

Normative Detachment Theory in Context

Literature Review 0.13

Date: 3 September 2026 **Canonical comparison authority:** Dataset 1.20 (48 records) **Data contract:** Data Contract 1.13
Source authority: Source Register 2.15 **Relation authority:** Taxonomy 2.2 **Status:** Unratified v52 RC3 candidate
projection; unmerged and unreleased

1. Comparison discipline

This review compares bounded claims at declared units, explananda, roles, explanatory levels, modal strengths, and terminal kinds. It does not score whole thinkers or treat shared vocabulary as agreement. Contradiction is a legitimate finding. A source is not compatible merely because it does not explicitly deny an NDT condition. Usable machinery and a rival or conflicting claim remain separately visible. Dataset 1.20 controls every exact relation; reader-facing language cannot rename it.

NDT's comparison target is usually the formation and cognitive type of an individual judgment occurrence. A source focused on behavior, moral truth, moral worth, domain content, lexical usage, or an empirical response profile may still matter, but the level difference must be recorded rather than converted into consilience.

2. Exact relation vocabulary

Governed relation: contradicts NDT

Meaning: A source-backed claim conflicts with an NDT constitutive or universal claim at the same explanatory level. Any separately usable component must be recorded as a distinct aspect rather than used to hide the conflict.

Current records: 4

Governed relation: empirical rival partition

Meaning: The source offers a competing empirical grouping or explanation for the same observations. Comparative tests must be declared in advance.

Current records: 4

Governed relation: tokenwise cross-classification

Meaning: The source and NDT can classify individual tokens differently because they use different typing criteria; this is not agreement merely because some outputs coincide.

Current records: 1

Governed relation: restricted decision-model family

Meaning: The external proposal occupies a defined region of NDT model space; NDT also contains cases outside that region.

Current records: 9

Governed relation: exact submodule

Meaning: One external formal operation is isomorphic to one NDT layer, not to the whole theory.

Current records: 2

Governed relation: implementation family

Meaning: The source supplies content or computation that may implement an NDT operation, but does not by itself determine output or moral type.

Current records: 10

Governed relation: empirical signature

Meaning: The source describes an observable judgment profile that may evidence, but does not constitute, an NDT process or type.

Current records: 6

Governed relation: output-only approximation

Meaning: The source can match labels, contents, or individual cases without recovering the judgment-producing derivation.

Current records: 1

Governed relation: scope boundary

Meaning: The source addresses a different target, level, or domain. The record states the boundary without turning non-contradiction into compatibility.

Current records: 5

Governed relation: source claim underdetermined

Meaning: The source language does not support a stable same-level classification under the current evidence. The record states what would be needed and does not simulate agreement or conflict.

Current records: 2

Governed relation: conceptual context

Meaning: The source offers philosophical or conceptual context but supplies neither an NDT computation nor a same-level empirical or constitutive rival.

Current records: 3

Governed relation: lineage

Meaning: The source is relevant to the development or history of an NDT idea without thereby establishing derivation, agreement, or explanatory equivalence.

Current records: 1

3. Findings that must remain paired

- **Rai and Fiske:** the four relational models can cross-classify individual tokens, but their universal relational-motive account also conflicts with NDT's valence-positioning-agency-and-provenance criterion.
- **Theory of Dyadic Morality:** dyadic content can be an implementation resource while harm as the fundamental moral template conflicts with NDT.
- **Turiel:** criterion-judgment tasks provide an empirical signature, while treating their response profile as constitutive of moral domain type conflicts with NDT.
- **Moral Foundations Theory:** CTX011 records prepared/related-module machinery as an implementation family; CTX060 separately records the five-foundation domain taxonomy as an empirical rival partition.

- **Nichols:** affect-plus-rule machinery can appear within NDT, while the sentimental-rules account supplies a rival empirical partition.

4. Bounded RC3 correction

Dataset 1.20 preserves every exact relation, external source binding, comparison decision, reader gloss, Proposition Registry pointer, and Empirical Audit pointer from Dataset 1.19. It refreshes only twelve current Formal Specification pointers from 1.6.1 to 1.6.2. It adds no search, source, external claim, record, relation, or novelty conclusion; the cell-level boundary is recorded in the Dataset 1.20 bounded-correction audit.

Standing-mediated positioning remains retired. Identity-mediated positioning includes shared, contrastive, and complementary-role forms; intimacy and direct encounter remain separate routes. Moral-expansiveness and empathy results remain indirect signatures, not direct measures of the current constructs.

5. Canonical comparison records

5.1 contradicts NDT

Response patterns do not define a domain / Criterion-judgment domain identification

Record: CTX013 **NDT claim:** 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. **Source-backed claim:** Social Domain Theory treats criterion judgments as definitional features that distinguish moral and conventional domains. **Exact relation:** contradicts NDT. **Why:** NDT's model-backed same-signature/different-provenance cases deny that criterion-judgment patterns identify token type. Medium confidence records the possibility that 'definitional' is intended diagnostically rather than constitutively. **Negation test:** A criterion-identification position predicts that rule independence, generalizability, and related response patterns settle a judgment's domain type. **Discriminating consequence:** Holding criterion judgments fixed while varying evaluative organization and provenance, NDT predicts tokens of different type with the same response pattern. **Primary sources:** Turiel, 1983-b

Cooperation does not constitute moral type / Cooperation constitution of morality

Record: CTX014 **NDT claim:** Cooperation can shape judgments, but NDT permits cooperative judgments that are nonmoral and moral judgments that are not cooperative solutions. **Source-backed claim:** Morality-as-Cooperation says morality consists of biological and cultural solutions to recurrent cooperation problems. **Exact relation:** contradicts NDT. **Why:** Morality-as-Cooperation constitutes the moral domain through recurrent cooperative problems, while NDT permits cooperative nonmoral tokens and moral-provenance tokens outside that taxonomy. The claims conflict at the domain-constitution level once NDT token classifications are aggregated to the same declared unit. The source does not state a necessary-and-sufficient judgment-token classifier, so no such classifier is attributed. **Negation test:** At the domain-constitution level, Morality-as-Cooperation predicts that biological and cultural solutions to recurrent cooperation problems exhaustively organize the moral domain. NDT denies that exhaustiveness after level-aligned aggregation; this test does not attribute a judgment-token classifier that the source does not state. **Discriminating consequence:** Compare the cooperation taxonomy with provenance at declared levels: aggregate token-level NDT classifications to the domain taxonomy, then test for cooperative nonmoral tokens and moral-provenance tokens not captured by the cooperation categories without treating granularity difference as automatic refutation. **Primary sources:** Curry, Mullins & Whitehouse, 2019

Harm templates do not constitute morality / Dyadic harm as moral basis

Record: CTX031 **NDT claim:** NDT permits moral judgments without a dyadic harm template and dyadic harm appraisals that lack qualifying moral provenance. **Source-backed claim:** The Theory of Dyadic Morality treats perceived harm as the fundamental basis by which moral judgments are made and maintained. **Exact relation:** contradicts NDT. **Why:** The source's basis claim conflicts with NDT's model-backed non-dyadic moral and dyadic-without-provenance witnesses. **Negation test:** A dyadic-basis position predicts that perceived intentional-agent/vulnerable-patient harm is necessary for,

and grounds, moral judgment. **Discriminating consequence:** NDT predicts moral tokens without the dyadic template and dyadic appraisals that remain nonmoral because qualifying provenance is absent. **Primary sources:** Schein & Gray, 2018-b

Four relational models do not exhaust moral qualification / Four relational models as universal moral basis

Record: CTX032 **NDT claim:** 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. **Source-backed claim:** Relationship Regulation Theory says people across cultures base moral judgment and behavior on the same four relational motives. **Exact relation:** contradicts NDT. **Why:** The universal-basis claim conflicts with NDT's two-way crossed classification between relational-model type and the Gate 2 relationship, subject, access, and exact local-contribution profile. **Negation test:** Negate NDT at the domain-constitution level by establishing that the four relational motives exhaust morally qualifying represented relationships and that model membership is sufficient for qualification. **Discriminating consequence:** NDT predicts qualifying relationships outside the four models and four-model instances that fail one or more independent valence, positioning, agency, subject, access, or contribution conditions. **Primary sources:** Rai & Fiske, 2011-b

5.2 empirical rival partition

Feeling can help build a moral judgment without defining morality / Sentimental rules account of the moral/conventional distinction

Record: CTX036 **NDT claim:** Affect-linked prohibitions may help form a moral judgment, but affect, prohibition, and their conjunction do not determine moral type without relationship provenance. **Source-backed claim:** Nichols explains the moral/conventional distinction through prohibited-action information together with an affective mechanism. **Exact relation:** empirical rival partition. **Why:** Nichols supplies affect-plus-prohibition machinery and a task-level account of the moral/conventional distinction. Those mechanisms may be usable within NDT, but the cited source does not state that their conjunction is necessary and sufficient for NDT's judgment-token type. The empirical rivalry is therefore limited to a preregistered comparison between the task partition and provenance, not a source-attributed constitutive identity. **Negation test:** At the task level, the sentimental-rules account predicts that changing prohibited-action information or the relevant affective mechanism changes performance on moral/conventional distinction tasks; this test does not attribute an NDT-style necessary-and-sufficient type criterion to Nichols. **Discriminating consequence:** Crossed tests can hold prohibition and affect fixed while varying relationship provenance, and hold provenance fixed while varying affect, then compare the preregistered task and provenance partitions. Task success alone does not establish moral cognitive type. **Primary sources:** Nichols, 2002

Can producing history unify moral judgment? / Fine-grained moral-judgment taxonomy

Record: CTX048 **NDT claim:** 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. **Source-backed 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. **Exact relation:** empirical rival partition. **Why:** The source proposes a finer-grained empirical partition of moral judgment, while NDT proposes a provenance-based unifier. Shared dissatisfaction with ordinary vocabulary is not evidence of agreement; the rival partitions must be compared on held-out intervention, transfer, process, and explanatory-stability tests. **Negation test:** A view that finds a stable moral kind using content, behavior, character, policy, or another independently specified partition denies the need for NDT's proposed unifier. **Discriminating consequence:** Compare pre-specified partitions on held-out intervention, transfer, process, and explanatory stability. Persistent failure or outperformance of the provenance partition counts against NDT's kind proposal. **Primary sources:** Sinnott-Armstrong & Wheatley, 2014

What kind of unity does moral judgment have? / Moral judgment as a natural kind

Record: CTX049 **NDT claim:** NDT locates unity in represented relationship organization and actual contributing support while allowing different contents, emotions, and implementations. **Source-backed 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. **Exact relation:** empirical rival partition. **Why:** The homeostatic-kind and provenance accounts propose different and independently variable unifiers for moral judgment. Neither is a component of the other; they are rival empirical partitions answerable to crossed and held-out tests. **Negation test:** A homeostatic belief-emotion partition that explains the relevant cases without any added gain from producing history challenges NDT's proposed unifier. **Discriminating consequence:** A crossed design can test whether provenance-based or belief-emotion and homeostatic classification produces the more stable causal and explanatory grouping. **Primary sources:** Kumar, 2016

Provenance and foundation taxonomies can classify moral cognition differently / Five-foundation moral-domain taxonomy

Record: CTX060 **NDT claim:** 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. **Source-backed claim:** Graham and colleagues describe five universally available, though variably developed, sets of moral intuitions: care, fairness, loyalty, authority, and purity. **Exact relation:** empirical rival partition. **Why:** The five-foundation proposal partitions the moral domain by recurrent intuition families, whereas NDT partitions individual judgment occurrences by actually contributing relationship provenance. The levels must be aligned rather than treated as an automatic contradiction: aggregate preregistered NDT token classifications to the same sampling unit before comparing coverage, predictive stability, and intervention sensitivity. This rival record is independent of CTX011's modular implementation claim. **Negation test:** At its stated domain level, the five-foundation taxonomy predicts that moral-intuition variation is organized by the five universally available foundation sets; an observation pattern requiring a non-foundation partition, after measurement and level alignment, denies that exhaustiveness claim. **Discriminating consequence:** Preregister foundation and provenance measures, classify held-out judgment tokens independently, then aggregate provenance results to the domain-level unit. Systematic same-foundation/different-provenance and different-foundation/same-provenance cases, plus comparative intervention and prediction tests, discriminate the partitions without treating a granularity mismatch as a result. **Primary sources:** Graham et al., 2011

5.3 tokenwise cross-classification***What makes a represented relationship morally qualifying / Relational-model individuation of moral motives***

Record: CTX015 **NDT claim:** 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. **Source-backed claim:** This construct distinguishes moral motives according to which of four basic kinds of social relationship they regulate. **Exact relation:** tokenwise cross-classification. **Why:** Individual tokens can receive both a four-way relational-model label and an NDT qualification, but the classifications vary independently. The external taxonomy neither fixes inherent valence, identifies an operative positioning bridge, establishes perceived agency, nor supplies the subject, access, and contribution conditions. **Negation test:** Negate NDT at the matched level by showing that the relational-model type itself exhaustively constitutes the individual token's morally qualifying relationship, regardless of the independent valence, positioning, agency, subject, access, and contribution conditions. **Discriminating consequence:** Cross relational-model type with positive-only, negative-only, mixed, and neither valence profiles; the three positioning routes and identity forms; perceived agency; subject-gate status; access; and exact local contribution. NDT predicts qualifying relationships outside the four models and four-model instances that fail qualification. **Primary sources:** Rai & Fiske, 2011-a

5.4 restricted decision-model family

Goal-and-means reasoning / Practical-inference guide structure

Record: CTX001 **NDT claim:** 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. **Source-backed claim:** Practical inference organizes ends or goods and available means into guidance about what to do; its terminus varies across Aristotle, Anscombe, and Davidson. **Exact relation:** restricted decision-model family. **Why:** The sources converge on prospective guidance from represented ends or goods and means, but differ over terminal kind. NDT selects judgment-derivation tokens as its analytic unit and makes no identity or enactment claim. **Negation test:** A noninferentialist or rationalization-only account denies that the premises generate prospective guidance rather than merely redescribing an independently produced action. **Discriminating consequence:** NDT admits this guide family while leaving encoded evaluation, hard requirements, moral provenance, and polyphony as NDT cases outside the match; it predicts nothing about enactment. **Primary sources:** Aristotle, n.d.-a; Anscombe, 1957; Davidson, 1978

Institution-based requirements / Institutional promise obligation

Record: CTX004 **NDT claim:** When an institution, its rule, and the conditions for that rule are represented as active, the institution can supply or shape a requirement. **Source-backed claim:** A promise made within an active institution can create an obligation because of the institution's rules. **Exact relation:** restricted decision-model family. **Why:** The promise case is exact only when the institution, uptake, feasibility, and absence of an undefeated exception are represented. **Negation test:** A model that denies any represented constitutive institutional bridge predicts that an institution and its active rule cannot generate an obligation from a promise. **Discriminating consequence:** The match holds only when the institution, constitutive rule, and uptake conditions are represented as active; institutional compliance still does not settle moral type. **Primary sources:** Searle, 1964

Outcome-maximizing guidance / Principle of utility

Record: CTX006 **NDT claim:** 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. **Source-backed claim:** The principle of utility judges actions by whether they increase or reduce the happiness of those affected. **Exact relation:** restricted decision-model family. **Why:** The principle of utility defines an outcome-ordered region of NDT model space. NDT also represents non-outcome guidance and remains silent on which first-order standard is correct. **Negation test:** A nonconsequentialist guide position predicts cases in which represented utility maximization does not supply the operative guidance. **Discriminating consequence:** The comparison predicts a bounded family of utility-maximizing NDT tokens alongside tokens guided by requirements, relationships, or other represented standards. **Primary sources:** Bentham, n.d.

Requirements that outcomes cannot outweigh / Formula of Humanity

Record: CTX007 **NDT claim:** An applicable requirement that has not been overridden can rule out an action even when that action would produce a better outcome. **Source-backed claim:** The Formula of Humanity requires treating rational nature as an end in itself. **Exact relation:** restricted decision-model family. **Why:** A represented humanity requirement can occupy NDT's hard-requirement family and disqualify an outcome-better action; NDT does not adjudicate Kant's first-order supremacy claim. **Negation test:** A compensatory outcome-only decision model denies that a represented humanity requirement can disqualify an action whose outcome scores higher. **Discriminating consequence:** Within this family, an applicable, undefeated humanity requirement can change the chosen action without any change in outcome ranking. **Primary sources:** Kant, n.d.

Overridden duties that still matter / Prima-facie duty and moral residue

Record: CTX008 **NDT claim:** 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. **Source-backed claim:** A prima-facie duty can be overridden while the unfulfilled duty still matters morally. **Exact relation:** restricted decision-model family. **Why:** The family maps to priority-sensitive requirements; NDT additionally

requires support, provenance, and well-formed priority grounds. **Negation test:** A no-persisting-conflict model, on which only the actual duty remains after priority is resolved, denies that an overridden duty can generate residue.

Discriminating consequence: Holding the chosen action fixed, NDT distinguishes an action with no persisting violated requirement from a licensed override that leaves violation and residue; this contrast does not posit a primitive permission output. **Primary sources:** Ross, 1930

Relationship-grounded guidance / Care as a moral framework

Record: CTX009 **NDT claim:** A represented relationship can create the requirement that guides action, rather than merely motivating someone to follow a rule fixed independently. **Source-backed claim:** Care can be treated as the wider moral framework within which justice should operate. **Exact relation:** restricted decision-model family. **Why:** Care-priority and justice-priority orderings can each be represented as guide families. NDT separately asks whether a represented relationship constructs the operative requirement. **Negation test:** A justice-priority position denies that care should supply the wider framework within which justice operates. **Discriminating consequence:** The comparison permits either priority ordering and separately tests whether relationship-grounded organization actually generated the guidance. **Primary sources:** Held, 1995

Complaint-sensitive guidance / Reasonable rejectability

Record: CTX010 **NDT claim:** NDT can represent cases in which complaints or reasonable rejection guide a decision, but those cases do not define all normative thinking. **Source-backed claim:** Reasonable rejectability asks whether conduct can be justified to others on terms they could not reasonably reject. **Exact relation:** restricted decision-model family. **Why:** Reasonable rejectability defines a complaint-sensitive family within the domain of what people owe to one another. A minimax model is one witness, not a definition of reasonableness. **Negation test:** Within the what-we-owe-to-each-other domain, a complaint-insensitive position denies that reasonable rejectability is necessary for permissible guidance. **Discriminating consequence:** NDT represents complaint-sensitive and other guide families without treating contractualist structure as exhaustive of normative cognition. **Primary sources:** Scanlon, 1998

Bargaining-based decision models / Resource-rational agreement approximation

Record: CTX022 **NDT claim:** NDT can represent bargaining-based standards and decision routes, but using them does not determine whether a particular judgment is moral. **Source-backed claim:** Resource-rational contractualism treats many action- and welfare-based moral judgments as efficient approximations of agreement-based judgment. **Exact relation:** restricted decision-model family. **Why:** Agreement-derived welfare standards map to outcome ordering, action standards to requirements, and explicit, simulated, or cached agreement processing to decision routes; NDT tests moral provenance separately. **Negation test:** An independent-mechanism account denies that a substantial class of action- and welfare-based judgments derives from agreement through resource-rational approximation. **Discriminating consequence:** NDT locates the agreement-approximation family without treating it as exhaustive or constitutive of moral type. **Primary sources:** Levine et al., 2026

Case-by-case guidance / Moral particularism

Record: CTX024 **NDT claim:** NDT allows guidance to depend on the particulars of a case while also allowing models that use general principles. **Source-backed claim:** Moral particularism holds that moral thought does not depend on exceptionless general principles. **Exact relation:** restricted decision-model family. **Why:** Particularist reasoning occupies the NDT family in which guidance need not derive from exceptionless general principles; NDT also permits principle-governed families and determines moral provenance independently. **Negation test:** A moral generalist position predicts that sound moral thought requires a sufficient stock of exceptionless general principles. **Discriminating consequence:** NDT permits case-sensitive and principle-governed model families without endorsing Dancy's wider holism, realism, or epistemology. **Primary sources:** Dancy, 2004

5.5 exact submodule

Cautious challenge resolution / Grounded argumentation semantics

Record: CTX005 **NDT claim:** 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. **Source-backed claim:** Grounded semantics resolves a network of attacks into one cautious set of accepted claims. **Exact relation:** exact submodule. **Why:** The fixed-point calculation is exact for the admitted attack graph; NDT adds admission, target typing, provenance, and moral classification. **Negation test:** A preferred-extension resolver on the same admitted attack graph denies that the unique least-fixed-point grounded extension is the selected result, instead returning maximal admissible extensions when the semantics diverge. **Discriminating consequence:** On graphs where the grounded extension is properly contained in at least one preferred extension, NDT's grounded reference semantics withholds arguments admitted under that preferred extension; when the unique preferred extension equals the grounded extension, the outputs coincide. **Primary sources:** Dung, 1995-a

NESS credit is one part of the contribution test / NESS minimal-sufficient-set membership

Record: CTX020 **NDT claim:** 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. **Source-backed claim:** The NESS condition identifies a cause as a necessary member of an actually instantiated set sufficient for the outcome. **Exact relation:** exact submodule. **Why:** The direct match is restricted to ProfileNESSOn_s: membership as a necessary element within one actually instantiated sufficient profile. It does not transfer to NDT's complete ActualContribution predicate or settle route realization, object binding, controlled effect, same-chain provenance, or moral type. **Negation test:** A simple but-for account denies that membership in one actually instantiated sufficient set can establish contribution when another sufficient route is also active. **Discriminating consequence:** Under redundant support, a relationship profile may receive ProfileNESSOn_s credit when it is necessary within at least one active sufficient profile; the remaining NDT contribution conditions must still be established independently. **Primary sources:** Wright, 1985

5.6 implementation family

Rule-based guidance / Input/output logic

Record: CTX003 **NDT claim:** 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. **Source-backed claim:** Input/output logic maps stated conditions to normative conclusions without treating the governing rule as an ordinary if-then proposition. **Exact relation:** implementation family. **Why:** Input/output logic supplies a non-truth-functional norm-output transformation that can implement part of NDT's rule-conditioned candidate formation. The cited source does not establish an isomorphism to a named NDT primitive, and it does not supply NDT's frame, source, addressee, admission, provenance, defeat, permission, or moral-type conditions. **Negation test:** An account that reduces normative output relations to ordinary material implication denies the distinctive non-truth-functional transformation that this implementation can supply. **Discriminating consequence:** Selected input/output and NDT rule-conditioned candidates may coincide, but no I/O verdict transfers without separate NDT checks for typing, source, addressee, admission, provenance, defeat, output survival, and moral type; this record adds no primitive permission output. **Primary sources:** Makinson & van der Torre, 2007

Content does not settle moral type / Prepared-module moral-intuition architecture

Record: CTX011 **NDT claim:** NDT denies that activation of a prepared moral-content module is enough to determine whether a particular judgment is moral; qualifying provenance remains necessary. **Source-backed 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. **Exact relation:** implementation family. **Why:** The cited prepared-module and related-module claims supply a possible content-generation and response-constraint implementation inside an NDT route. They do not determine judgment-token moral type, establish whole-theory compatibility, or support

the five-foundation domain taxonomy, which is separately represented in CTX060. **Negation test:** An architecture with no recurrent prepared foundation-like content or processing family denies this proposed implementation resource without thereby denying NDT's provenance criterion. **Discriminating consequence:** Independently measure foundation-linked activation and NDT provenance. Foundation-linked modules may predict content or processing while tokens with similar activation differ in moral type when their actually contributing provenance differs; the domain-taxonomy comparison remains separately testable in CTX060. **Primary sources:** Haidt & Joseph, 2004; Graham et al., 2013

Representing a harmful dyad / Dyadic harm template

Record: CTX012 **NDT claim:** 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. **Source-backed claim:** The dyadic harm template represents harm as an intentional actor damaging a vulnerable patient. **Exact relation:** implementation family. **Why:** The dyadic template supplies agent, patient, harm, and causation content that NDT can represent, while NDT's provenance conditions determine type independently. **Negation test:** A valence-only or role-collapsing harm representation denies that distinguishing intentional agent, vulnerable patient, and causal damage adds a necessary representational structure. **Discriminating consequence:** NDT can preserve the full dyadic template as appraisal content without treating that content as necessary or sufficient for moral type. **Primary sources:** Schein & Gray, 2018-a

Learning other people's moral weights / Hierarchical Bayesian moral-weight inference

Record: CTX016 **NDT claim:** 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. **Source-backed claim:** A hierarchical Bayesian model uses observed behavior to estimate how strongly other people weight abstract moral principles. **Exact relation:** implementation family. **Why:** Hierarchical Bayesian inference supplies latent moral-weight estimates that can populate NDT outcome evaluation, while NDT's provenance test remains separate. **Negation test:** A behaviorist-only learner that maps observed behavior directly to prediction without inferring latent principles or weights denies the Bayesian construct. **Discriminating consequence:** The learned weights can alter NDT's outcome ordering, while identical weights can occur in tokens of different moral type. **Primary sources:** Kleiman-Weiner, Saxe & Tenenbaum, 2017

Structural models of actual causation / Contingent-dependence structural actual causation

Record: CTX033 **NDT claim:** NDT needs a way to test actual contribution under redundant support, but it does not require one unique structural-equation definition of causation. **Source-backed claim:** Halpern–Pearl actual causation tests dependence under selected structural-model contingencies, making model choice part of the analysis. **Exact relation:** implementation family. **Why:** Structural-equation contingency analysis can implement NDT's actual-contribution check, but model and contingency choices are substantive and do not determine moral type. **Negation test:** A model-free or purely regularity-based causal account denies that selected structural contingencies are needed to determine actual causation. **Discriminating consequence:** Different causation implementations may agree on NDT contribution status; disagreement exposes an implementation-sensitive case rather than a change in moral-type criteria. No Halpern–Pearl verdict transfers until the structural model and contingency scheme are registered and shown adequate for the NDT target. **Primary sources:** Halpern & Pearl, 2005

A relationship route must be available to the operative process in order to contribute, but access does not confer moral type / Global-workspace broadcast availability

Record: CTX044 **NDT claim:** 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. **Source-backed claim:** The global-workspace framework describes selected and amplified contents becoming available to multiple receiving systems. **Exact relation:** implementation family. **Why:** Global-workspace broadcasting is one possible realization of functional availability. NDT permits local and nonconscious access models and separately tests whether the accessed typed edge actually contributed. **Negation test:** The comparison would become exact only if global-workspace broadcast were necessary and sufficient for NDT functional access and token-level local actual contribution across implementations. **Discriminating consequence:** Hold the relationship core and downstream edge fixed while enabling or blocking the declared workspace path. With no alternative path, blocking access prevents contribution through that path; restoring

access does not itself supply valence, positioning, agency, a bender effect, or moral type. **Primary sources:** Dehaene & Naccache, 2001

Shared identity is one identity-mediated way to position an object relative to a subject / Common-ingroup representation

Record: CTX045 **NDT claim:** 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. **Source-backed claim:** Gaertner and colleagues compared one-group, two-group, and separate-individual representations and found different patterns of bias reduction. **Exact relation:** implementation family. **Why:** Common-ingroup representation supplies a candidate mechanism and transfer signature for the shared form of identity-mediated positioning. A one-group manipulation or reduced bias does not validate the NDT route, establish moral type, or show token-level actual contribution; separate-individual results also bar a universal identity gate. **Negation test:** The relation would be exact only if common-ingroup representation supplied NDT's independently specified exact subject-object bridge and its contribution test, rather than a candidate shared-form implementation. **Discriminating consequence:** Hold objective membership fixed while manipulating the operative representation. The shared form predicts transfer tied to the one-group 'we'; contrastive or complementary-role identity, intimacy, or direct encounter may preserve positioning when shared identity is absent, and token-level contribution remains separately required. **Primary sources:** Gaertner et al., 1989

Intimacy positions an object through a particular relationship / Interpersonal-process intimacy

Record: CTX046 **NDT claim:** 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. **Source-backed claim:** Laurenceau and colleagues tested interaction-level relations among disclosure, perceived partner responsiveness, and experienced intimacy. **Exact relation:** implementation family. **Why:** The interpersonal-process model supplies relationship-specific candidate indicators that may implement IntimacyRP_s. Disclosure, responsiveness, and experienced intimacy do not by themselves validate the NDT route, establish moral type, or show token-level actual contribution. **Negation test:** The relation would be exact only if the interpersonal-process variables were necessary and sufficient for NDT intimacy positioning and its local contribution across domains. **Discriminating consequence:** Hold valence, agency, identity-mediated and encounter routes, familiarity, and downstream benders fixed while varying disclosure or responsiveness. A candidate intimacy effect should be partner-specific; another positioning route may remain, and contribution evidence is separately required. **Primary sources:** Laurenceau et al., 1998

A local difference-maker test for actual contribution / Interventionist epistemic basing

Record: CTX053 **NDT claim:** 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. **Source-backed claim:** McCain uses interventionist causation to distinguish a reason's nondeviant basing of a belief from mere causal ancestry, including under overdetermination. **Exact relation:** implementation family. **Why:** Interventionist machinery can implement a test separating causal ancestry from contribution. Its epistemic target does not supply NDT's frames, relationship conditions, bindings, or moral classification; the relation is limited to that declared implementation job. **Negation test:** A noncausal account of epistemic basing denies that an interventionist causal condition is necessary for a reason to be a belief's basis. **Discriminating consequence:** A representation that remains causally upstream but fails the registered intervention should not count under an NDT implementation that requires local difference-making; another sufficient route may still support the judgment. **Primary sources:** McCain, 2012

A moral judgment can be produced before the reasons later offered for it / Social Intuitionist Model

Record: CTX058 **NDT claim:** 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. **Source-backed 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. **Exact relation:** implementation family. **Why:** The model is a historical and mechanistic neighbor for separating producing support from later explanation. Its default sequence can implement NDT provenance questions but does not

provide NDT's moral-type criterion or make intuition necessary. **Negation test:** A report-identity account predicts that the reasons a person later offers are the producing basis whenever they match the judgment's content, regardless of process or intervention evidence. **Discriminating consequence:** Hold judgment content fixed while manipulating intuitive cues, opportunity for prior reasoning, later justification prompts, and social influence. Process and intervention evidence should distinguish producing from post-hoc routes without assigning moral type from sequence alone. **Primary sources:** Haidt, 2001

5.7 empirical signature

Response patterns can help identify moral and nonmoral normativity / Social Domain moral/conventional criterion-judgment paradigm

Record: CTX038 **NDT claim:** 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. **Source-backed claim:** Social Domain studies distinguish moral rules from conventions through patterns such as authority-independence, generalizability, and judged wrongness or seriousness across cases. **Exact relation:** empirical signature. **Why:** The paradigm supplies elicitation tasks and observable response profiles that can help test NDT's computational distinction between moral and nonmoral normativity. It remains a component rather than a restricted model family because NDT permits dissociations: the same criterion profile can arise from different provenance, and a moral token need not display every canonical response. CTX013 separately records the conflict that arises only when the profiles are treated as constitutive definitions of domain type. **Negation test:** A null identification result would find no reliable relation between independently measured relationship provenance and Social Domain response profiles once content, rule familiarity, authority cues, and demand characteristics are controlled. **Discriminating consequence:** A registered study can hold surface content constant, vary or measure qualifying relationship organization and actual contribution, and test whether Turiel-style criterion judgments add out-of-sample information about moral versus nonmoral normative provenance while allowing principled false positives and false negatives. **Primary sources:** Turiel, 1983-a; Turiel, 1983-b

Direct encounter positions the particular other through presence, expression, testimony, address, or condition / Empathy-induced particularized altruism

Record: CTX047 **NDT claim:** 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. **Source-backed claim:** Batson and colleagues induced empathy toward a particular person and found target-favoring allocations that could conflict with an independently recognized justice principle. **Exact relation:** empirical signature. **Why:** The experiment supplies a particularized-empathy signature potentially relevant to DirectEncounterRP_s. Empathy, allocation behavior, and judged partiality are not direct measures of the NDT route, moral type, or token-level actual contribution. **Negation test:** The relation would be an implementation family only if the manipulation independently realized the direct subject-object positioning bridge, rather than supplying an outcome signature that might arise through another process. **Discriminating consequence:** Hold access, identity-mediated positioning, intimacy, the relationship core, and the fairness rule fixed while varying empathic concern for one target. Direct encounter predicts target-specific weighting and possible conflict with a justice route; actual contribution and any bender edge remain separately required. **Primary sources:** Batson et al., 1995

The same action can be understood in different ways / Flexible moral construal

Record: CTX050 **NDT claim:** 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. **Source-backed 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. **Exact relation:** empirical signature. **Why:** The same-stimulus design supplies an observable signature and design pattern for flexible construal. It does not identify NDT's frame boundaries, relationship builders, or actual contribution, so no component-level compatibility is inferred. **Negation test:** A stimulus-fixed classifier predicts no systematic type-relevant change when the active evaluative goal changes while the described action remains fixed. **Discriminating consequence:** A stronger NDT study would independently measure the active relationship representation and its contribution, then test whether they explain classification beyond moral-versus-pragmatic instructions. **Primary sources:** Van Bavel et al., 2012

Nonmoral representations can help build a moral judgment / Nonmoral representations that shape moral judgment

Record: CTX051 **NDT claim:** Causal, intentional, and other domain-general representations may help build or bend a moral judgment without becoming what makes the judgment moral. **Source-backed claim:** Cushman and Young explain familiar patterns of moral judgment through nonmoral causal and intentional representations, including action-omission and means-side-effect contrasts. **Exact relation:** empirical signature. **Why:** The account provides observable evidence that domain-general causal and intentional representations shape moral judgment. That signature does not establish compatibility and does not distinguish moral from nonmoral normative judgment in NDT. **Negation test:** An architecture on which every representation influencing a moral judgment must itself be specifically moral denies this component relation. **Discriminating consequence:** Causal and intentional representations can be held fixed while qualifying relationship contribution varies, or varied while the moral relationship remains fixed. **Primary sources:** Cushman & Young, 2011

Character inferences supply a signature; the rival partition is NDT's test / Person-centered moral character inference

Record: CTX057 **NDT claim:** 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. **Source-backed 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. **Exact relation:** empirical signature. **Why:** Uhlmann and colleagues support person-centered character inference as an empirical phenomenon and observable signature; they do not supply a constitutive partition distinguishing moral from nonmoral cognitive type. The character-based rival partition used under the provenance-based kind-unity hypothesis is an NDT-constructed preregistered test, not a claim attributed to the source. **Negation test:** A strictly act-centered account predicts that judgments will not vary with diagnostic implications for an agent's moral character once act features and outcomes are held fixed. **Discriminating consequence:** Use source-grounded character-inference measures as candidate indicators, then independently predefine the NDT-constructed character and provenance partitions, hold out persons or scenarios, control complexity, and compare predictive and intervention stability. A character partition that matches or outperforms provenance without added provenance gain counts against the provenance-based kind-unity hypothesis. **Primary sources:** Uhlmann et al., 2015

Moral concern can extend across a wider or narrower range of entities / Moral Expansiveness Scale

Record: CTX059 **NDT claim:** 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. **Source-backed claim:** Crimston and colleagues measure variation in the breadth and depth with which people extend moral concern across human, nonhuman, and environmental entities. **Exact relation:** empirical signature. **Why:** Moral expansiveness is a useful observable signature and possible indicator family for inclusion breadth. It is not a direct measure of Relational Positioning, does not identify a route or identity form, and does not establish inherent valence or token-level actual contribution. **Negation test:** The relation would become a direct operational match only if the scale independently identified exact subject-object positioning, route type, and token-level contribution rather than reported breadth and depth of concern. **Discriminating consequence:** Validate expansiveness against independently identified Relational Positioning, route/form, valence, and contribution. Correlation without discriminant validity, intervention sensitivity, and process evidence does not identify the construct; systematic divergence counts against using the scale as its indicator. **Primary sources:** Crimston et al., 2016

5.8 output-only approximation

Association-based moral prediction / Association-based moral-relevance prediction

Record: CTX023 **NDT claim:** Association networks can approximate moral relevance or labels, but they do not reconstruct the decision process or the provenance of a particular judgment. **Source-backed claim:** The Moral Association Graph uses word associations to predict human moral relevance and moralization, with tested performance against baseline models. **Exact relation:** output-only approximation. **Why:** MAG predicts labels and relevance from association structure. NDT can

use such outputs as salience or evaluative inputs without treating them as a derivation of token type. **Negation test:** An association-noninformative or language-model-only position predicts that word-association structure adds no reliable moral-relevance prediction beyond the stated baselines. **Discriminating consequence:** MAG and NDT can agree on a label while differing on the decision route and provenance; prediction does not recover derivation. **Primary sources:** Ramezani & Xu, 2025-b

5.9 scope boundary

Generation–status separation / Normative and metaethical status

Record: CTX025 **NDT claim:** NDT's account of how a judgment is produced does not by itself decide whether that judgment is true, justified, correct, or objectively valid. **Source-backed claim:** Normative and metaethical status concerns whether moral judgments are true, justified, or objectively valid; Mackie's objective-value denial is one instance. **Exact relation:** scope boundary. **Why:** This is a narrow generation/status boundary, not a claim that NDT is compatible with every metaethical theory, theory of moral states, or theory of motivation. Such theories may conflict with NDT's production commitments even though no production result by itself settles truth, justification, correctness, or objective status. **Discriminating consequence:** No result about how a judgment is generated licenses an inference to its truth, justification, correctness, or objective status, and vice versa. **Primary sources:** Mackie, 1977

No ought from evaluatively empty inputs / Evaluatively empty is–ought barrier

Record: CTX030 **NDT claim:** When a description contains no evaluative ordering, standard, or requirement, NDT derives no prescriptive conclusion from that description alone. **Source-backed claim:** Hume's is–ought barrier requires explaining any move from ordinary descriptive propositions to a new ought-relation. **Exact relation:** scope boundary. **Why:** The exact proposition compared is limited to an evaluatively empty input not licensing a prescriptive conclusion by itself. Hume's passage motivates that barrier, but this reconstruction does not claim textual equivalence, a direct whole-theory match, or broader psychological agreement. **Negation test:** A descriptive-entailment position predicts that a prescriptive conclusion can follow from inputs stipulated to be evaluatively empty without any additional evaluative bridge. **Discriminating consequence:** For inputs stipulated to be evaluatively empty, the reconstructed barrier blocks the prescriptive transition; cognitively evaluative representations and broader claims about motivation fall outside the compared proposition. **Primary sources:** Hume, n.d.-b

Being available is not the same as helping form the judgment / Reasons one has, reasons why, and reasons for which

Record: CTX052 **NDT claim:** NDT separates considerations that were available or causally upstream from representations that actually helped form this particular judgment. **Source-backed 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. **Exact relation:** scope boundary. **Why:** Neta studies epistemic basing, whereas NDT uses local contribution to classify a judgment's moral type. The shared availability-versus-basis distinction is not enough to establish a component or analogy relation; the record states the difference in target and level. **Negation test:** An account that identifies an attitude's basis with any possessed reason that causally helps explain it collapses distinctions maintained by both views. **Discriminating consequence:** Hold judgment content fixed and compare classifications by availability, causal ancestry, reasons-for-which, and NDT local contribution on held-out intervention and process evidence. **Primary sources:** Neta, 2015; Neta, 2019

One contributing route must do both jobs / Prime View of believing for a normative reason

Record: CTX054 **NDT claim:** 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. **Source-backed 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. **Exact relation:** scope boundary. **Why:** The Prime View concerns normative-reason basing, whereas NDT concerns the provenance of a judgment's cognitive type. Rejecting a split witness is a shared structural concern, but it does not make one view a component of the other. The record therefore states the explanatory-level boundary. **Negation test:** The

Composite View predicts that independently satisfying the explanatory and normative clauses is sufficient.

Discriminating consequence: If one consideration explains the judgment while a different, inert consideration supplies moral-looking standing, NDT withholds moral type because the required work is split across routes. **Primary sources:** Sylvan & Lord, 2019

Does moral classification really require actual support? / No general basing requirement for doxastic justification

Record: CTX055 **NDT claim:** 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. **Source-backed 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. **Exact relation:** scope boundary. **Why:** Silva addresses doxastic justification, while NDT addresses descriptive moral type. The source supplies a challenge that motivates a separate NDT rival test, but does not itself become an NDT component or same-level analogue. **Negation test:** A universal basing-demand account predicts that doxastic justification always requires belief on an epistemically appropriate basis. **Discriminating consequence:** A parallel NDT rival would classify judgments by appropriately represented relationship content whether or not it contributed. If that simpler partition matches or outperforms provenance on pre-specified tests, NDT's contribution requirement loses comparative support. **Primary sources:** Silva, 2015

5.10 source claim underdetermined

What supplies motivation / Passional source requirement for motivation

Record: CTX002 **NDT claim:** 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. **Source-backed claim:** Hume holds that reason alone cannot motivate: an impulse must come from passion, including calm passions known mainly through their effects. **Exact relation:** source claim underdetermined. **Why:** A detachable-state reading and a broad functional reading classify NDT's encoded evaluative control differently. Specification 1.6.2 §§2.2–2.3 defines the encoded/attribution boundary, while §9 leaves its empirical identification and measurement open. **Negation test:** An anti-Humean motivational account on which a representative evaluative state can motivate without any passion denies the passional-source requirement. **Discriminating consequence:** The record requires an independently justified criterion for passion or desire before it can classify encoded evaluative control as agreement or conflict. **Primary sources:** Hume, n.d.-a; Hume, n.d.-c

Evaluation can be built into a representation's control role / Strict representation–passion separability

Record: CTX034 **NDT claim:** 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. **Source-backed 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. **Exact relation:** source claim underdetermined. **Why:** A strict universal separability thesis would conflict with NDT's encoded route, but the cited Humean material concerns motivation and allows calm passions known through their effects. It does not independently fix whether every evaluatively structured representation must receive its role from a functionally distinct passion. The previous contradiction label therefore exceeded the source; the record states the missing criterion instead. **Negation test:** Strict representation–passion separability predicts that every evaluatively action-guiding representation requires a functionally distinct passion to supply its motivating role; no encoded evaluative realization is possible. **Discriminating consequence:** A governed case with no separately tokened desire, policy, verdict, or functionally distinct passion, but with an independently identified constitutive evaluative control role, is admitted by NDT's encoded route and excluded by strict universal separability. This comparison does not classify every broader Humean psychology. **Primary sources:** Hume, n.d.-a

5.11 conceptual context

Source-sensitive judgments and source-sensitive moral credit / Kant's source-sensitive moral-worth contrast

Record: CTX035 **NDT claim:** 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. **Source-backed claim:** Kant's shopkeeper illustrates that outwardly honest conduct can differ in moral worth according to why it is performed. **Exact relation:** conceptual context. **Why:** Both views reject surface conformity as sufficient, but they classify different things. Kantian moral worth neither implements NDT nor rivals its token-type criterion. The comparison is retained only as conceptual context and does not attribute NDT's provenance architecture to Kant. **Negation test:** An outward-conformity account that treats matching conduct as sufficient denies the source-sensitive contrast shared by the two views. **Discriminating consequence:** Holding words and behavior fixed, NDT token type can vary with relationship contribution; after NDT type is held fixed, Kantian moral worth can still vary with motive. **Primary sources:** Kant, 1785

A precise model need not preserve every use of the word moral / Family-resemblance structure in concept use

Record: CTX037 **NDT claim:** 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. **Source-backed claim:** Wittgenstein's family-resemblance model explains how uses can overlap without sharing one feature common to every case. **Exact relation:** conceptual context. **Why:** Family resemblance helps explain semantic variation in the word moral, but supplies neither an NDT computation nor a same-level cognitive classifier. It is conceptual context for a lexical problem, not a component within NDT's architecture. **Negation test:** A strict lexical-preservation requirement predicts that a theory of moral cognition is inadequate whenever its classifications diverge from an established use of moral. **Discriminating consequence:** NDT can acknowledge family-resemblance variation while remaining answerable to independent paradigm cases, contrast cases, and discriminating empirical consequences. **Primary sources:** Wittgenstein, 1953

Direct encounter without prior identity or intimacy / Levinasian face-to-face ethical address

Record: CTX056 **NDT claim:** 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. **Source-backed 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. **Exact relation:** conceptual context. **Why:** Levinas supplies phenomenological and normative context for direct encounter, not an empirical mechanism, an NDT lineage claim, or evidence that encounter is the uniquely required positioning route. The comparison remains outside the operational-rival grid. **Negation test:** Not applicable as a contradiction test: conceptual context would be misclassified if treated as a psychological implementation, empirical validation, or lineage source. **Discriminating consequence:** The comparison keeps open direct encounter without prior identity or intimacy while requiring separate evidence to identify a psychological route, any downstream bender edge, and actual contribution to a judgment. **Primary sources:** Levinas, 1961/1969

5.12 lineage

Many-voiced normative cognition / Bakhtinian polyphony

Record: CTX026 **NDT claim:** NDT represents multiple coherent normative frames without forcing them into one voice and compares them only when their targets can be aligned. **Source-backed claim:** Bakhtinian polyphony describes independent voices and consciousnesses that are not merged into a single viewpoint. **Exact relation:** lineage. **Why:** The term migrates from literary form to cognitive organization. NDT adds formal frame-event identity, comparison readiness, alignment, and typed outcomes; none is attributed to Bakhtin. **Primary sources:** Bakhtin, 1984

6. What this review establishes

Across 48 governed records, this map finds no exact same-level match for NDT's full package. That bounded result does not establish global novelty, proof, measurement validity, or a natural kind.