

Labels Are Not Arguments

Credence, Identity, and the Misclassification of Rational Agents

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

A broad identity label can be useful shorthand, but it cannot by itself warrant a verdict on a person's rationality or an argument's quality. Unlike a claim-specific profile—the proposition, confidence, reasons, and revision conditions—one label can fit many opposing epistemic states; dismissal therefore needs an additional, reliable relevance bridge. Explicit self-definitions and dependable group patterns may remain evidence, but only when connected to the case and kept open to correction.

Preview:

Label offered	Missing bridge	Individuating test	Rational use
"Atheist," "Christian," "liberal," or another identity term.	Show why that category predicts the premise, credibility, motive, or error attributed here.	Ask for the exact claim, confidence, reasons, and what would change them.	Use the label as orientation or defeasible background evidence, not as a verdict.

The target is label-alone adjudication, not every provisional use of categories or group information.

Visual preview:

Labels vs. Arguments: Moving from Shorthand to Substance

Labels are low-resolution summaries that often function as argument-substitutes to dismiss opinions without engagement. This infographic contrasts the binary switch of labeling with the gradient dial of rational credence, advocating for a shift toward inquiry-based dialogue.

The Trap of Labels (Low-Resolution Shortcuts)



Labels are not Arguments

They compress complex beliefs into low-resolution social tags, often serving as permission to stop listening.



The "Identity-Protective" Shield

Labels turn movable beliefs into rigid identities, making counter-evidence feel like a personal attack on belonging.



Testimonial Injustice

Credibility is unfairly deflated based on a person's category rather than the merit of their claim.



BINARY SWITCH

The Credence Profile (High-Resolution Reality)



Belief is a Dial, Not a Switch

Because evidence arrives in degrees, rational belief should be represented as a gradient (0–100%).



Prioritize "Update Conditions"

Shift the question from "What are you?" to "What evidence would move your confidence?"



Mental Framework Comparison

Feature	Label-Based Thinking	Credence-Based Thinking
Logic	Binary (Yes/No)	Gradient (Degrees of Confidence)
Goal	Social Sorting / Dismissal	Rational Inquiry / Understanding
Flexibility	Static (Protect the Identity)	Dynamic (Update for New Evidence)

Classificatory Humility

Use labels for initial orientation, but never as a final verdict on a person's reasons.



The infographic contrasts identity-protective binary sorting with a claim-specific credence profile. Its dial represents confidence under uncertainty; the comparison does not imply that every categorical description is false or useless.

Abstract

Public disagreement often turns a social label into a premature epistemic verdict. Once a speaker is classified as atheist, Christian, liberal, conservative, skeptic, activist, or fundamentalist, the label is made to explain the speaker's motive, determine credibility, and dispose of the argument. This paper argues that such *label substitution* is unwarranted. A broad identity label normally underdetermines the exact proposition affirmed, the degree of confidence, the supporting reasons, and the conditions for revision. Unless the user supplies a reliable and context-relevant bridge from the category to the property attributed, the label cannot by itself show that a person is irrational or that an argument fails.

The case combines conceptual analysis with work on degrees of belief, social categorization, essentialism, identity-protective cognition, and epistemic injustice. The central inference is deductive: a classification can support a specific verdict only by entailing the relevant property or by carrying a sufficiently reliable association in the present context; a broad label offered alone supplies neither. The surrounding psychology explains why the substitution is attractive and how it can become resistant to correction. A constructive alternative asks for the claim, confidence, evidence, defeaters, and update conditions before adjudication.

The conclusion is deliberately bounded. Categories are cognitively necessary, self-ascribed labels can communicate genuine commitments, and accurate group patterns can rationally affect initial expectations. Those concessions do not license movement from a group-level tendency to an unqualified judgment about one person. Labels should therefore remain provisional, relevance-tested, and defeasible: useful places to begin inquiry, never sufficient permission to end it.

Keywords: labels; credence; social identity; stereotyping; epistemic injustice; rationality; argument substitution

1 From Shorthand to Verdict

Human beings cannot reason without categories. We group objects, events, and people because finite attention cannot start from an unstructured world. Categories make communication possible and let past experience guide present expectation. The classic psychology of prejudice already recognized this cognitive economy: the problem is not classification as such but what is inferred from it and how rigidly the inference is held (Allport 1954). Research on judgment under uncertainty likewise shows why an easier classificatory question can replace a harder evidential one (Tversky and Kahneman 1974). “What sort of person is this?” is fast. “Does this premise follow, how strong is the evidence, and what alternatives remain?” is slow.

The danger begins when shorthand becomes a verdict. A speaker offers an argument about prayer, abortion, climate policy, moral realism, or economic inequality. Instead of addressing the claim, a hearer replies that the speaker is an atheist, evangelical, conservative, progressive, activist, or ideologue. The category may be true. It may even help explain why the speaker encountered the issue. Yet none of that shows that the inference is invalid, the evidence weak, or the conclusion false. Social location is not logical refutation.

Call this move *label substitution*: the hearer evaluates a category at the point where the claim, evidence, or speaker-specific credibility should have been evaluated. Label substitution is broader than name-calling. It can occur in calm academic prose and in sophisticated community sorting. The decisive feature is inferential: the label is made to carry a verdict that has not been connected to the property under assessment.

This paper concerns *epistemic* rationality—roughly, whether a person's confidence and reasoning fit the relevant evidence and respond properly to counterevidence. It does not claim that a complete assessment of practical or moral rationality can be reduced to a probability. Nor does it assume that

every ordinary belief must be stated numerically. The point is simpler: when uncertainty matters, a rational assessment needs more resolution than broad camp membership supplies.

2 What a Label Leaves Undetermined

An identity label can be accurate while remaining radically incomplete. Four dimensions of incompleteness matter for rational assessment.

2.1 Propositional content

First, the same word can cover different propositions. *Atheist* may mean lacking belief in deities, affirming that no deities exist, rejecting one religious tradition, or adopting a social identity. *Christian* may indicate creedal assent, sacramental membership, cultural inheritance, communal practice, or some combination. Political terms shift across countries, institutions, and decades. Until the relevant content is fixed, no definite premise follows from the label.

Context sometimes supplies that content. Membership in an organization with an explicit platform may raise the probability that a member accepts parts of it. A self-description offered in answer to a precise question may communicate more than a label imposed by an opponent. The lesson is not that meanings are forever inaccessible. It is that the user must establish which meaning governs before drawing a conclusion.

2.2 Degree of confidence

Second, threshold terms compress different degrees of belief. Two people may both cross the ordinary threshold for calling themselves Christian while one is tentative and the other nearly certain. Two people may both decline that label while assigning very different probabilities to central Christian claims. One may be undecided; another may regard the claims as extremely unlikely. The categorical boundary hides these differences.

Formal epistemology represents uncertainty through credences, or degrees of belief. Accuracy-based arguments show why coherent partial belief matters to epistemic rationality, while work on non-ideal agents explains how human reasoning can approximate such norms without exact numerical precision (Joyce 1998; Staffel 2019). The present argument needs no commitment to one complete theory of credence. It needs only the modest fact that confidence can be stronger or weaker and that this variation changes what criticism is appropriate.

Ordinary outright belief can coexist with degrees of confidence. A person may say “I believe the train will arrive” while being only moderately confident, because the context permits action above a threshold. The mistake is therefore not using a binary sentence. It is treating the sentence or identity term as if it revealed the hidden confidence, evidence, and stability of the attitude.

2.3 Reasons and evidential routes

Third, a label does not identify why the belief is held. People who share a conclusion may arrive through testimony, direct experience, argument, deference, habit, fear, community loyalty, or combinations of these. Rationality attaches partly to that route. The same conclusion can be responsibly or irresponsibly held depending on the evidence available to the person and the use made of it.

This is why a worldview attributed wholesale to an identity is usually a straw target. *Atheist* does not entail materialism, nihilism, scientism, or moral error theory. *Christian* does not entail one theory of scripture, atonement, politics, or science. Even when two commitments often travel together, frequency is not entailment. The connecting evidence must be stated rather than smuggled into the label.

2.4 Revision conditions

Fourth, a label does not show what would change the person's mind. One agent may treat the identity as a revisable summary. Another may experience change as betrayal of family, status, or community. Their public labels can be identical while their responsiveness to evidence differs sharply. Because responsiveness is central to epistemic appraisal, any verdict that ignores update conditions risks classifying social continuity rather than rationality.

The Label-Alone Argument

- P1.** A category can warrant a specific verdict about a person or argument only if it entails the relevant property or is connected to that property by a sufficiently reliable association in the present context.
- P2.** A broad identity label, considered by itself, encodes neither an entailment from membership to the attributed error, motive, or credibility nor a case-specific reliability relation. Any such relation must be established by further semantic, statistical, or behavioral evidence.
- P3.** What supplies neither route cannot by itself warrant the verdict.
- P4.** Therefore a broad identity label, offered alone, cannot warrant a verdict that the person is epistemically irrational or the argument defective.

The inference is deductively valid if the premises are granted. Premise 1 permits strict and probabilistic evidence. Premise 2 is not stipulative: broad identity membership does not itself contain the needed bridge. Statistics, institutional commitments, conduct, or self-definition may add one; then the verdict no longer rests on the label alone. The argument blocks automatic adjudication without denying that further evidence can make a category relevant.

3 Why Labels Become Epistemic Armor

The core error is conceptual, but psychology helps explain its attraction. Social labels reduce complexity and coordinate belonging. Social identity theory describes how group membership enters self-conception and intergroup comparison (Tajfel and Turner 1979). Once an identity organizes belonging, contrary evidence can threaten more than a proposition. It can threaten a network of relationships and a valued account of oneself.

Cognitive dissonance theory helps explain the pressure to reduce conflict between commitments and unwelcome information (Festinger 1957). One available strategy is revision. Another is to redescribe the source of the challenge as the kind of person whose testimony need not count. The label then becomes an epistemic shield: it protects confidence by lowering the standing of counterevidence before its content is assessed.

Kahan's synthesis of identity-protective cognition describes selective crediting and dismissal of information in ways aligned with group commitments (Kahan 2017). This pattern does not mean that every group member reasons badly or that identity is the sole cause of disagreement. It identifies a mechanism by which intelligence can serve coalition protection rather than correction. Label substitution is especially useful to that mechanism because it makes rejection feel evidential: the challenge has been classified, so it appears to have been answered.

Essentialist thinking adds another layer. Social categories are often represented as if they express a stable inner nature. Haslam and colleagues distinguish dimensions such as perceived naturalness,

discreteness, uniformity, and informativeness in essentialist beliefs about social groups (Haslam et al. 2000). When a political or religious category is treated as deeply uniform and highly informative, the distance from “belongs to this group” to “therefore has this motive and reasoning defect” begins to disappear. A defeasible summary has become an imagined essence.

These mechanisms explain why reciprocal dismissal is common. An apologist may hear *atheist* and infer rebellion, moral evasion, or materialism. A skeptic may hear *Christian* and infer indoctrination, fear of death, or hostility to science. Either story may happen to fit a particular person. Neither is licensed by the identity term alone. Symmetry matters: a method that excuses one’s own group from individuation while diagnosing opponents by category is not a neutral standard of inquiry.

4 When Labels Really Are Evidence

A defensible critique must not define every generalization as false or every use of group information as wrong. Beeghly argues for a non-moralized conception of stereotypes precisely because moral or epistemic defect should not be built into the definition before the case is assessed (Beeghly 2015). Jussim’s review presses a related empirical caution: some perceptions of groups correspond substantially to measured group differences, and accuracy at one level should not be prejudged by a theory of pervasive error (Jussim 2012).

These points improve rather than defeat the label-alone argument. A true group-level average is not automatically a warranted verdict about one member. Group means can overlap widely; base rates can change by context; the relevant person may be an exception; and a correlation with one property may say nothing about the argument presently offered. Moving from a group tendency to an individual judgment requires attention to predictive strength, variation, stakes, and available individuating information.

Some labels are more evidentially disciplined than others. *Cardiologist* is tied to training and licensure, so it can rationally raise initial confidence about a person’s competence in a specified medical domain. *Treasurer* may identify an institutional role and responsibility. A party officer may be reasonably expected to know the party platform. Even here the inference remains bounded: a cardiologist is not therefore correct about every diagnosis, and office does not prove honesty.

Self-ascription can also be informative. If a person says “By Christian I mean that I affirm these creeds,” the stipulated content creates an evidential bridge. If the same word is imposed from outside or used across contexts with different thresholds, the bridge weakens. The responsible question is not “Can labels ever provide evidence?” but “What exactly does this label predict here, how reliably, and what would defeat that prediction?”

Table 1 separates four uses that are often blurred.

The table exposes a common slide. A label that is serviceable for orientation is treated as though it already explains motive; a weak predictor becomes a verdict; a group pattern becomes an individual essence. The cure is not label-blindness but inferential bookkeeping. Each step must pay its own evidential cost.

5 Credibility, Prejudice, and Epistemic Injustice

Labels have moral as well as logical consequences when they govern who receives a hearing. Fricker’s central case of testimonial injustice involves a credibility deficit caused by identity prejudice: a speaker is wronged specifically in the capacity to contribute knowledge (Fricker 2007). This is narrower than any credibility adjustment based on background information. Expertise, track record, conflict of interest, and demonstrated dishonesty can rationally alter credibility. The injustice arises when a prejudicial identity association depresses credibility without the warrant the situation requires.

Use	What the label does	Evidence needed	Calibrated status
Orientation	Locates a topic, community, or rough region of belief.	Speaker meaning and conversational context.	Usually legitimate as a prompt for inquiry.
Prediction	Raises or lowers an initial expectation about a commitment or behavior.	A reliable association, relevant reference class, and known variation.	Defeasible background evidence, not certainty.
Explanation	Attributes a reason or motive for the person’s view.	Case-specific causal evidence beyond membership.	Often speculative unless independently supported.
Adjudication	Declares the person’s reasoning or claim defective.	The actual premise, inference, evidence, or demonstrated unreliability.	The label alone is insufficient.

Table 1: Four epistemic uses of a label. Increasing inferential ambition increases the evidence required; the table does not imply that orientation or probabilistic prediction is automatically prejudicial.

Label substitution can produce that wrong. “Of course she says that; she is an ex-Christian” converts a potentially relevant biography into a reason not to examine the testimony. “He is an evangelical, so the historical argument is propaganda” replaces source criticism with identity classification. The speaker is not owed agreement, unlimited attention, or immunity from background evidence. The speaker is owed a verdict responsive to the actual epistemic grounds when the label does not settle them.

The loss falls on both sides. The speaker may be excluded or misread; the hearer loses access to information that could correct error. At an institutional scale, repeated sorting can create credibility economies in which communities pre-authorize in-group sources and discount out-group sources. Members then experience the inherited filter as common sense rather than as a policy that itself requires evidence.

Accuracy and justice therefore converge on individuation. If diagnostic information about the particular claim or person is available, it should usually outrank a loose social proxy. A real correlation may justify looking more carefully; it does not justify refusing to look.

6 A Claim-Specific Alternative

The constructive alternative is a credence-first inquiry. *Credence-first* does not mean asking ordinary speakers for exact percentages before conversation can begin. It means refusing to let identity shorthand hide the epistemically relevant structure: what is believed, how firmly, for what reasons, against which alternatives, and under what conditions it would change.

Table 2 gives a compact protocol.

This protocol preserves the efficiency of shorthand while making its limits visible. A person can still say “I am an atheist” or “I am a Christian.” The next question determines what rational conversation follows. In some contexts, the identity will be central. In others, it will turn out to be almost irrelevant to the claim.

The protocol also clarifies the genetic fallacy. Learning where a belief came from can matter: a source may be unreliable, testimony dependent, or a conflict of interest undisclosed. What does not follow is that a claim is false or an argument invalid merely because of its origin. Genealogy becomes epistemically relevant only through a bridge to reliability, evidence, or inference. Labels can point toward that inquiry; they cannot complete it.

Question	What it recovers	Why it matters
What is the exact claim?	Proposition and scope.	Prevents a broad identity from importing a complete worldview.
How confident are you?	Degree or ordinary range: tentative, moderate, firm, near-certain.	Distinguishes shared labels that conceal different epistemic states.
What supports it?	Evidence, testimony, experience, argument, and dependence among sources.	Makes criticism attach to reasons rather than affiliation.
What counts against it?	Defeaters, alternatives, and acknowledged uncertainty.	Reveals whether counterevidence is allowed to matter.
What would change your view?	Update conditions and intellectual risk.	Distinguishes a revisable summary from identity-protective closure.
Is the label relevant here?	The proposed entailment or probabilistic bridge.	Forces the classifier to show rather than assume the label's evidential role.

Table 2: A claim-specific inquiry protocol. It does not guarantee agreement or demand numerical precision; it recovers the information required for a rationally targeted disagreement.

7 Objections and Replies

7.1 Conversation cannot proceed without simplification

Correct. The conclusion is not abolition but subordination. A subway map omits buildings because its purpose is route finding. It becomes misleading only when its limited geometry is used to settle a question it was not designed to answer. Identity labels can coordinate conversation, research, and community life. They fail when low-resolution orientation is made to adjudicate high-resolution questions.

7.2 People choose labels and should bear their implications

Self-ascription is evidence, especially when the label has an explicit public definition. Yet implication must still be distinguished from association. A person who joins an organization can reasonably be asked about its platform; membership does not entail endorsement of every member, policy, motive, or historical act. If the identity is being used to attribute a specific commitment, the commitment should be checked when the stakes justify it.

7.3 Some categories reliably predict dangerous commitments

Sometimes they do. High stakes can justify precaution before complete individuation, just as risk management can use base rates. But precaution is not a proof of guilt or irrationality. The stronger the sanction or dismissive verdict, the stronger the required connection between category and case. Emergency action may precede full inquiry; a final epistemic judgment should not pretend that the inquiry already occurred.

7.4 Credence profiles are artificial and unavailable

Exact numerical confidence is often unavailable, and self-report can be noisy. The alternative does not require false precision. Ordinary contrasts—tentative versus firm, central versus peripheral, revisable versus fixed—already recover information that a label suppresses. When numbers are useful, they can clarify thresholds; when they are not, qualitative degrees still improve the target.

7.5 Motives matter even when arguments are valid

Yes. Motives can explain selective evidence, predict behavior, or bear on trust. They should be treated as hypotheses supported by conduct, incentives, history, and response to correction. A valid argument does not purify a bad motive, and a good motive does not repair an invalid argument. Keeping those questions separate prevents psychological diagnosis from masquerading as logical assessment.

7.6 Individuation imposes unlimited conversational labor

No one owes a complete hearing to every claim. Attention is scarce, and people may rely on trusted filters. The thesis concerns the strength of the verdict, not an unlimited duty to engage. One may rationally decline a conversation for lack of time. What one may not infer from that practical choice is that the unheard argument has thereby been refuted or the speaker shown irrational.

8 Conclusion

Labels are not arguments because classification and adjudication are different epistemic acts. A broad identity term can locate a person in a rough social or intellectual region. It normally cannot reveal the exact proposition, confidence, evidence, defeaters, or conditions for revision that determine whether a particular belief is responsibly held. Without an entailment or a reliable context-specific association, the move from label to verdict is a non sequitur.

The surrounding evidence explains why the move persists. Categories save effort; identities organize belonging; essentialist habits exaggerate uniformity; and identity protection can turn out-group classification into a defense against correction. These mechanisms do not prove every label false. They show why the missing relevance bridge is easily overlooked.

The rational discipline is modest: use labels for orientation, test them as possible evidence, and never let them end inquiry at a resolution they cannot support. Ask what the person believes, how strongly, why, what counts against it, and what would change it. Apply the same standard across group boundaries. The result may still be sharp disagreement, but it will be disagreement with an agent and an argument rather than with a category-shaped substitute.

A Formal Reconstruction of the Central Thesis

1. A classificatory fact can warrant a specific epistemic verdict only if it either entails the property attributed or is connected to that property by a sufficiently reliable association in the relevant context.
2. A broad identity label, considered by itself, neither entails that the person is irrational or the argument defective nor encodes a reliable context-specific association to that verdict; establishing such an association requires additional evidence.
3. Therefore a broad identity label, offered alone, cannot warrant the verdict that the person is epistemically irrational or the argument defective.
4. Claim-specific evidence or a demonstrated relevance bridge may add warrant that the label alone lacks; the conclusion does not exclude those additions.

The inference is modus tollens on premise 1's necessary condition. That premise permits entailment and probabilistic support; premise 2 concerns what the label itself establishes and is supported by the four dimensions above. Statistics or other evidence can supply the missing bridge, so statement 4 preserves the scope limit. Psychology explains prevalence, not validity.

References

- Allport, Gordon W. (1954). *The Nature of Prejudice*. Cambridge, MA: Addison-Wesley.
- Beeghly, Erin (2015). “What Is a Stereotype? What Is Stereotyping?” In: *Hypatia* 30.4, pp. 675–691. DOI: [10.1111/hypa.12170](https://doi.org/10.1111/hypa.12170).
- Festinger, Leon (1957). *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press. DOI: [10.1093/acprof:oso/9780198237907.001.0001](https://doi.org/10.1093/acprof:oso/9780198237907.001.0001).
- Haslam, Nick, Louis Rothschild, and Donald Ernst (2000). “Essentialist Beliefs about Social Categories.” In: *British Journal of Social Psychology* 39.1, pp. 113–127. DOI: [10.1348/014466600164363](https://doi.org/10.1348/014466600164363).
- Joyce, James M. (1998). “A Nonpragmatic Vindication of Probabilism.” In: *Philosophy of Science* 65.4, pp. 575–603. DOI: [10.1086/392661](https://doi.org/10.1086/392661).
- Jussim, Lee (2012). *Social Perception and Social Reality: Why Accuracy Dominates Bias and Self-Fulfilling Prophecy*. New York: Oxford University Press. DOI: [10.1093/acprof:oso/9780195366600.001.0001](https://doi.org/10.1093/acprof:oso/9780195366600.001.0001).
- Kahan, Dan M. (2017). *Misconceptions, Misinformation, and the Logic of Identity-Protective Cognition*. Public Law Research Paper 605. Yale Law School. DOI: [10.2139/ssrn.2973067](https://doi.org/10.2139/ssrn.2973067).
- Staffel, Julia (2019). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford: Oxford University Press. DOI: [10.1093/oso/9780198833710.001.0001](https://doi.org/10.1093/oso/9780198833710.001.0001).
- Tajfel, Henri and John C. Turner (1979). “An Integrative Theory of Intergroup Conflict.” In: *The Social Psychology of Intergroup Relations*. Ed. by William G. Austin and Stephen Worchel. Monterey, CA: Brooks/Cole, pp. 33–47.
- Tversky, Amos and Daniel Kahneman (1974). “Judgment under Uncertainty: Heuristics and Biases.” In: *Science* 185.4157, pp. 1124–1131. DOI: [10.1126/science.185.4157.1124](https://doi.org/10.1126/science.185.4157.1124).