Moral Coherence and the Boundary of the Moral Community

Animal Killing, Institutional Concealment, and Necrocapitalism

Flyxion
Independent Researcher

November 2026

Abstract

Moral principles are seldom applied to every being they could govern. They are partitioned by categories inherited from custom, law, and economic organization, and the local rules on each side of a partition may differ. This essay asks what must be true of such a partition for the local rules to belong to one coherent principle. It defines accountability, the requirement that a moral judgment depend only on properties the principle itself admits as relevant, and shows that two accountable local judgments, one for humans and one for nonhuman animals, extend to a single accountable judgment if and only if they agree on every relevance profile that both domains realize. It then shows that adding a further relevant property removes a disagreement only for the cases that the property distinguishes, and removes it vacuously when the property merely separates the two domains. The framework is applied to the killing of animals for food. Institutional concealment is analyzed as a loss of information that prevents a community from implementing its own judgments, with no claim about psychological effects. Necrocapitalism is given a two-tier definition that distinguishes revenue contingent on death from accumulation that depends on the deliberate killing of organisms, and the formal results are brought to bear on the second tier. Exceptions to a prohibition are analyzed next, and causal contribution to suffering is separated from the symbolic normalization of animal consumption. The thesis is conditional on an explicit commitment to nonarbitrary moral distinctions and does not depend on any empirical claim about meat consumption: given that commitment, a moral community cannot secure the coherence of its principles by institutionalizing exceptions, and must justify the criteria by which exceptions are made.

Keywords: moral coherence, accountability, gluing condition, moral community, institutional concealment, necrocapitalism

1 Accountability, Compatibility, and the Gluing Condition

Many moral principles, among them the principle that one ought to be kind, are stated without restriction and applied with restriction. Kindness is demanded toward neighbors, strangers, and children, while the treatment of animals follows other rules, under which killing for food is routine in most societies and the killing of companion animals is condemned. A division of this kind is not by itself incoherent, because a principle that assigns different weights to different circumstances can consistently yield different verdicts in different places. The question taken up here is narrower and logically prior to any dispute about which animals matter or how much they matter: what must be true of a partition of a principle's domain of application for the local rules on the two sides of it to belong to one coherent principle at all. The answer proposed is a condition on overlaps, and it is the condition familiar from the gluing of local data into global data.

The analysis is relative to premises and is intended to remain so. It takes a principle as given, together with a specification of the properties that principle admits as morally relevant, and asks whether a family of judgments is consistent with it. Nothing in what follows entails that animals have the same moral standing as humans, that any particular practice is wrong, or that any premise is correct. The formal results are compatible with a view on which animal suffering counts for very little, provided the discount is accountable to a stated property that does the work claimed for it. What the formalism can establish is inconsistency relative to specified premises, and what it can locate is the place where a justification is owed.

Let A be a set of action types, I a set of individuals who may be affected by actions, and C a set of circumstances. A *case* is a triple $x = (a, i, c)$, and $X = A \times I \times C$ is the set of cases. Let $\mathcal{P}$ be a set of properties of cases, including properties of the action (that it kills its patient), of the patient (that it is sentient, that it is vulnerable), and of the circumstances (that the action is necessary to the agent's survival). A function π assigns to each case the set $\pi(x) \subseteq \mathcal{P}$ of properties it has. The *relevance set* $R \subseteq \mathcal{P}$ is the set of properties that the principle under examination admits as morally relevant, and the *relevance profile* of a case is $\rho_R(x) = \pi(x) \cap R$. A *judgment* on a set of cases $X' \subseteq X$ is a function $J : X' \to \{0, 1\}$, where 1 means permissible. The binary codomain is a convenience, and every statement below holds when it is replaced by any ordered set of verdicts, so that graded weights are covered. Local judgments carry subscripts, as in J_H and J_N, and J without a subscript denotes a judgment on an arbitrary set of cases, in particular the global judgment of Proposition 1.3. The set of individuals is partitioned as $I = H \sqcup N$, humans and nonhuman animals, and X_H and X_N denote the cases whose patients lie in H and N respectively. This division is the socially inherited categorization under examination. Real practice is finer, distinguishing companion animals, livestock, pests, wild animals, and laboratory animals, and the construction applies to any finite cover, with the conditions below imposed on each pair of cells.

Definition 1.1 (Accountability)**.** A judgment J on X' is *accountable to* R if for all $x, y \in X'$, $\rho_R(x) = \rho_R(y)$ implies $J(x) = J(y)$.

Accountability is the requirement that like cases, in the respects the principle recognizes, be judged alike, and it is weaker than any requirement of equal treatment. Cases with different profiles may be judged differently, and the difference in weight may be as large as the principle allows. A defender who holds that a unit of human suffering outweighs any amount of animal

suffering satisfies the definition so long as the profile records the property on which the difference depends. What the definition rules out is a difference in verdict between two cases that agree on every property the principle admits as relevant. This keeps the formalism from treating every difference in moral weight as an inconsistency, which would make the analysis trivially hostile to every nonidentical treatment.

Given accountability, each local judgment factors through its profiles. If J_D is accountable on X_D, there is a function $\bar{J}_D$ defined on the set $\rho_R(X_D)$ of profiles realized in D with $J_D = \bar{J}_D \circ \rho_R$. The two domains share no individuals and therefore no cases, so the overlap between them cannot be an overlap of cases. It is an overlap of profiles.

Definition 1.2 (Overlap and compatibility)**.** The *overlap* of the two domains is $O_R = \rho_R(X_H) \cap \rho_R(X_N)$. Accountable judgments J_H and J_N are *compatible* if $\bar{J}_H(p) = \bar{J}_N(p)$ for every $p \in O_R$.

Proposition 1.3 (Gluing)**.** *Let J_H and J_N be judgments on X_H and X_N, each accountable to R. There is a judgment J on $X_H \cup X_N$ that is accountable to R and restricts to J_H on X_H and to J_N on X_N if and only if J_H and J_N are compatible. When such a J exists it is the only judgment on $X_H \cup X_N$ that restricts to the two given ones.*

Proof. Suppose the judgments are compatible. Because $X_H \cap X_N = \emptyset$, the function J that equals J_H on X_H and J_N on X_N is well defined. To show that it is accountable, take x, y with $\rho_R(x) = \rho_R(y)$. If both lie in the same domain, the accountability of the local judgment gives $J(x) = J(y)$. If $x \in X_H$ and $y \in X_N$, the common profile belongs to O_R, and compatibility gives $\bar{J}_H(p) = \bar{J}_N(p)$, so $J(x) = J(y)$. Conversely, suppose an accountable extension J exists, and let $p \in O_R$, so that $p = \rho_R(x) = \rho_R(y)$ for some $x \in X_H$ and $y \in X_N$. Accountability of J gives $J(x) = J(y)$, and since J extends the local judgments this says $\bar{J}_H(p) = \bar{J}_N(p)$. Uniqueness holds because an extension is determined on each cell by its restriction to that cell. □

Remark 1.4. In sheaf-theoretic language, accountability makes each local judgment a section over the set of realized profiles of the sheaf of $\{0,1\}$-valued functions on the profile space, and Proposition 1.3 is the gluing axiom for the cover $\{\rho_R(X_H), \rho_R(X_N)\}$ of $\rho_R(X_H \cup X_N)$ [6]. The mathematics is elementary, and the vocabulary matters only for what it makes explicit: coherence is a condition on overlaps, local rules that are each internally consistent can fail to define a global rule, and the failure is located at specific shared profiles and not in either rule taken alone.

Definition 1.5 (Obstruction)**.** The *obstruction* to gluing is the set $\mathrm{Ob} = \{p \in O_R : \bar{J}_H(p) \neq \bar{J}_N(p)\}$. The two judgments are coherent exactly when Ob is empty.

A nonempty obstruction does not by itself condemn either local rule. It says that some human case and some animal case agree on every property the principle admits as relevant and receive different verdicts, so that either the verdicts must change or the principle must recognize a further property that separates the cases. The second option is the one that defenders of a partition normally take, and the next results describe what it can and cannot accomplish. Let $R' = R \cup Q$ be a refinement of the relevance set by a set Q of additional properties disjoint from R. A local judgment that is accountable to R is accountable to every refinement of R, since equality of R'-profiles implies equality of R-profiles.

Proposition 1.6 (Refinement)**.** *Let J_H and J_N be accountable to R and let $R' = R \cup Q$ with $Q \cap R = \emptyset$. Then J_H and J_N are compatible relative to R' if and only if every pair $(x, y) \in X_H \times X_N$ with $\rho_R(x) = \rho_R(y)$ and $J_H(x) \neq J_N(y)$ satisfies $\pi(x) \cap Q \neq \pi(y) \cap Q$.*

Proof. Note first that $\rho_{R'}(x) = \rho_R(x) \cup (\pi(x) \cap Q)$, so $\rho_{R'}(x) = \rho_{R'}(y)$ holds exactly when $\rho_R(x) = \rho_R(y)$ and $\pi(x) \cap Q = \pi(y) \cap Q$. Suppose the judgments are compatible relative to R' and let (x, y) be a disagreeing pair with equal R-profiles. If $\pi(x) \cap Q = \pi(y) \cap Q$, then the R'-profiles are equal, the common profile lies in $O_{R'}$, and compatibility forces $J_H(x) = J_N(y)$, a contradiction. Conversely, suppose every such disagreeing pair differs on Q, and let $x \in X_H$, $y \in X_N$ have equal R'-profiles. Then their R-profiles are equal, and if $J_H(x) \neq J_N(y)$ the pair would be a disagreeing pair that agrees on Q, contrary to hypothesis. So the judgments agree on every shared R'-profile. □

Proposition 1.6 states the sense in which the burden falls on the defender of a partition. A proposed justifying property q removes the obstruction only for those disagreeing pairs on which it differs, and a set of properties removes the whole obstruction only if every disagreeing pair is distinguished by some member of it. The condition is joint, since Q is a set: no single property need distinguish all the pairs, and what must be shown is that no disagreeing pair is left undistinguished by every member of the set taken together. Which pairs a given property distinguishes is a question of fact and argument that the formalism leaves open. It is also the structure of the argument from marginal cases, which holds that every property proposed to justify the difference in treatment is absent from some humans or present in some animals [10, 9]. Stated as a conditional about pairs, the argument says that if some disagreeing pair agrees on every candidate property, the obstruction persists for that pair. Whether such pairs exist for the candidates offered, such as language, rational agency, or reciprocity, is the substantive dispute, and the formal result neither settles it nor pre-judges it.

There is a second respect in which refinement can fail, and it concerns not the pairs a property fails to distinguish but a property that distinguishes everything. Say that a property q *separates* the domains if $q \in \pi(x)$ for every $x \in X_H$ and $q \notin \pi(y)$ for every $y \in X_N$. The membership property of the patient, that it belongs to H, is one such property.

Proposition 1.7 (Collapse)**.** *If some property $q \in \mathcal{P}$ separates the domains, then for any pair of local judgments accountable to R, the refinement $R' = R \cup \{q\}$ renders them compatible.*

Proof. Every $x \in X_H$ has q in its R'-profile and no $y \in X_N$ does, so no human case shares an R'-profile with an animal case. Hence $O_{R'} = \emptyset$, and compatibility holds vacuously. □

Proposition 1.7 shows that the gluing condition has no content if every refinement is admissible. A rule that treats the partition as itself a morally relevant property cannot be inconsistent with anything, for the same reason that it cannot be tested by anything. Coherence therefore constrains a partition only when the admissible refinements are restricted, and the restriction must come from outside the formalism. Three requirements can be distinguished. A property is *independent* if it is specified without reference to the domain label, so that its application to a case can be determined without first asking which domain the case belongs to. It is *consistently applied* if every case in either domain that has the property is treated as having it. And it is *substantively justified* if there is a moral reason, not derived from the partition, for treating it as relevant. The first two are formal or procedural conditions, and a property can satisfy both without being defensible, since an arbitrary biological characteristic might distinguish two groups perfectly while supplying no adequate ground for different treatment. The third is normative, and it is the point at which the formalism stops. Independence in particular is a weak screen: it excludes

only properties defined by the partition itself, and a property that is merely coextensive with the partition, such as species membership determined biologically, passes it. Such a property is excluded, if at all, by the substantive requirement, so the third requirement is the one that does the non-trivial work, and the first two are best read as procedural screens that leave it intact. Call a refinement *principled* when each of its properties meets all three requirements. A principled property that happens to separate the domains is not excluded by this definition, but it is no longer trivial, since its justification must then be shown to apply to the individual cases and not merely to the label.

Assumption 1.8 (Nonarbitrariness)**.** A community that states a principle as general commits itself to admitting differences in treatment under that principle only through principled refinements.

This commitment is not derived from the gluing theorem. It is an additional normative premise, the demand for nonarbitrary moral distinctions, and the argument of this essay is conditional on it. Given the premise, the formal results show where it bites: the community owes a principled property that distinguishes the disagreeing pairs, and an exception drawn on the partition itself does not discharge the debt. Without the premise the results remain true and lose their normative force. Coherence is necessary for a system that claims principled consistency, but it is not sufficient for moral adequacy, since a community can construct an internally consistent system around distinctions that are arbitrary or cruel, and the gluing condition by itself cannot exclude them.

2 The Origin and Maintenance of the Boundary

If partitions of a principle's domain were located by the principle, the question of their coherence would not arise, since the boundary would sit exactly where the relevance set places it. The boundaries that actually govern the treatment of animals are located by other means. Societies treat species unevenly, protecting some, consuming others, and exterminating a third group as pests, and the pattern of that unevenness varies across cultures and periods in ways that appear to follow familiarity, appearance, and disgust at least as closely as they follow any property that a principle of kindness would name. That claim is offered as a hypothesis whose support is partial, and the argument here does not depend on its strongest form. It depends only on the weaker point that the location of a boundary can be explained without reference to the principle it divides, and that when this is so the boundary has never been tested against the principle.

Experimental work on categorization is consistent with this picture. Participants who read that an animal was one that local people ate appear to have accorded it less moral standing than the same animal when it was not so described, so that the classification of an animal as food seems to alter the moral weight it is perceived to have [4]. The motivational account of the meat paradox proposes that people who value animals and also eat them reduce the resulting conflict through a family of strategies, among them avoidance of information, dissociation of the product from the animal, and denial of the animal's mental capacities [3]. In the vocabulary of Section 1, these strategies can be read as moves that present a separating property, here the status of being food, as though it were a principled one. They do not establish that the property is irrelevant. These findings concern how a community's working relevance set is shaped, and they leave open whether the shaping is principled, since perceived standing and membership in the relevance set

that a principle licenses are distinct. They show how a partition can persist when its justification has not been supplied, and they explain why the absence of a justification is not noticed from inside the practice.

History offers a further consideration, which should be stated cautiously. Nikiforuk's central claim concerns slavery and fossil energy, namely that the abolition of slavery in the nineteenth century had an invisible ally in coal and oil [8], and the historical relation between fossil fuels and abolition is contested, so the account is relevant here as an interpretation and not as an established causal explanation. Extending any such pattern to animal-derived products would be an analogy and is not attributed to that book. The claim that the present argument needs is weaker and safer: technological substitution can change the economic cost of abandoning a practice, and so can alter the conditions under which moral objections to it become politically effective. If that is so, the location of a boundary at a given time reflects what is economically feasible to give up as well as what is morally defensible, and a boundary that moves with feasibility was not set only by the principle.

The conclusion of this section is correspondingly modest. The persistence of a partition without a stated justification does not show that the partition is unjustified, and the observation that its location can be explained by habit, categorization, and cost does not supply the missing justification. It does mean that the gluing condition of Section 1 has not been satisfied by default. A community cannot assume that its divisions between those protected by a principle and those excluded from it have already been made on principled grounds, since the mechanisms that locate such divisions operate whether or not such grounds exist.

3 Concealment and Implementability

A principle can be violated by those who understand it perfectly and can also be unimplementable by those who cannot observe the properties it makes relevant. Many consumer products reach their purchasers with their production histories removed. Slaughter is separated physically from retail, processed products carry names that do not refer to animals, and packaging frequently presents imagery of pastoral rearing that may or may not describe the conditions of production. Each of these practices can be modeled as a loss of information at the point where a judgment is made, and the model allows a precise statement of what the loss does and does not show.

Let σ be a *presentation map* that assigns to each product-related case x the information $\sigma(x)$ available to the person who must decide whether to purchase or consume. Two notions need to be separated. Say that σ involves *relevance-information loss* relative to R if there are cases x and y with $\rho_R(x) \neq \rho_R(y)$ and $\sigma(x) = \sigma(y)$, so that the relevance profile cannot be recovered from the presentation. The definition is relative to R because an omission counts only when it removes a property that the principle admits as relevant. A label that omits the color of the barn loses nothing, and a label that omits the method of killing loses a great deal if the method is in R. Say further that the loss is *decision-obstructing* for a judgment J if there are such cases with $J(x) \neq J(y)$, so that the verdict cannot be recovered. The first concerns the inability to recover the profile and the second the inability to recover the verdict, and only the second prevents a rule from being carried out, since omitted information that would not have changed the verdict is harmless to implementation.

Proposition 3.1 (Implementability)**.** *Let J be accountable to R. A decision procedure that depends only on $\sigma(\cdot)$ can agree with J on every case if and only if the information loss of σ is not decision-obstructing for J.*

Proof. Suppose the loss is not decision-obstructing, and let $d(s) = J(x)$ for any case x with $\sigma(x) = s$. This is well defined: if $\sigma(x) = \sigma(y)$, then either $\rho_R(x) = \rho_R(y)$, in which case accountability gives $J(x) = J(y)$, or the profiles differ, in which case the hypothesis gives $J(x) = J(y)$. The procedure $d \circ \sigma$ then agrees with J. Conversely, if the loss is decision-obstructing, there are cases x, y with $\sigma(x) = \sigma(y)$ and $J(x) \neq J(y)$. A procedure that depends only on σ returns the same verdict on both, and so disagrees with J on at least one of them. □

The procedure in the proof is specified extensionally, and it exists whether or not a finite agent can compute it; a version for finite agents would require σ to be finite or canonically presented. The proposition is almost trivial, and its force lies in what it separates. It establishes that decision-obstructing concealment can prevent a community from carrying out its own judgments at the point of choice, so that a person who accepts the principle and wishes to act on it cannot do so reliably. It does not establish that concealment changes what people believe, how they reason, or how they develop. Whether presentation shapes moral perception in practice is an empirical question to which the categorization experiments cited above bear indirectly and which the present argument does not need to answer. Research on children's moral reasoning about eating animals shows that the question is not idle. Children aged six to ten who became vegetarian against the practice of their families reported reasons concerning animal welfare, the suffering and death involved, whereas none of the children who ate meat gave welfare reasons for avoiding particular meats and cited health and taste [5]. The sample was small and self-selected, and the finding concerns reported reasons, not necessarily underlying moral judgment. That finding indicates that young children can regard the matter as moral and that the information given to them about production is therefore not morally neutral. It does not indicate anything about developmental outcomes for children who object and are overruled, and the literature cited does not address that question. Anecdotal reports of children's distress on learning that a familiar product is made from an animal are common, and they are treated here as illustrations of the information structure, not as evidence for a causal claim.

Concealment is also neither necessary nor sufficient for the incoherence analyzed in Section 1. A fully transparent system could still fail the gluing condition, since people can know every property of every case and apply different verdicts to cases with equal profiles. A concealed system could satisfy the condition if the judgments it produced were in fact compatible. The role of concealment is a protective one. It prevents the cases that would reveal an obstruction from being compared at the point where a person decides, and so it allows a partition to persist without being tested. That is a statement about institutional function, and it is separable from any claim about the mental effects of the practice on those exposed to it.

4 Necrocapitalism: A Two-Tier Definition

The term necrocapitalism has been used for several related purposes. In critical theory it extends the idea of necropolitics, the exercise of sovereign power over who may live and who must die [7], to the organization of accumulation, where it has been applied to forms of enterprise that

depend on the dispossession and death of vulnerable populations [2]. A recent collective work uses it to analyze capitalist responses to the COVID-19 pandemic [1], and a further book uses it for a systemic account of capitalism's connection to ecological and civilizational decline [11], an account the same author's shorter writing extends to descriptions of capitalism as a funeral business that enriches itself through extinction. These usages overlap without coinciding, and the breadth of the term makes it difficult to use as an analytic concept. A distinction is proposed here that narrows it without discarding the cases the broader usages have in view.

Let an enterprise E earn revenue r_E from the goods and services it sells, and let X denote deaths that occur independently of its operation. The definitions are stated in terms of revenue, since profit depends on costs that the classification does not need.

Definition 4.1 (Death-contingent revenue)**.** An enterprise has *death-contingent revenue* if its operation produces no deaths and its revenue is increasing in X.

Definition 4.2 (Death-productive accumulation)**.** An enterprise engages in *death-productive accumulation* when the deliberate killing of organisms is a constitutive operation in the production of goods from which it earns revenue.

The second definition is stated in terms of the production process and not of the product. It does not depend on whether a technological substitute exists, so cultivated meat, which is intended to supply a meat product without slaughter, falls outside it, while the existing process of slaughtering animals for meat falls inside it whatever alternatives may later become available. It also excludes enterprises whose operations cause deaths only as an unwanted consequence of increased output while their revenue comes from something unrelated to those deaths, since the killing must be a constitutive and deliberate step in producing what is sold.

Funeral services are the clearest instance of the first tier. Revenue arises from deaths, and the enterprise neither causes them nor is paid to cause them. Industrial slaughter is the clearest instance of the second, because the killing of the animal is a constitutive step in producing the meat that is sold. Fur obtained by trapping or farming and the traditional products of whaling also belong to it. Life insurance is a boundary case and not an instance of the first tier on the definitions as stated. An insurer receives premiums, pays out on death, and earns its margin from the difference, so that its profit rises when fewer deaths occur than were priced. It prices mortality without profiting from the occurrence of death, and the broader usage that groups it with the funeral industry rests on the idea that it commodifies death, which is a different feature from the one captured by either definition. The distinction does not require the broader usage to be rejected. It makes clear which cases engage the analysis of this essay directly.

The two tiers differ in their relation to the coherence condition of Section 1. A death-contingent enterprise raises moral questions, about the exploitation of grief, about pricing, and about access, that do not concern the killing of any patient, and its operation produces no cases with the property of killing. A death-productive enterprise generates such cases as a matter of its structure, and when the patients are animals it generates cases in X_N with a specific profile. Only the second tier therefore bears on the partition problem, and the argument of the earlier sections applies to it: either the cases it generates are judged compatibly with cases of equal relevance profile in the human domain, or a principled property distinguishes them. Accordingly the results of Sections 1 and 3 bear on the second tier, through the cases it generates and the presentation of its products, and they do not bear on the first. A further feature of death-productive enterprises is

that their persistence does not require anyone's intention. Revenue depends on demand, demand is distributed across consumers, retailers, and restaurants whose local conduct is accountable to local incentives, and the aggregate pattern is accountable to no stated principle. This is the institutional analogue of the gluing failure of Section 1, in which each local rule is internally consistent and the global pattern follows from none of them.

5 Exceptions, Causal Responsibility, and Objections

Exceptions to a prohibition are a normal feature of moral reasoning, and a framework that classified them as incoherent would be wrong. Killing in self-defense is excused, a necessity that endangers life licenses acts otherwise forbidden, and duress reduces responsibility. The question for the present analysis is when an exception preserves coherence. Let $k \in R$ be the property that an action kills its patient, and let $n \in R$ be a property of circumstances such as that the action is necessary to avert the agent's death or to escape grave coercion. Consider the rule e that modifies a local judgment J_D to J_D^e by setting $J_D^e(x) = 1$ whenever $k, n \in \rho_R(x)$ and $J_D^e(x) = J_D(x)$ otherwise.

Proposition 5.1 (Exception invariance)**.** *If J_H and J_N are accountable to R and compatible, then J_H^e and J_N^e, obtained by applying the same rule e to both, are accountable to R and compatible.*

Proof. Whether x falls under the rule depends only on $\rho_R(x)$, so cases with equal profiles either both fall under it, receiving verdict 1, or both do not, keeping the verdicts of the unmodified judgment, which are equal by accountability. For compatibility, let $p \in O_R$. If $k, n \in p$, both modified judgments assign 1 to p. Otherwise each assigns the value of its unmodified counterpart, and these agree by hypothesis. □

The proposition is close to immediate from accountability, and its content is negative: exceptions as such are not the source of incoherence. The substantive case is the exception applied on one side only.

Corollary 5.2 (One-sided exception)**.** *Let J_H and J_N be accountable to R and compatible, and let the rule e be applied to J_N alone. Then J_H and J_N^e are compatible if and only if $\bar{J}_H(p) = 1$ for every $p \in O_R$ with $k, n \in p$.*

Proof. For $p \in O_R$ with $k, n \in p$, the modified judgment assigns 1 to p, so agreement with J_H requires $\bar{J}_H(p) = 1$. For every other $p \in O_R$ the modified judgment equals J_N, which agrees with J_H by hypothesis. □

The corollary says that if necessity to the agent's survival licenses killing an animal, then coherence requires that the same necessity license killing a human patient on every shared profile, or that some principled property in R keep such profiles out of the overlap. Most moral views supply such a property, in the form of the greater moral weight of persons or of a right not to be killed that attaches to them, and the formal point is that the exception is then accountable to that property and not to the partition. The coercion exception, which applies when abstention from an expected practice would cause grave social or physical harm to the agent, has the same structure. It is admissible to the extent that the harm to the agent is a property in R and is weighed on the same terms in every case where it occurs.

Two principles are commonly run together in discussions of animal consumption, and the analysis requires that they be kept apart. Let $\kappa(a)$ denote the expected marginal increase in animal suffering or killing that is causally attributable to an action a, and let $\nu(a)$ denote the extent to which a expresses or reinforces an attitude on which animals are consumable objects. The quantity ν is introduced as a conceptual variable with no specified scale, and nothing here assumes that it can presently be measured. The normalization principle is accordingly not yet representable in the framework of Section 1: ν would have to be a property in R, applied to cases in a contestedly determinable way, before it could enter the gluing analysis, and the principle is included here only to mark where it diverges from the causal principle. The *causal principle* makes the permissibility of consuming an animal product depend on κ, and the *normalization principle* makes it depend on ν. The two can diverge sharply. Meat that has been discarded and would otherwise be destroyed has κ close to zero, provided that taking it generates no additional demand and sends no signal to producers. The proviso matters, since gifts, for example, can affect future purchasing decisions even when no money changes hands, so that gifted meat meets the condition only when no such effect is expected. The proviso is also an idealization, since any act of consumption is observable and may, in aggregate, send signals, so that under realistic conditions κ need not be close to zero even for salvage. Its ν is contested, since consuming it may be read as accepting the status of the animal as consumable. Under the causal principle such consumption is permissible on the stated conditions, and under the normalization principle its status is unsettled. The essay adopts neither principle. It requires that an argument for a verdict on salvage consumption say which principle it relies on, and that the principle be one the community accepts for the other cases it judges. Each principle also has a difficulty to face. The causal principle meets the objection that the marginal contribution of a single consumer is negligible or probabilistic, which accounts based on expected effects try to answer. The normalization principle faces the difficulty of specifying what makes an act an expression of an attitude, and of showing that the attitude is itself in the relevance set and not merely a restatement of the partition.

Several objections remain. It may be said that the moral weight of a being scales with cognitive complexity and that this justifies the difference in treatment. If cognitive complexity belongs to R, the framework accommodates the claim, and the question becomes whether cases with equal complexity are treated alike, including those humans whose capacities fall below those of some animals. The answer cannot be derived from the formalism, and the argument from marginal cases cited above is the usual place where it is pursued. It may be said that membership in a community of reciprocal obligation is relevant. In that case the property belongs to R, and the framework asks whether it is a principled property whose rationale applies independently of the partition, or whether it functions as a separating property in the sense of Proposition 1.7. It may be objected that the formalism proves nothing about which premises ought to be adopted. That is correct, and the claim is only that relative to the premises a community states, certain divisions among its judgments are consistent with them and others are not. Finally, no claim has been made or relied upon concerning the effects of these arrangements on the mental health or development of children, and the question appears not to have been systematically studied.

The problem that motivates the analysis is pedagogical. A child taught kindness is not receiving isolated behavioral instructions but is being introduced to a system of reasons that explains why some acts are permitted and others prohibited. When adults teach kindness

as a general principle while exempting familiar practices from its application, they incur an explanatory obligation. The child may reasonably ask why suffering counts in one circumstance and not in another, and the answer owed is a principled property of the kind described in Section 1, not an appeal to the familiarity of the practice. This claim carries no empirical commitment. It does not assert that the child will be confused or harmed, only that an educational practice that presents its principles as general is answerable for the exceptions it teaches alongside them.

The thesis that remains is conditional and general. A moral community that states a principle of kindness, and that divides the domain of the principle by categories whose location can be explained without reference to it, owes a justification for the division in the form of a principled property that distinguishes the cases on which its local rules disagree. It cannot meet this obligation by institutionalizing exceptions, by concealing the production histories that would reveal the disagreement, or by organizing profit around the practices the exceptions cover. Coherence is secured when the criteria for exceptions are themselves justified, and until that is done the boundary of the moral community remains a fact about practice that the principle has not yet been shown to support.

References

[1] M. I. Asma, a collective of authors (2021). *On Necrocapitalism: A Plague Journal*. Montreal: Kersplebedeb.

[2] Banerjee, S. B. (2008). Necrocapitalism. *Organization Studies*, 29(12), 1541–1563.

[3] Bastian, B., and Loughnan, S. (2017). Resolving the Meat-Paradox: A Motivational Account of Morally Troublesome Behavior and Its Maintenance. *Personality and Social Psychology Review*, 21(3), 278–299.

[4] Bratanova, B., Loughnan, S., and Bastian, B. (2011). The Effect of Categorization as Food on the Perceived Moral Standing of Animals. *Appetite*, 57(1), 193–196.

[5] Hussar, K. M., and Harris, P. L. (2010). Children Who Choose Not to Eat Meat: A Study of Early Moral Decision-Making. *Social Development*, 19(3), 627–641.

[6] Mac Lane, S., and Moerdijk, I. (1992). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. New York: Springer.

[7] Mbembe, A. (2003). Necropolitics. *Public Culture*, 15(1), 11–40.

[8] Nikiforuk, A. (2012). *The Energy of Slaves: Oil and the New Servitude*. Vancouver: Greystone Books.

[9] Regan, T. (1983). *The Case for Animal Rights*. Berkeley: University of California Press.

[10] Singer, P. (1975). *Animal Liberation*. New York: New York Review of Books.

[11] Tsakraklides, G. (2023). *In the Grip of Necrocapitalism: The Making and Breaking of a Psychonomy*. Kindle edition, Amazon (ASIN B0CJXDSKZX); no publisher is stated in the listing.