \(S40\)](#).

CTX023. Association-based moral prediction

Retained exact relation: **output-only approximation**.

Objects compared: Associative moralization / Association-based moral-relevance prediction.

NDT: Association networks can approximate moral relevance or labels, but they do not reconstruct the decision process or the provenance of a particular judgment.

External claim: The Moral Association Graph uses word associations to predict human moral relevance and moralization, with tested performance against baseline models.

Sources: [Ramezani & Xu, 2025-b \(S55\)](#).

CTX024. Case-by-case guidance

Retained exact relation: **restricted decision-model family**.

Objects compared: Tokenwise contextual variation / Moral particularism.

NDT: NDT allows guidance to depend on the particulars of a case while also allowing models that use general principles.

External claim: Moral particularism holds that moral thought does not depend on exceptionless general principles.

Qualification: The retained restricted-family classification concerns guidance without dependence on exceptionless principles. It does not classify Dancy's wider holism, realism, or epistemology, or resolve every technical characterization of particularism.

Sources: [Dancy, 2004 \(S15\)](#).

CTX025. Generation–status separation

Record class: **scope boundary**; exact relation field: **null**.

Objects compared: Generation–status separation / Normative and metaethical status.

NDT: NDT's account of how a judgment is produced does not by itself decide whether that judgment is true, justified, correct, or objectively valid.

External claim: Normative and metaethical status concerns whether moral judgments are true, justified, or objectively valid; Mackie's objective-value denial is one instance.

Sources: [Mackie, 1977 \(S45\)](#).

CTX026. Many-voiced normative cognition

Record class: **lineage**; exact relation field: **null**.

Objects compared: Normative polyphony / Bakhtinian polyphony.

NDT: NDT represents multiple coherent normative frames without forcing them into one voice and compares them only when their targets can be aligned.

External claim: Bakhtinian polyphony describes independent voices and consciousnesses that are not merged into a single viewpoint.

Sources: [Bakhtin, 1984 \(S44\)](#).

CTX030. No ought from evaluatively empty inputs

Retained exact relation: **scope boundary**.

Objects compared: Evaluatively empty input barrier / Evaluatively empty is–ought barrier.

NDT: When a description contains no evaluative ordering, standard, or requirement, NDT derives no prescriptive conclusion from that description alone.

External claim: Hume's is–ought barrier requires explaining any move from ordinary descriptive propositions to a new ought-relation.

Sources: [Hume, n.d.-b \(S48\)](#).

CTX031. Harm templates do not constitute morality

Retained exact relation: **contradicts NDT**.

Objects compared: Dyadic-harm–moral-type independence / Dyadic harm as moral basis.

NDT: NDT permits moral judgments without a dyadic harm template and dyadic harm appraisals that lack qualifying moral provenance.

External claim: Schein and Gray describe perceived harm as fundamental to moral judgment, while qualifying the central prediction statistically.

Qualification: Read with §5's distinction between a constitutive boundary and a statistical generalization. No empirical refutation is asserted.

Sources: [Schein & Gray, 2018-b \(S52\)](#).

CTX032. Four relational models do not exhaust moral qualification

Retained exact relation: **contradicts** NDT.

Objects compared: Gate 2 moral relationship beyond four relational models / Four relational models as universal moral basis.

NDT: NDT permits a represented relationship to qualify outside communal sharing, authority ranking, equality matching, and market pricing. Conversely, fitting one of those models is insufficient when the independent inherent-valence, Relational Positioning, perceived-agency, subject, access, or actual-contribution conditions fail.

External claim: Rai and Fiske's universal-basis formulation places moral judgment and behavior across cultures within four relational motives.

Qualification: The retained contradiction targets the universal-basis formulation. A universal classification of moral motives does not alone establish that every occurrence of a relational-model type is sufficient for moral cognition. That stronger inference remains unestablished.

Sources: [Rai & Fiske, 2011-b \(S54\)](#).

CTX033. Structural models of actual causation

Retained exact relation: **implementation family**.

Objects compared: Actual-contribution implementation / Contingent-dependence structural actual causation.

NDT: NDT needs a way to test actual contribution under redundant support, but it does not require one unique structural-equation definition of causation.

External claim: Halpern–Pearl actual causation tests dependence under selected structural-model contingencies, making model choice part of the analysis.

Sources: [Halpern & Pearl, 2005 \(S39\)](#).

CTX034. Evaluation can be built into a representation's control role

Retained exact relation: **source claim underdetermined**.

Objects compared: Encoded Evaluative Realization / Strict representation–passion separability.

NDT: NDT permits one CET route in which evaluative structure is built into a representation's control-sensitive role rather than supplied only by a separately represented desire, policy, or verdict. This is a possible realization of CET, not a condition on every normative detachment.

External claim: On a strict separability reading of Hume, a representation can direct action but cannot itself supply the evaluative or motivating impulse; a distinct passion must do so.

Sources: [Hume, n.d.-a \(S02\)](#).

CTX035. Source-sensitive judgments and source-sensitive moral credit

Retained exact relation: **conceptual context**.

Objects compared: Judgment provenance and moral worth / Kant's source-sensitive moral-worth contrast.

NDT: NDT classifies the cognitive type of a judgment by what actually helped form it; it does not assess the moral credit of an action or the purity of an agent's motive.

External claim: Kant's shopkeeper illustrates that outwardly honest conduct can differ in moral worth according to why it is performed.

Sources: [Kant, 1785 \(S57\)](#).

CTX036. Feeling can help build a moral judgment without defining morality

Retained exact relation: **empirical rival partition**.

Objects compared: Affect-linked norms as an input, not a type criterion / Sentimental rules account of the moral/conventional distinction.

NDT: Affect-linked prohibitions may help form a moral judgment, but affect, prohibition, and their conjunction do not determine moral type without relationship provenance.

External claim: Nichols explains the moral/conventional distinction through prohibited-action information together with an affective mechanism.

Qualification: The retained rivalry concerns task-level organization and a proposed comparison with provenance; it is not a source-attributed necessary-and-sufficient moral-token definition.

Sources: [Nichols, 2002 \(S58\)](#).

CTX037. A precise model need not preserve every use of the word moral

Retained exact relation: **conceptual context**.

Objects compared: Moral cognition as an explanatory construct / Family-resemblance structure in concept use.

NDT: NDT offers a precise explanatory construct for moral cognition rather than a dictionary rule that must reproduce every ordinary, disciplinary, or philosophical use of the word moral.

External claim: Wittgenstein's family-resemblance model explains how uses can overlap without sharing one feature common to every case.

Sources: [Wittgenstein, 1953 \(S59\)](#).

CTX038. Response patterns can help identify moral and nonmoral normativity

Retained exact relation: **empirical signature**.

Objects compared: Empirical identification of moral versus nonmoral normativity / Social Domain moral/conventional criterion-judgment paradigm.

NDT: NDT distinguishes moral from nonmoral normative judgments by their represented organization and actual contributing provenance. Authority-independence, generalizability, seriousness, and related response patterns can be used as fallible evidence about that distinction when they are part of an independently specified identification model; no response pattern settles a token's type by itself.

External claim: Social Domain studies distinguish moral rules from conventions through patterns such as authority-independence, generalizability, and judged wrongness or seriousness across cases.

Sources: [Turiel, 1983-a \(S18\)](#); [Turiel, 1983-b \(S53\)](#).

CTX044. A relationship route must be available to the operative process in order to contribute, but access does not confer moral type

Retained exact relation: **implementation family**.

Objects compared: Functional access separated from relationship qualification / Global-workspace broadcast availability.

NDT: A valence, positioning, agency, or bender representation must be available to the process that uses its edge, but functional access alone supplies none of those roles and does not confer moral type.

External claim: The global-workspace framework describes selected and amplified contents becoming available to multiple receiving systems.

Sources: [Dehaene & Naccache, 2001 \(S77\)](#).

CTX045. Shared identity is one identity-mediated way to position an object relative to a subject

Retained exact relation: **implementation family**.

Objects compared: Shared identity form of identity-mediated Relational Positioning / Common-ingroup representation.

NDT: Shared identity is one form of identity-mediated Relational Positioning through an operative shared 'we'. Mere category co-membership or role possession is insufficient, and intimacy or direct encounter can position the object without shared identity.

External claim: Gaertner and colleagues compared one-group, two-group, and separate-individual representations and found different patterns of bias reduction.

Sources: [Gaertner et al., 1989 \(S78\)](#).

CTX046. Intimacy positions an object through a particular relationship

Retained exact relation: **implementation family**.

Objects compared: Intimacy Relational Positioning route / Interpersonal-process intimacy.

NDT: Intimacy is one Relational Positioning route involving relationship-specific history, disclosure, responsiveness, trust, dependence, mutual recognition, or knowledge. It is neither necessary nor sufficient for moral type.

External claim: Laurenceau and colleagues tested interaction-level relations among disclosure, perceived partner responsiveness, and experienced intimacy.

Sources: [Laurenceau et al., 1998 \(S79\)](#).

CTX047. Direct encounter positions the particular other through presence, expression, testimony, address, or condition

Retained exact relation: **empirical signature**.

Objects compared: Direct-encounter Relational Positioning route / Empathy-induced particularized altruism.

NDT: Direct encounter is one Relational Positioning route. Empathic concern is a possible process or signature, not the definition of the route and not a direct measure of moral type or actual contribution.

External claim: Batson and colleagues induced empathy toward a particular person and found target-favoring allocations that could conflict with an independently recognized justice principle.

Sources: [Batson et al., 1995 \(S80\)](#).

CTX048. Can producing history unify moral judgment?

Retained exact relation: **empirical rival partition**.

Objects compared: Provenance-kind unity as a taxonomic hypothesis / Fine-grained moral-judgment taxonomy.

NDT: NDT accepts that judgments called moral are heterogeneous in content, feeling, function, and implementation, but proposes a narrower class unified by a qualifying represented relationship and the support that actually formed the judgment.

External claim: Sinnott-Armstrong and Wheatley argue that moral judgments are not unified by familiar candidates such as content, phenomenology, force, form, function, or a shared neural mechanism, and recommend studying more carefully defined groups.

Sources: [Sinnott-Armstrong & Wheatley, 2014 \(S63\)](#).

CTX049. What kind of unity does moral judgment have?

Retained exact relation: **empirical rival partition**.

Objects compared: Provenance-kind unity against homeostatic unity / Moral judgment as a natural kind.

NDT: NDT locates unity in represented relationship organization and actual contributing support while allowing different contents, emotions, and implementations.

External claim: Kumar argues that moral judgment is a natural kind whose components are held together by homeostatic mechanisms and develops a hybrid account involving belief and emotion.

Sources: [Kumar, 2016 \(S64\)](#).

CTX050. The same action can be understood in different ways

Retained exact relation: **empirical signature**.

Objects compared: Frame-relative token construal / Flexible moral construal.

NDT: NDT classifies a particular judgment through the interpretation active on that occasion, not through a stimulus description alone. An interpretation matters only through the support that actually helps form the judgment.

External claim: Van Bavel and colleagues found that people could shift between moral and nonmoral evaluation of the same actions when their evaluative goal changed from trial to trial.

Sources: [Van Bavel et al., 2012 \(S62\)](#).

CTX051. Nonmoral representations can help build a moral judgment

Retained exact relation: **empirical signature**.

Objects compared: Domain-general support within moral judgment / Nonmoral representations that shape moral judgment.

NDT: Causal, intentional, and other domain-general representations may help build or bend a moral judgment without becoming what makes the judgment moral.

External claim: Cushman and Young explain familiar patterns of moral judgment through nonmoral causal and intentional representations, including action-omission and means-side-effect contrasts.

Sources: [Cushman & Young, 2011 \(S61\)](#).

CTX052. Being available is not the same as helping form the judgment

Retained exact relation: **scope boundary**.

Objects compared: Availability–contribution separation / Reasons one has, reasons why, and reasons for which.

NDT: NDT separates considerations that were available or causally upstream from representations that actually helped form this particular judgment.

External claim: Neta distinguishes a reason one has, a reason why one believes, and a reason for which one believes; a consideration can be available and causally explanatory without being the basis on which the belief is held.

Sources: [Neta, 2015 \(S65\)](#); [Neta, 2019 \(S66\)](#).

CTX053. A local difference-maker test for actual contribution

Retained exact relation: **implementation family**.

Objects compared: Interventionist model of local contribution / Interventionist epistemic basing.

NDT: NDT requires more than a path leading toward the judgment: the relevant representation must belong to an actual sufficient route and make a local difference to the evaluation-relevant organization.

External claim: McCain uses interventionist causation to distinguish a reason's nondeviant basing of a belief from mere causal ancestry, including under overdetermination.

Sources: [McCain, 2012 \(S67\)](#).

CTX054. One contributing route must do both jobs

Retained exact relation: **scope boundary**.

Objects compared: Same-chain moral provenance / Prime View of believing for a normative reason.

NDT: NDT withholds moral classification when one route merely contains a qualifying relationship while a different route alone forms the evaluation. The relationship must contribute on the path that forms this particular judgment.

External claim: Sylvan and Lord reject treating believing for a reason and that reason's normative standing as two independently sufficient facts; their Prime View requires the first to obtain because the second obtains.

Sources: [Sylvan & Lord, 2019 \(S68\)](#).

CTX055. Does moral classification really require actual support?

Retained exact relation: **scope boundary**.

Objects compared: Pressure test for actual support / No general basing requirement for doxastic justification.

NDT: NDT says that representing an appropriate relationship is not enough: that relationship must actually help form the judgment before it can make the judgment moral.

External claim: Silva challenges the claim that a belief can be doxastically justified only if it is based on an epistemically appropriate reason, while limiting the argument to doxastic justification.

Sources: [Silva, 2015 \(S69\)](#).

CTX056. Direct encounter without prior identity or intimacy

Retained exact relation: **conceptual context**.

Objects compared: Direct encounter without prior identity or intimacy / Levinasian face-to-face ethical address.

NDT: NDT allows a particular other's presence, expression, testimony, address, or condition to position that other relative to the subject without prior identity-mediated positioning or intimacy. Actual contribution remains separate.

External claim: Levinas describes the face-to-face relation as an ethical address that resists being reduced to the subject's power or prior assimilation of the other.

Sources: [Levinas, 1961/1969 \(S81\)](#).

CTX057. Character inferences supply a signature; the rival partition is NDT's test

Retained exact relation: **empirical signature**.

Objects compared: Person-centered character-inference signature and NDT-constructed H18 rival / Person-centered moral character inference.

NDT: NDT treats character inference as a possible influence and observable signature, while constructing a separate character-based rival partition for a fair test against provenance. The source does not itself define that partition or make character inference constitutive of moral type.

External claim: Uhlmann, Pizarro, and Diermeier argue that moral judgment often centers on what an act reveals about an agent's moral character, not only on the act's isolated permissibility or outcome.

Qualification: A character-based rival partition is an NDT-constructed comparison proposal. Uhlmann and colleagues' character-inference account is not itself classified as supplying that token partition.

Sources: [Uhlmann et al., 2015 \(S82\)](#).

CTX058. A moral judgment can be produced before the reasons later offered for it

Retained exact relation: **implementation family**.

Objects compared: Intuition-first and post-hoc production sequence / Social Intuitionist Model.

NDT: NDT distinguishes the support that actually formed a judgment from reasoning or explanation supplied afterward. It does not require every judgment to be intuition-first or deny that reasoning can sometimes contribute.

External claim: Haidt's Social Intuitionist Model gives rapid moral intuition a central judgment-producing role and treats much moral reasoning as a later, socially situated attempt to justify or influence judgments.

Sources: [Haidt, 2001 \(S83\)](#).

CTX059. Moral concern can extend across a wider or narrower range of entities

Retained exact relation: **empirical signature**.

Objects compared: Relational-positioning breadth signature / Moral Expansiveness Scale.

NDT: Patterns in the breadth and depth of moral concern may indicate which objects are relationally positioned, but they do not directly identify a positioning route, inherent valence, or contribution to a particular judgment.

External claim: Crimston and colleagues measure variation in the breadth and depth with which people extend moral concern across human, nonhuman, and environmental entities.

Sources: [Crimston et al., 2016 \(S84\)](#).

CTX060. Provenance and foundation taxonomies can classify moral cognition differently

Retained exact relation: **empirical rival partition**.

Objects compared: Provenance partition versus foundation-domain taxonomy / Five-foundation moral-domain taxonomy.

NDT: NDT classifies individual judgment occurrences by the relationship provenance that actually helped form them; it does not require the moral domain to be exhausted by a fixed list of foundation-linked intuition sets.

External claim: Graham and colleagues' 2011 article develops a five-foundation account and measure while allowing foundations to be added, divided, or combined. The five dimensions are not an exhaustive final inventory for mature MFT.

Qualification: The retained empirical-rival relation concerns the source-bounded foundation partition. Comparisons of mature MFT must also retain the open inventory, learning, and later equality/proportionality distinction discussed in §4. No five-foundation exhaustiveness finding is asserted.

Sources: [Graham et al., 2011 \(S16\)](#).

Appendix B. Source key for the comparison reference

Source keys distinguish passages used for different claims within the same publication. Repeated works are not separate theories. Exact passage locators are supplied where available; a work's complete page range is not an exact claim locator. The Master Bibliography supplies consolidated publication entries. The mature-source discussion also cites Atari and colleagues (2023) and Gray and colleagues (2022) directly in the main text; those additions supply developmental context without adding a comparison record.

S01. Aristotle, n.d.-a

Aristotle, *De Motu Animalium*, Part 7 (trans. Farquharson). [Source](#).

S02. Hume, n.d.-a

David Hume, *A Treatise of Human Nature*, 2.3.3.3, SBN 414. [Source](#).

S03. Makinson & van der Torre, 2007

David Makinson & Leendert van der Torre (2007), "What is Input/Output Logic?" [Source](#).

S04. Searle, 1964

John R. Searle (1964), "How to Derive Ought from Is." [Source](#).

S05. Dung, 1995-a

Phan Minh Dung (1995), "On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games," *Artificial Intelligence* 77(2), 321–357, Definition 20, p. 329. [Source](#).

S06. Bentham, n.d.

Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation*, ch. I, §II. [Source](#).

S07. Kant, n.d.

Kant, I. (1785). *Groundwork of the Metaphysics of Morals*. Formula of Humanity, Academy 4:429. The link identifies the registered online text; translation-specific wording should be checked against the edition used. [Source](#).

S08. Ross, 1930

W. D. Ross (1930), *The Right and the Good*, ch. 2, p. 28. [Source](#).

S10. Held, 1995

Virginia Held (1995), “The Meshing of Care and Justice.” [Source](#).

S11. Scanlon, 1998

T. M. Scanlon (1998), What We Owe to Each Other, ch. 5, p. 189. [Source](#).

S15. Dancy, 2004

Jonathan Dancy (2004), Ethics Without Principles, ch. 1, p. 7. [Source](#).

S16. Graham et al., 2011

Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, 101(2), 366–385. The registered five-foundation statement is on p. 366; the functional definition, factor comparison, and revisable inventory are on pp. 368, 372–373, and 382. [Source](#).

S17. Schein & Gray, 2018-a

Schein, C., & Gray, K. (2018). The theory of dyadic morality: Reinventing moral judgment by redefining harm. *Personality and Social Psychology Review*, 22(1), 32–70. Template statement: p. 32. [Source](#).

S18. Turiel, 1983-a

Elliot Turiel (1983), *The Development of Social Knowledge: Morality and Convention*, ch. 4, p. 61. [Source](#).

S19. Curry, Mullins & Whitehouse, 2019

Curry, Mullins & Whitehouse (2019), “Is It Good to Cooperate?” [Source](#).

S20. Rai & Fiske, 2011-a

Rai & Fiske (2011), “Moral Psychology Is Relationship Regulation,” *Psychological Review* 118(1), 57–75, p. 58. [Source](#).

S21. Kleiman-Weiner, Saxe & Tenenbaum, 2017

Kleiman-Weiner, Saxe & Tenenbaum (2017), “Learning a Commonsense Moral Theory,” *Cognition* 167, 107–123, author manuscript p. 2. [Source](#).

S38. Wright, 1985

Richard W. Wright (1985), ‘Causation in Tort Law,’ *California Law Review* 73(6), 1735–1828, quote p. 1774. [Source](#).

S39. Halpern & Pearl, 2005

Joseph Y. Halpern & Judea Pearl (2005), 'Causes and Explanations: A Structural-Model Approach. Part I: Causes,' *British Journal for the Philosophy of Science* 56(4), 843–887. [Source](#).

S40. Levine et al., 2026

Sydney Levine, Nick Chater, Joshua B. Tenenbaum & Fiery Cushman (2026), 'Resource-rational contractualism: A triple theory of moral cognition,' *Behavioral and Brain Sciences* 49, e133, manuscript p. 13. [Source](#).

S44. Bakhtin, 1984

Mikhail Bakhtin (1984), *Problems of Dostoevsky's Poetics*, ed. and trans. Caryl Emerson, ch. 1, p. 6. [Source](#).

S45. Mackie, 1977

J. L. Mackie (1977), *Ethics: Inventing Right and Wrong*, ch. 1, §2, p. 17. [Source](#).

S46. Anscombe, 1957

G. E. M. Anscombe (1957), *Intention*, §33, p. 60. [Source](#).

S47. Davidson, 1978/2006

Davidson, D. (2006). Intending. In E. Lepore & K. Ludwig (Eds.), *The essential Davidson* (pp. 122–137). Oxford University Press. Original essay published 1978; cited passage p. 132 of the 2006 collection. [Publication record](#).

S48. Hume, n.d.-b

David Hume, *A Treatise of Human Nature*, 3.1.1.27, SBN 469–470. [Source](#).

S49. Hume, n.d.-c

David Hume, *A Treatise of Human Nature*, 2.3.3.8, SBN 417. [Source](#).

S50. Haidt & Joseph, 2004

Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, 133(4), 55–66. Prepared-module comparison: p. 60. [Source](#).

S51. Graham et al., 2013

Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. In P. Devine & A. Plant (Eds.), *Advances*

in experimental social psychology (Vol. 47, pp. 55–130). Academic Press. The registered modularity passage is on p. 63; additional main-text locators identify the relevant qualifications. [Source](#).

S52. Schein & Gray, 2018-b

Schein, C., & Gray, K. (2018). The theory of dyadic morality: Reinventing moral judgment by redefining harm. *Personality and Social Psychology Review*, 22(1), 32–70. Fundamental-basis statement: p. 32. Central statistical qualification: p. 43. [Source](#).

S53. Turiel, 1983-b

Elliot Turiel (1983), *The Development of Social Knowledge: Morality and Convention*, ch. 4, p. 64. [Source](#).

S54. Rai & Fiske, 2011-b

Tage Shakti Rai & Alan Page Fiske (2011), “Moral Psychology Is Relationship Regulation,” *Psychological Review* 118(1), 57–75, p. 58. [Source](#).

S55. Ramezani & Xu, 2025-b

Aida Ramezani & Yang Xu (2025), “Moral Association Graph: A Cognitive Model for Automated Moral Inference,” *Topics in Cognitive Science* 17(1), 120–138, §5.1. [Source](#).

S57. Kant, 1785

Kant, I. (1785). *Groundwork of the Metaphysics of Morals*. Section I, shopkeeper contrast, Academy 4:397. The link identifies the registered online text; translation-specific wording should be checked against the edition used. [Source](#).

S58. Nichols, 2002

Shaun Nichols (2002), “Norms with Feeling: Towards a Psychological Account of Moral Judgment,” *Cognition* 84(2), 221–236. [Source](#).

S59. Wittgenstein, 1953

Ludwig Wittgenstein (1953), *Philosophical Investigations*, §§65–67 (family resemblance). [Source](#).

S61. Cushman & Young, 2011

Fiery Cushman & Liane Young (2011), “Patterns of Moral Judgment Derive From Nonmoral Psychological Representations,” *Cognitive Science* 35(6), 1052–1075. [Source](#).

S62. Van Bavel et al., 2012

Jay J. Van Bavel, Dominic J. Packer, Ingrid J. Haas, & William A. Cunningham (2012), “The Importance of Moral Construal,” PLOS ONE 7(11), e48693. [Source](#).

S63. Sinnott-Armstrong & Wheatley, 2014

Walter Sinnott-Armstrong & Thalia Wheatley (2014), “Are Moral Judgments Unified?” *Philosophical Psychology* 27(4), 451–474, abstract (p. 451). [Source](#).

S64. Kumar, 2016

Victor Kumar (2016), “The Empirical Identity of Moral Judgment,” *Philosophical Quarterly* 66(265), 783–804. [Source](#).

S65. Neta, 2015

Ram Neta (2015), “Basing Is Conjuring,” *Minds Online*. [Source](#).

S66. Neta, 2019

Ram Neta (2019), “The Basing Relation,” *The Philosophical Review* 128(2), 179–217, abstract (p. 179). [Source](#).

S67. McCain, 2012

Kevin McCain (2012), “The Interventionist Account of Causation and the Basing Relation,” *Philosophical Studies* 159(3), 357–382, abstract (p. 357). [Source](#).

S68. Lord & Sylvan, 2020

Lord, E., & Sylvan, K. (2020). Prime time (for the basing relation). In J. A. Carter & P. Bondy (Eds.), *Well-founded belief: New essays on the epistemic basing relation* (pp. 141–173). Routledge. Cited locations: authors’ abstract; §2.2, pp. 148–150. The publisher gives this author order and a 2020 copyright date; the source unit remains indexed as Sylvan and Lord (2019). [Author source](#); [publication record](#).

S69. Silva, 2015

Paul Silva Jr. (2015), “Does Doxastic Justification Have a Basing Requirement?” *Australasian Journal of Philosophy* 93(2), 371–387. [Source](#).

S77. Dehaene & Naccache, 2001

Stanislas Dehaene & Lionel Naccache (2001), ‘Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework,’ *Cognition* 79(1–2), 1–37, §6.2, p. 30. [Source](#).

S78. Gaertner et al., 1989

Samuel L. Gaertner, Jeffrey Mann, Audrey Murrell & John F. Dovidio (1989), 'Reducing Intergroup Bias: The Benefits of Recategorization,' *Journal of Personality and Social Psychology* 57(2), 239–249, abstract, p. 239. [Source](#).

S79. Laurenceau et al., 1998

Jean-Philippe Laurenceau, Lisa Feldman Barrett & Paula R. Pietromonaco (1998), 'Intimacy as an Interpersonal Process: The Importance of Self-Disclosure, Partner Disclosure, and Perceived Partner Responsiveness in Interpersonal Exchanges,' *Journal of Personality and Social Psychology* 74(5), 1238–1251, discussion, p. 1247. [Source](#).

S80. Batson et al., 1995

C. Daniel Batson, Tricia R. Klein, Lori Highberger & Laura L. Shaw (1995), 'Immorality from Empathy-Induced Altruism: When Compassion and Justice Conflict,' *Journal of Personality and Social Psychology* 68(6), 1042–1054, abstract, p. 1042. [Source](#).

S81. Levinas, 1961/1969

Emmanuel Levinas (1961/1969; 1991 Springer edition), 'Ethics and the Face,' in *Totality and Infinity: An Essay on Exteriority*, trans. Alphonso Lingis, p. 198. [Source](#).

S82. Uhlmann et al., 2015

Eric Luis Uhlmann, David A. Pizarro & Daniel Diermeier (2015), 'A Person-Centered Approach to Moral Judgment,' *Perspectives on Psychological Science* 10(1), 72–81, abstract, p. 72. [Source](#).

S83. Haidt, 2001

Jonathan Haidt (2001), 'The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment,' *Psychological Review* 108(4), 814–834, central claim, p. 817. [Source](#).

S84. Crimston et al., 2016

Daniel Crimston, Paul G. Bain, Matthew J. Hornsey & Brock Bastian (2016), 'Moral Expansiveness: Examining Variability in the Extension of the Moral World,' *Journal of Personality and Social Psychology* 111(4), 636–653, abstract, p. 636. [Source](#).

Exact passage locators remain unavailable in the source register for the input/output introduction, institutional promise article, care-and-justice article, cooperation article, structural actual-causation article, sentimental-rules article, nonmoral-representation article, flexible-construal article, natural-kind article, Neta's 2015 discussion, and Silva's basing article. Their entries identify works and bounded claims, not a completed passage-by-passage verification of every source.