

Split-Domain Cognition in Animal Rights and Veganism

Canonical variant. Form of follow: alignment. Version 1.0, April 2026.

Preface

About this document. This is a domain-specific variant of Split-Domain Cognition, addressed to animal ethics and the welfare-rights collapse. It works out what the SDC failure mode looks like inside a political domain — where descriptive welfare language is reliably made to do categorical rights work — and what a non-collapsed version of the same conversation would have to hold apart. Its philosophical underpinning is the *worth-is-not-hierarchical* position paper (Prayas Abhinav, February 2026), extended outward to non-human animals. Koher's *animalRightsLens* is the first tool partner that follows this variant.

Where this sits in the corpus. The canon home is splitdomaincognition.org. The variants index is at [/variants/](#); this variant in particular at [/variants/animal-rights-and-veganism/v1/](#). The principle this variant interprets is articulated long-form in *Split-Domain Cognition* and short-form in *A Principle, Not a Pattern*. The protocol by which this variant was derived is at [/derivation-protocol-v1/](#).

Authority and version. Canonical, v1.0. April 2026. The website is the source of record. If this PDF and the website disagree, follow the website.

Use. Openly citable. See [Governance](#) for the one-person canon and how variants are admitted.

Author. Prayas Abhinav.

Opening

Animal ethics, as practised in contemporary public discourse, is a domain in which a very specific cognitive collapse has been institutionalised. The collapse is not accidental and not confined to a few bad actors. It is the form in which most everyday conversation about animals and food is conducted. The collapse is between two kinds of thinking that look similar enough to be fused: the language of *welfare* — which describes, measures, compares, and contextualises the treatment of animals — and the language of *rights* — which issues categorical verdicts about whether certain treatments are acceptable at all. These two kinds of thinking are not two degrees of the same thing. They are two different things, with different grammars, different yardsticks, and different criteria for being right.

When the two are fused, a particular rhetorical object appears: the welfare-shaped phrase that does rights-shaped work. "Humane meat" is the archetypal example. "Ethical leather", "grass-fed beef", "free-range eggs", "compassionate farming" — each is a welfare modifier attached to a practice that a rights framework would refuse categorically. The phrase does two jobs at once. It describes a practice (welfare work). It implies that the practice is acceptable (rights work). The implication is smuggled, because the fusion hides it. The consumer encounters a descriptive adjective and receives an ethical verdict, without ever having to examine the verdict against any explicit rights framework. This is Split-Domain Cognition's failure mode appearing in a political domain: an evaluative claim dressed as a descriptive one, produced by a single channel that should have been two.

Veganism, as a lived practice rather than a dietary preference, is the refusal to accept that fusion. It says: the welfare question and the rights question are separable; they must be separated; and when separated, the rights question produces a verdict the welfare question was never going to produce. SDC in this domain is not a technical intervention. It is the architecture of an ethical position that has been obscured by a confused vocabulary.

The language work in this domain

The language work in this domain is the literal description of what is done to sentient beings and under what conditions. What is killed, what is confined, what is bred into existence to be used, what is separated from its offspring, what is cut, what is milked, what is slaughtered, what is skinned, what is displayed. Plain nouns and plain verbs. The description names the act. It does not dress the act in modifiers that carry silent judgement in favour of its continuation.

This is a narrower understanding of the language layer than most public discourse about animals offers. Most public discourse substitutes welfare vocabulary for description — "free-range", "humane", "grass-fed", "ethically raised", "responsibly sourced", "small-scale", "compassionate" — and calls the substitution description. The substitution is not language work in the SDC sense. It is judgement work wearing descriptive clothing. A phrase that pre-loads the acceptability of the practice it names is not describing the practice; it is already issuing a verdict in favour of the practice, delivered at the level where the reader does not yet know a verdict is being issued. Welfare vocabulary is this kind of verdict-in-disguise. It is the reason the collapse catalogued in this variant is so hard to see.

The language layer the rights architecture requires is therefore different from what the surrounding discourse treats as "neutral reporting". It consists of investigative footage that shows what happens in the slaughterhouse, the dairy farm, the chick-culling line, and the live-export ship without narration; first-person accounts from workers inside industries who report what they did with their own hands; supply-chain documentation that traces which animal became which product; species-specific ethology describing what the animal does when it is not being used; and literal language about what each practice involves — breeding for use, confinement, forced impregnation, separation, killing, processing, eating, wearing. Good language work in this domain is that which refuses the welfare-to-verdict shortcut and lets the literal facts be seen before any verdict is issued.

Welfare science, properly located, is not a neighbour operating at the language layer of a shared architecture. It is a parallel framework with its own judgement layer folded into its apparent descriptions. Welfare science asks "how well is this animal being treated within a practice whose continuation is presupposed?" The presupposition of continuation is judgement-work done silently, before the first measurement is taken. The measurements that follow are not neutral description — they are evaluation under a framework that has already decided the use is acceptable and is now optimising within that decision. Welfare science and rights work are not two layers of a single architecture. They are two incompatible architectures. The rights architecture refuses the welfare architecture's prior assumption; the welfare architecture makes the rights architecture's question invisible. This variant sits inside the rights architecture. From inside it, welfare research is not a collaborator at a different layer — it is a framework that has absorbed the judgement work into its descriptions and carries the absorption forward under a name ("welfare") that signals neutrality while doing none.

Language work alone is insufficient. This is the central point the rest of the variant rests on. Accurate literal description of what is done to animals does not, by itself, produce the rights verdict. A reader can look at footage of a slaughterhouse, read a first-person account of a dairy separation, and trace a supply chain from breeding facility to plate, and still shrug. The language layer makes the facts available; it does not close the question of whether the facts are acceptable. Closure requires the judgement layer, and the judgement layer requires substantive content — a framework that says what is and is not permitted, grounded in a reason. Without that framework, language work produces either numbing ("this is what people eat") or sentimentality ("this is sad"), neither of which has the categorical force the rights position requires. SDC in this domain therefore does not end at the separation. It requires the judgement layer to do its work, and the judgement layer must issue its verdict explicitly or the whole architecture collapses back into welfare-flavoured description with slightly better sourcing.

The judgement work in this domain

The judgement work in this domain is the rights verdict. This variant does not hide the verdict's substance behind the variant's architecture. The structural claim (that welfare and rights are two incompatible frameworks, not two layers of one) is true but it is not enough. The rights framework has substantive content, and the content must be stated, or the variant does the same thing the welfare vocabulary does — issues a verdict by implication while refusing to say what the verdict is.

The grounding. The rights framework in this variant is grounded in the argument worked out in

`/home/prayas/Dropbox/personal_projects/gabor.mate/worth-is-not-hierarchical.md` (Prayas Abhinav, 23 February 2026 / 11 April 2026).

That argument shows, without reliance on any prior moral commitment, that worth does not admit of hierarchical ordering across domains of capacity. The attempt to rank one being's worth above another's fails at the level of the common unit: there is no feature of the domains of capacity that would settle the weighting required to combine them. What survives the collapse of comparative measurement is *intrinsic worth* — the property a being has considered in itself, outside any ranking. Intrinsic worth is not a sentimental posit. It is the residual concept when the arithmetic of hierarchical worth is followed through to its end.

The substantive claim. A sentient being possesses worth in the non-comparative sense. Using a sentient being as a means to human ends — as food, as a product, as entertainment, as a test subject, as an aesthetic object, as a status signal, as an object of affection that is still owned — treats the being as if its worth were instrumental, realisable only through serving another being's purposes. That is precisely the kind of treatment the non-hierarchical account of worth refuses. A being with worth is not a resource. Practices that treat a being as a resource are categorically refused by the framework. The refusal does not depend on whether the being was treated humanely inside the practice; the practice itself is the violation.

What the framework refuses. The refused practices include, but are not limited to:

- Killing animals for food. This refuses meat, fish, seafood, and any slaughter-dependent product regardless of species, method, or scale.
- Confining or breeding animals to extract products from their bodies. This refuses dairy, eggs, honey, silk, wool (as currently produced), and every supply chain that depends on ongoing use of living animals.
- Taking skin, fur, feathers, horn, bone, or other body parts, whether from the living animal or from animals killed for this purpose or any other.
- Testing products, chemicals, or procedures on animals.
- Breeding animals into existence for human convenience, human entertainment, or human aesthetic preference — including the breeding of companion species to pedigree specifications that compromise the animal's own welfare, and the breeding of animals for zoos, circuses, aquariums, and racing.
- Using animals in entertainment, display, or ritual — rodeo, bullfighting, horse racing, dog fighting, elephant rides, temple elephants, circuses, performing-animal acts, roadside zoos.
- Habitat destruction conducted to clear land for animal agriculture, which is the same framework violation at a different layer of the supply chain.

Each of these practices treats a sentient being with intrinsic worth as a means. Each is refused categorically. The refusal is not probabilistic, not contextual, not a matter of degree. It is binary.

What the framework does not refuse. The framework does not refuse companionship with rescued animals who would not otherwise survive. It does not refuse conservation work oriented to allowing wild populations to live as they would without human

interference. It does not refuse scientific observation conducted without confinement, manipulation, or harm. The framework refuses *instrumental use*. It does not refuse all contact between humans and animals; it refuses contact that treats the animal as a means.

Why welfare modifiers do not reach this framework. "Humanely raised" does not cancel "killed" because "killed" is the practice the framework refuses. "Free-range" does not cancel "confined-to-extract-eggs" because the extraction is the refused practice. "Small-scale" does not cancel "slaughter" because the slaughter is the refused practice. "Ethical leather" does not cancel "skin taken from an animal" because the taking is what is refused. A welfare modifier can only register a difference in severity inside a practice. It cannot touch the categorical question of whether the practice should exist.

This variant is not architecturally neutral about the refusal. The framework refuses specific practices for specific reasons, each traceable to the single underlying claim that a being with intrinsic worth is not a resource. It refuses **meat** (including fish and seafood) because meat requires killing a sentient being; killing a being with intrinsic worth for another being's taste, nutrition, or convenience uses the killed being as a means. It refuses **dairy** because dairy requires ongoing forced impregnation of the cow, separation of the calf from the mother, killing of male calves who cannot produce milk, and slaughter of the cow at the end of her productive life — the cow's existence has been organised from birth around extraction for human consumption, which is use-as-means at every stage. It refuses **eggs** because egg production requires breeding hens for laying, killing of male chicks at birth (they produce no eggs), and slaughter of the hen when laying declines. It refuses **leather, fur, and skins** because taking skin, coat, or body parts from a being's body treats the body as a human resource and requires killing or mutilating the being to acquire them. It refuses **wool** because wool production requires ongoing confinement, forced breeding, and eventual slaughter of sheep. It refuses **silk** because silk requires killing silkworms inside their cocoons (typically by boiling) to preserve the thread intact — aesthetic preference for an intact fibre purchased at the cost of a sentient being's life. It refuses **honey** because bees are treated as production units whose colony structure, queen's reproduction, and food stores are appropriated for human consumption, with the bees' purposes subordinated to the human's. It refuses **animal testing** because testing subjects sentient beings, who cannot consent, to procedures for another being's experimental or commercial interest.

In each case the practice treats the animal as a means to a human end. In each case the worth argument — the claim that worth does not admit hierarchical ordering and that a being has worth considered in itself, not derived from usefulness to another — refuses the

treatment. The refusal is not a preference the framework could be argued out of. It is the substantive content of what the framework is.

It follows that the industries that depend on these practices are not acceptable, that their products are not made acceptable by better welfare, better sourcing, smaller scale, more compassionate intent, or cleaner marketing, and that the practice of veganism — the consistent refusal to consume or use products of animal use — is the ethical consequence of taking the worth argument seriously across species lines. Veganism, in this reading, is not a preference and not a dietary style. It is the substantive content of the judgement layer being applied in practice.

The variant could have been written without this section. It would then have had the architectural shape of a rights-layer argument without the content of one. The reader would have been shown that welfare and rights are different kinds of thing, and left to supply the rights content themselves. That is the move this variant refuses. A rights architecture without a rights verdict is welfare vocabulary with tidier punctuation. The verdict has to be issued, and this section issues it.

The collapse, with examples

The collapse in this domain is dense, widespread, and difficult to see because it is woven into the marketing vocabulary that surrounds most animal products. A partial inventory:

"Humane meat." A welfare adjective ("humane") attached to a practice-name ("meat") that rights thinking categorically rejects. The phrase uses welfare language to close the rights question. The consumer reads the phrase and understands that the ethical question has been addressed, when in fact the ethical question has been avoided by replacement.

"Ethical leather." Same structure. Leather is skin. The rights question is about taking skin at all. "Ethical" is welfare-flavoured vocabulary about how the skin was taken. The adjective does not reach the question it appears to answer.

"Grass-fed beef", "free-range eggs", "organic dairy." Each is a welfare descriptor ("grass-fed", "free-range", "organic") that signals improved conditions within the practice. None addresses whether the practice is acceptable. The consumer is given the language of better treatment and assumes the rights question has been handled.

"Compassionate farming." A welfare orientation ("compassion") presented as a complete ethical stance. The compassion is genuine at the welfare level — the farmer may well be doing everything welfare science recommends. The compassion does not cross the rights boundary. It is a description of the farmer's relationship to the animal, not a verdict on the practice.

"Small-scale", "local", "artisanal." Scale and provenance descriptors that carry the aesthetics of ethical production without making an ethical claim. A local small-scale artisanal slaughter is still a slaughter. The descriptors describe the conditions around the act, not the act.

"Responsibly sourced." The verb "sourced" has been selected so that the act at the end of the chain disappears. The adjective "responsibly" is welfare-level. The phrase does rights-work while using only descriptive-work vocabulary.

"Animal welfare organisations that do not question use." Institutions that accept the use of animals as a given and work to improve conditions within use. This is a legitimate activity at the welfare layer and is a category confusion when presented as a complete ethical position. Welfare reform does not replace rights-level verdicts; it operates inside a rights framework that has already decided the question of use.

Each of these fusions works the same way. A welfare-level modifier is attached to a practice. The modifier appears to do ethical work. The rights question is quietly not asked. The consumer or reader absorbs the ethical reassurance without having been given any rights-level argument to examine.

The collapse is architectural. It is not a marketing error to be fixed by better copy. It is the fusion of two domains of thinking into a single rhetorical object, and it produces its effect precisely because the fusion hides the two operations. A reader who could see that "humane meat" is welfare language doing rights work would not be persuaded by the phrase; the persuasion depends on the fusion being invisible.

The cost of the collapse

Several costs follow from accepting the collapse as the domain's normal vocabulary.

The rights question disappears from public consciousness. When welfare language occupies the space that rights language should occupy, the rights question is not

debated, not refuted, not answered — it is simply not asked. A population can discuss animal ethics for decades without encountering a single rights-level verdict, because every conversation is deflected into welfare-level modification.

Consumers cannot calibrate. A vegan choosing between products labelled "humane", "ethical", "compassionate", and "responsibly sourced" has no way to compare them, because none of the labels answers the question the vegan is asking. The labels answer a different question (how well was the animal treated?) and the consumer is asking a different one (should this animal have been used at all?).

Welfare research is treated as though it were neutral reporting when it is not. Welfare research operates inside a framework that has already accepted the use of animals and is measuring within that acceptance. When welfare findings are cited in public as though they were neutral descriptions of conditions, the prior acceptance is carried forward silently. The rights question — whether the use should exist at all — is not refuted; it is made invisible by the citation of numbers that presuppose a "yes" answer to it. The cost is epistemic: the public loses the ability to recognise that the most frequently cited "evidence" in animal debates is evidence from inside a framework the rights position rejects.

The movement drifts toward welfare-reform strategy because welfare is politically palatable. Welfare-reform work can be funded, legislated, and marketed without disturbing the underlying use of animals. Rights-level work cannot. Under pressure — funding pressure, reputational pressure, coalitional pressure — organisations that began with a rights commitment often migrate toward welfare-reform language and welfare-reform wins, and describe the migration as pragmatism. The drift is predictable and it is the collapse doing its work at the institutional level. Without an explicit architecture separating the two frameworks, the harder framework loses ground to the easier one over time, and the movement forgets what it had originally been refusing.

The person who lives the rights position cannot find a vocabulary to defend it in public. Every phrase available has been colonised by welfare-flavoured modification. A vegan explaining why they do not eat "humanely raised" dairy is accused of not recognising the welfare improvement, as though the improvement answered the question. It does not; but the language has no way to say so without first doing the architectural work of separating the two layers. The variant you are reading is that architectural work, done for one domain.

Intervention in the industry is confused. Activists, legislators, NGOs, and funders operating under welfare vocabulary produce welfare outcomes and call them rights

outcomes, or produce rights outcomes and describe them as welfare reforms. Over time the activity of the movement drifts toward whichever vocabulary is more politically palatable, which is welfare. The rights-level work becomes harder to fund, harder to explain, and harder to sustain, because the domain's vocabulary no longer holds a clear place for it.

The separation, in this domain's language

Applying SDC in this domain means writing the three layers explicitly, in the domain's terms, rather than letting one collapsed vocabulary carry all of them. The separation is not a compromise between welfare and rights. It is the internal architecture of the rights position itself.

Layer 1 — literal description (language work). What is happening to this sentient being? What is the practice, what does it consist of, what does the animal undergo, what does the human do? The inputs to this layer are literal: investigative footage viewed unnarrated, first-person accounts from workers inside the industries, supply-chain documentation tracing which animal becomes which product, species-specific ethology describing what the animal would do if not used, plain-language naming of each step in the practice. Welfare-flavoured modifiers are excluded from this layer because they carry judgement-in-disguise and contaminate the description. A good Layer 1 tells the reader what is there, in the literal sense, without pre-closing the question of whether what is there is acceptable.

Layer 2 — rights framework with substantive content (judgement work). The categorical framework against which the described practice is assessed. In this variant the framework is the one specified in the judgement-work section above: sentient beings possess non-hierarchical intrinsic worth, instrumental use of a being with worth is categorically refused, and the list of practices that count as instrumental use is specific and public. The framework is prior to the case. It can be revised only as a public act, by returning to the grounding in `worth-is-not-hierarchical.md` and the substantive claim. Applied to a Layer 1 description, the framework issues a binary verdict — use-as-means fails the framework; the practice is refused. The binary is not a rhetorical flourish. It is what a rights framework is for; a rights framework that issued probabilistic verdicts would have collapsed back into welfare.

Layer 3 — narration of the refusal (explanation in readable prose). A statement of the verdict that references the description (Layer 1) and the framework (Layer 2) and shows the reader how the verdict follows. Layer 3 is where vegan discourse usually lands in public, and usually does all three layers' work simultaneously, which is why it is often experienced as moralistic or confusing. A good Layer 3 narration reports a verdict that has already been issued at Layer 2 against a description produced at Layer 1 — the verdict is not being improvised in the moment of narration. Layer 3 does not smuggle evaluation into description, does not hide the framework, and does not pretend the refusal is a matter of taste.

Layer 1 alone is insufficient. Layer 2 alone — a framework without a described case — issues no verdict. The three layers must be held together, in order, or the architecture does not work. A practice described honestly at Layer 1 is still only described; until Layer 2 is applied, no refusal has been issued. A framework stated at Layer 2 without Layer 1 is a slogan. A narration at Layer 3 that has no Layer 2 behind it is the vegan's "moralising" that the surrounding culture complains about — a verdict delivered without the architecture that would let the verdict be re-entered.

This is the architecture *animalRightsLens* implements in software, and it is the architecture that makes the lived position of veganism defensible in public without requiring the vegan to perform moral outrage. The verdict lives in the framework. The vegan is reporting that the framework has been applied to a literal description of the practice in question.

The philosophical grounding

The separation above rests on a deeper structural claim about whether worth admits of hierarchical ordering at all. That claim is worked out in

`/home/prayas/Dropbox/personal_projects/gabor.mate/worth-is-not-hierarchical.md` (Prayas Abhinav, 11 April 2026), and it grounds the rights framework this variant operates under.

The argument in that document proceeds in several steps. It shows that a hierarchy of worth requires a common unit in which different domains of excellence can be compared. It shows that the domains of human excellence (intellectual, moral, aesthetic, relational, craft-level) do not in fact share such a unit. It shows that the weighting required to combine them into a single ranking is arbitrary in the strict sense — there is no feature of

the domains themselves that would settle the weighting. It shows that the retreat to domain-local rankings concedes the original question: there is no hierarchy of worth as such, only hierarchies of capacity in particular domains, none of which aggregates into a hierarchy of worth. And it shows that the qualification required to perform a global ranking is self-defeating: the cross-domain structural faculty that alone could warrant the judge's seat discovers, in its exercise, that the seat is unoccupiable.

What follows is the concept of *intrinsic worth*, understood not as a sentimental posit but as what remains after comparative measurement has dissolved. Intrinsic worth is the property a thing has considered in itself, outside any ranking. It is what survives when the attempt at comparison is followed through to its end.

This argument does not, by itself, tell us which entities have intrinsic worth. But it removes the restriction that worth is a coherent concept only where human-metric excellence is demonstrable. A cow does not accomplish anything by human metrics. Neither does a chicken. Neither does a child before they have achieved anything. These are not deficiencies measured against the right scale; they are differences that the argument shows to be incommensurable with the human scale in exactly the way different domains of human excellence are incommensurable with each other. The same structural observation that prevents the argument's author from feeling insufficient before Einstein extends the range of moral concern to beings whose excellences are not legible in human terms. The ethics follows from the logic.

This is why veganism, as practised in Prayas's life, is not a dietary preference and not a matter of taste. It is the ethical application of a structural argument about worth. The rights framework at Layer 2 of the SDC variant does not need to assert that animals are *equal* to humans on some ranking — that assertion would fail for the same reason hierarchical rankings of humans fail. It needs to assert only that worth does not admit of hierarchical ordering across domains, and that animals' worth is real in exactly the sense in which human worth is real: as a non-comparative property of the being considered in itself. From that follows the rights-level verdict that the practice of using animals for food, clothing, or entertainment fails the framework, because the framework does not recognise the ranking that would make the practice permissible.

Two consequences for the SDC variant.

First, the argument in `worth-is-not-hierarchical.md` is the deepest explanation of why the welfare / rights distinction matters structurally, not only as a matter of vocabulary hygiene. The distinction matters because the conflation of the two is a special case of the conflation of incommensurable domains — the very kind of

conflation that the worth argument shows to be incoherent at the global level. SDC in this domain is not a fresh invention; it is an application of a more general argument that Prayas had already made about persons, now extended to species.

Second, the argument is itself an instance of SDC. The worth document separates the descriptive work (identifying domains of excellence and their yardsticks) from the judgement work (the claim that a global ranking across those domains is structurally impossible) from the narration (the closing pages about the residual concept of intrinsic worth). The three moves are held apart. The document could have fused them and delivered a single moralistic assertion; it does not. This is SDC in prose, and its prose-level faithfulness to the architecture is what makes its argument survivable. A moralistic fusion would have collapsed under its own rhetoric. The separation is what lets the argument stand.

For these reasons, the worth document is not an optional appendix to this variant. It is the variant's load-bearing philosophical ground. Anyone building on the animal-rights variant — in a tool, a grant, a position statement, an artwork — should read the worth document first. The variant is the application; the worth document is the foundation.

Relationship to Koher output

The animal-rights variant has an unusually concrete landing in the Koher practice, compared to some other domains where the variant is mostly intellectual.

animalRightsLens. The tool operationalises the three-layer architecture for this domain. Stage 1 produces a structured description of a menu, dish, product, or practice. Stage 2 applies a rights-based rule framework to the description and issues a verdict. Stage 3 narrates the verdict. The tool is the software expression of this variant. Its partnership is with Open Paws (contact: Nicole), and its grant posture is via SFF Thematic Round (Animal Welfare) with Open Paws as fiscal sponsor.

Open Paws partnership. Nicole at Open Paws is Koher's primary collaborator in this domain. The partnership accepts the rights-layer framing and uses the tool for purposes that are rights-level work, even while operating in the broader ecosystem of welfare-reform organisations. The partnership is compatible with this variant precisely because Open Paws does not collapse welfare and rights — they use welfare tactics under rights-level commitments, which is what clean SDC architecture permits.

SFF Thematic Round submission. The submission (via Open Paws as fiscal sponsor) proposes extending the tool's coverage of rights-level verdicts. The grant narrative draws on this variant for its theoretical framing, without requiring reviewers to read the variant in full.

AI x Animals RFP submission (8 April 2026). Submitted via Granti. \$25,000 over six months. The rights-framework commitment of the tool is central to the proposal.

The worth document. Written by Prayas on 23 February 2026 in personal-relational form and rewritten on 11 April 2026 as a philosophical position. It is the deepest published reason the variant exists.

Manjushree Abhinav's role. Prayas's wife is a vegan activist and has served as an annotator on animalRightsLens. The lived practice of veganism in the household is a primary source of the variant's specificity. The variant is not written from outside the position; it is written from inside a shared daily practice of rights-level living.

Museum of Vestigial Desire and the earlier writing. The archive is not a primary source for this variant's structural claim — that claim is grounded in the 2026 worth argument — but it is the older layer of Prayas's work in which the *ethical consequence* of the structural claim was already present, before the structural form of the argument had been written. The clearest evidence is `desires/civilisation/animal` (29 December 2012, 914 words):

"The animal is the one pure state of consciousness that we can observe ... We will accept that the progression of evolution was reversed. That humanity is really the deep, far end and that the animals are ahead of us in the game. To accept that will be painful. The thought violates religion, folklore and everything that led us to have the privilege of wilfully stooping lower than ever thought possible, the privilege of developing a culture, of claiming autonomy, of being lost but having the insolence of driving on at full blaze."

The passage inverts the human-supremacy ranking directly — not as a rhetorical flourish, as the conclusion of an argument the essay is willing to follow — and names the inversion as painful and as violating inherited frames (religion, folklore, the privilege of culture). It does not yet have the 2026 worth argument's structural claim that hierarchical worth across domains is incoherent at the level of the common unit; what it has is the *ethical consequence* the 2026 argument will later license. Fourteen years of practice sits between the two, and the consistency is not rhetorical — it is the same person working the same position through different instruments.

The archive also has a quieter cross-link worth registering: MoVD tags `desires/civilisation/animal` with `distinction` and tags `analysis/who-is-a-user` (2021) with `distinction`. The archive's own tagging places the animal and the user in the same conceptual neighbourhood — both are categories whose unexamined acceptance the archive treats as a failure of attention. This variant's architectural move (welfare and rights are not two layers of one architecture; they are two incompatible architectures) is the SDC-era articulation of the distinction the archive was already drawing.

Several other MoVD pages touch species and instrumental-use concerns in prefigurative form (see `../movd-as-precursor.md`, passage 9, for the annotation), but `desires/civilisation/animal` is the load-bearing case and the one this variant would quote in a published position paper, a grant application, or a talk that wanted to show the variant is not a 2026 invention. It is the variant's older ground, and it is in Prayas's own voice.

What this variant makes possible

Three things become possible when the variant is written down.

A **vocabulary becomes available** for distinguishing what vegans have been trying to say from what welfare-reform language says. The distinction is no longer a matter of emphasis or passion; it is a matter of architecture. A vegan can point to Layer 2 of this variant and say: this is the question I am asking. The question is not identical to the welfare question, and the framework under which it is answered is explicit.

A **tool-and-grant strategy becomes legible** to funders who might otherwise read animalRightsLens as welfare-reform software. The tool is rights-framework software. Funders operating under rights-level commitments can see this clearly; funders operating under welfare-reform commitments can see why the tool is complementary to their work rather than contradictory.

A **pedagogy becomes possible**. Practitioners, students, and activists entering the domain can read this variant and see what the domain's collapse is, where the rights question lives, and why the welfare modifier is not a substitute for it. They can teach the distinction rather than transmit it tacitly through modelling. The learning curve for new

entrants to the rights-level position is shortened because the architecture has been made legible.

The variant does not argue that welfare-reform work is wrong. It argues that welfare-reform work operates at a different layer than rights work, that the collapse between them harms both, and that SDC's architectural refusal of the collapse gives both layers room to do their own work. The refusal is what the variant protects, and what the tool, the partnership, the grants, and the lived practice together make real.

Version 1.0 — 16 April 2026. Read `../sdc.md` and `README.md` before working on this variant. Pair-ref:

`/home/prayas/Dropbox/personal_projects/gabor.mate/worth-is-not-hierarchical.md` .