§ THE DIALECTIC REVIEW

Special Supplement · 2025 Cycle · 2,200 Submissions

PRIZE ’25

The Dialectic Prize

2025 Cycle · 2,200 Submissions · 6 pieces

Six essays selected from 2,200 submissions to the second annual Dialectic Prize, judged blind by the full advisor board against a published rubric: Argument (40%), Responsiveness to Objections (25%), Evidentiary Rigor (20%), Prose (15%).

ISSN 3049-7202

thedialecticreview.org

CONTENTS

1ST PLACE · PHILOSOPHY OF MIND 3

The Machine That Cannot Be Wrong: Trust, Testimony, and Artificial Authority

Sofia Marchetti — Liceo Classico Manzoni, Milan

2ND PLACE · ETHICS 7

Promises to the Dead

Arjun Mehta — Cathedral and John Connon School, Mumbai

3RD PLACE · POLITICAL PHILOSOPHY 11

The Tyranny of the Marginal Voter

Clara Fontaine — Lycée International, Geneva

HONORABLE MENTION · EPISTEMOLOGY 14

Silence as Speech: Rethinking the Right Not to Testify

Kwame Asante — Achimota School, Accra

HONORABLE MENTION · ETHICS 17

What the Trolley Problem Forgets

Yuki Tanaka — Azabu High School, Tokyo

HONORABLE MENTION · ETHICS 20

Owning a Mistake You Didn't Make

Emilia Novák — Gymnázium Jana Nerudy, Prague

1ST PLACE · PHILOSOPHY OF MIND

The Machine That Cannot Be Wrong: Trust, Testimony, and Artificial Authority

As automated systems increasingly deliver verdicts — diagnoses, credit decisions, translations — with no visible uncertainty, this essay asks what happens to human testimonial norms when one party to every exchange no longer performs doubt.

Sofia Marchetti Liceo Classico Manzoni, Milan

When a person tells you something, their manner ordinarily carries information alongside the content: a hedge, a pause, a raised eyebrow of their own uncertainty, all of which help you calibrate how much weight to put on the claim. A doctor who says “it’s probably a sprain, but let’s get it scanned to be sure” is giving you two things at once — a diagnosis and a confidence level — and competent listeners use both. This essay is about what happens to this arrangement when one party to the exchange is a system that, by construction, delivers every output in the same flat register of apparent certainty, whether the underlying answer is extremely well-supported or barely better than a guess, and argues that this uniformity is not a minor interface quirk but a genuine threat to how testimonial trust is supposed to work.

Human testimonial trust is calibrated, whether we notice it or not, by a long apprenticeship in reading confidence signals: tone, hedging language, willingness to say “I don’t know,” visible discomfort when pressed on a weak point. These signals are not decoration; they are

part of the actual evidence a hearer uses to set their credence in a claim, and centuries of face-to-face testimony have built our instincts around expecting them to correlate, reasonably reliably, with the speaker's actual epistemic position. A confidently delivered claim, from a human, really is some evidence the claim is well-supported, because confident delivery is costly to fake convincingly and sincere speakers' confidence usually does track their actual evidence.

Many automated systems break this correlation entirely, not through any intent to deceive, but through simple architectural fact: a language model or a classification system typically produces its most and its least reliable outputs in the identical surface register — the same fluent, assertive sentence structure, the same absence of hedging — because the system's training does not specifically optimize for making its own uncertainty visible in its output style the way a well-socialized human speaker does. A system asked a question deep within its training distribution and a system asked a question at the ragged edge of what it has ever seen reliable data about can produce answers that are, to a human reader applying ordinary testimonial instincts, indistinguishable in confidence — both fluent, both assertive, both giving no surface signal of the enormous difference in actual reliability between them.

This is a distinct problem from the more commonly discussed issue of a system simply being wrong, because a system that is wrong but signals its own low confidence is, in an important sense, testimonially honest — a hearer applying ordinary norms would appropriately discount the claim, exactly as they would discount a human expert who says "I'm really not sure about this one." The problem this essay is concerned with is narrower and more insidious: a system that is unreliable *and gives no signal of this*, because its architecture does not link internal uncertainty to surface confidence the way evolved and socialized human communication reliably (if imperfectly) does. The

hearer’s testimonial instincts, built for a world where confident delivery correlates with actual reliability, are not merely unhelpful here; they are actively miscalibrated, because the correlation they were built to exploit does not hold for this kind of speaker at all.

One response is to demand that systems be engineered to express calibrated uncertainty — to hedge more when they are less reliable, mimicking the human signal. This is a real and valuable engineering goal, and some systems now attempt exactly this, through techniques that surface a model’s own estimated confidence alongside its answer. But it faces a structural limitation worth taking seriously: a system’s *expressed* confidence, even when engineered to track something, is tracking the system’s internal estimate of its own reliability, which is itself produced by the same kind of process that produced the potentially unreliable answer in the first place — a system can be simultaneously wrong about a claim and wrong, in a correlated way, about how confident it should be in that claim, in a way a well-socialized human, whose confidence is calibrated by an entirely separate developmental process (social feedback on being embarrassingly overconfident, for instance) is comparatively less likely to be. Engineered confidence expression narrows the problem; it does not obviously dissolve it, because it inherits the same basic uncertainty about whether the system’s self-model is any more reliable than its object-level model.

What follows, I think, is not a case for abandoning trust in automated testimony, which would be as overcorrected as naive full trust, but a case for a specific and unfamiliar discipline: since the hearer cannot rely on the traditional signal (does this feel confident), they need to substitute a different, more effortful check — asking not how the system said it, but what kind of question this is, and whether this is a domain with the dense, verifiable training data that makes any output, confident-sounding or not, actually likely to be reliable. This

is a harder habit to build than reading tone, because it requires knowing something about the system's training and validation, not just attending to how an answer is delivered, and it inverts centuries of testimonial habit that let confident delivery itself serve as usable evidence. The machine that cannot be wrong, in the sense that matters here, is not a machine that is actually always right. It is a machine that has never learned to sound any less certain when it is not, and a hearer who has not learned to stop listening for the difference is trusting exactly the wrong thing.

2ND PLACE · ETHICS

Promises to the Dead

A deathbed promise binds the living to someone who can no longer release them from it, benefit from its keeping, or be wronged by its breaking in any ordinary sense. This essay asks what, then, the promise is actually binding.

Arjun Mehta Cathedral and John Connon School, Mumbai

A daughter promises her dying mother she will care for the family garden for the rest of her life. The mother dies. Decades later, the daughter, now old herself, wonders whether she is still bound. Nothing about the standard machinery of promissory obligation applies cleanly here: the mother cannot release her from the promise, cannot be disappointed or relieved by what happens to the garden, cannot have any of her interests set back by neglect, since the dead, on most accounts of harm, cannot be harmed by events occurring after their death. And yet the pull most people feel toward honoring such a promise is, if anything, stronger and more solemn than the pull attached to an ordinary promise to a living person who could simply release you from it with a word. This essay asks what a promise to the dead is actually binding, given that nearly every standard account of promissory obligation seems to have nothing left to grip once the promisee is gone.

The most natural first move is to relocate the obligation from the dead person to their prior, living interests: perhaps what is owed is not owed to the currently-dead mother at all, but was owed to the once-living mother’s interest, held while she was alive, in her garden’s future — an interest that does not need her to currently exist in order

to have been a real interest she held, and does not need her to currently exist in order for its having been respected or betrayed to matter. This is the shape of Joel Feinberg's account of posthumous harm, which argues that a person can have interests extending beyond their death — in their reputation, in their wishes being carried out, in projects they cared about continuing — and that these interests can be set back by post-mortem events, even though the person is no longer around to experience the setback.[1] On this view, breaking the promise would genuinely harm the mother, not now, but retroactively, by defeating an interest she held while alive.

This account has real appeal but faces a hard metaphysical problem that Feinberg himself wrestled with: harm, on almost any standard analysis, requires a subject who is made worse off, and it is deeply unclear how an event occurring after a person's death can make that person — who, by hypothesis, no longer exists to be in any state at all, worse or better — worse off in any literal sense. We can make sense of a living person's interest being frustrated because there is a subject around whose life goes worse as a result; there is no obvious parallel subject for whom the garden's neglect makes anything go worse, once the mother has ceased to exist altogether. Feinberg's own solution — that the interest can be harmed "retroactively," changing the truth-value of how well her life, viewed as a completed whole, actually went — is ingenious but strains our ordinary concept of harm considerably, requiring us to accept that events after a life's end can reach backward and alter facts about that life's success, which is a strange kind of causation-in-reverse for a merely evaluative fact to undergo.

I want to suggest a different, less metaphysically strained location for the obligation, one that does not require positing harm to the dead at all: the promise binds the daughter not to her mother, who cannot now be wronged, but to the person the daughter herself has become by virtue of having made the promise — specifically, to the kind of

person who keeps solemn commitments made at moments of maximal seriousness, a commitment that was not merely about the garden but was, at the moment of making it, a statement about what kind of daughter, and what kind of promise-keeper, she was choosing to be for the remainder of her own life. On this view, breaking the promise after the mother’s death is not a wrong done to the mother — there is no longer anyone there to be wronged in that direction — but a wrong done to the coherence of the daughter’s own moral identity, a betrayal of the specific, solemn self-conception she committed herself to at her mother’s bedside, in the presence of the last words she would ever say to her.

This reframing has a real cost, which honesty requires naming: it makes the promise, in the end, a duty to oneself rather than a duty to the dead, which will strike some readers as deflationary, as though it quietly concedes that promises to the dead were never really about the dead at all. I think this concession is correct, and that it is not actually deflationary once we notice what solemn promise-making at a deathbed was always doing, psychologically and morally, even for people who never asked the metaphysical question explicitly. The daughter kneeling at the bedside was not merely transferring an instruction; she was constituting herself, in that moment, as a particular kind of continuing agent, and it is that act of self-constitution — not a debt owed to a person who no longer exists to collect it — that a broken promise to the dead actually violates. This does not make such promises less serious. It explains, better than the alternative, why they so often feel more serious than promises to the living: because what is actually at stake was never releasable by the other party in the first place, precisely because it was never, in the end, only about the other party at all.

NOTES

1. Joel Feinberg, “The Nature and Value of Rights,” and later work on posthumous harm collected in *Harm to Others* (1984).

3RD PLACE · POLITICAL PHILOSOPHY

The Tyranny of the Marginal Voter

Elections are won and lost at the margin, and campaigns rationally target the voters least committed to any position. This essay asks what this does to the incentives facing everyone else.

Clara Fontaine Lycée International, Geneva

In a competitive election, the votes that determine the outcome are, definitionally, the votes of whichever small population of persuadable, weakly-committed voters happens to sit at the margin — the people genuinely undecided, or weakly attached to either side, whose choice will actually swing the result, as distinct from the much larger population of firmly committed voters on either side whose votes are, in the relevant strategic sense, already accounted for. This essay is not about the familiar complaint that campaigns cater excessively to the political center; it is about a subtler and less discussed consequence of this dynamic — that a system optimized to win over the marginal voter systematically rewards the specific psychological profile of the marginal voter, and asks what happens to democratic deliberation once this reward structure is fully internalized by political actors.

The marginal voter, as extensive political science research on persuadable and swing populations has documented, differs systematically from the committed voter not merely in their conclusions but in the cognitive style that produces those conclusions: lower baseline engagement with policy detail, higher sensitivity to personality, tone, and immediate emotional resonance

relative to programmatic consistency, and a documented tendency to weight recent, vivid events far more heavily than earlier, better-evidenced ones. None of this is a moral criticism of marginal voters as people; it is simply a description of why they are marginal in the first place — voters with strong, stable, well-examined commitments are, almost by definition, less movable, and the population that remains movable is disproportionately composed of people for whom the relevant considerations were never fully settled to begin with.

The strategic consequence follows almost mechanically. A rational campaign, seeking to win, will allocate disproportionate persuasive effort toward exactly the considerations that move the marginal voter — tone, personality, vivid recent events, emotional resonance — rather than the considerations that would matter most to a well-informed, carefully reasoning citizenry, because effort spent on the latter, while civically admirable, is strategically almost wasted: it will not move the committed voters, who are already decided, and it is comparatively less effective at moving the marginal voters than the former set of appeals. This is not a claim that campaigns are run by cynics indifferent to substance; it is a claim about what any campaign, however well-intentioned, is rationally incentivized to prioritize, given the actual, empirically documented psychology of the population whose behavior determines the outcome.

This produces what I want to call the tyranny of the marginal voter: not that marginal voters have illegitimate influence, which they do not — they are voting exactly as legitimately as anyone else — but that the entire apparatus of political communication, across all sides of a competitive system, is structurally bent by rational actors toward the cognitive style of a population whose defining feature is comparatively weak engagement with exactly the kind of information a well-functioning democracy is supposed to reward.

The tyranny is structural rather than personal: no single actor is doing anything illegitimate, and yet the aggregate effect, across an entire competitive system optimizing simultaneously toward the same marginal population, is a communicative environment increasingly calibrated to reach people who are, again by simple definition of what makes them persuadable, less anchored by the kind of stable, considered reasoning democratic theory usually assumes citizens bring to the polls.

The obvious objection is that this overstates the case: campaigns also need to mobilize their committed base, whose enthusiasm and turnout matter as much as marginal persuasion, which should pull communication back toward substantive appeals that reinforce, rather than merely attract, existing commitment. This is a real countervailing force, and it means the tyranny described here is a tendency rather than an iron law, its strength varying with how close and how turnout-sensitive a given election actually is. But the tendency is real and measurable wherever elections are genuinely competitive at the margin, which is precisely the condition democratic theory treats as healthiest — a landslide election, ironically, faces less of this pressure than a closely contested one, since there is no large, movable marginal population worth specifically targeting. This suggests an uncomfortable possibility that deserves more attention than it usually gets: that the closeness and competitiveness we generally prize in elections, as evidence of a genuinely contested, responsive democracy, is exactly the condition that most strongly incentivizes political communication to bend toward the psychology of its least deliberative participants, and that this cost, real and structural, is rarely weighed against the benefits of competitiveness we otherwise, and rightly, prize.

HONORABLE MENTION · EPISTEMOLOGY

Silence as Speech: Rethinking the Right Not to Testify

The right against self-incrimination is usually justified as protection from coercion. This essay argues it is better understood as protecting a specific kind of communicative act that speech itself cannot perform.

Kwame Asante Achimota School, Accra

The right against self-incrimination — the right to remain silent rather than answer questions that might expose oneself to criminal liability — is standardly justified as protection against a cruel dilemma: forcing someone to choose between perjury, self-accusation, and contempt of court is a form of coercion the state should not be permitted to impose. This essay argues that this justification, while correct as far as it goes, misses something distinctive about silence itself as a communicative option, and that recognizing this changes how we should think about contexts, well beyond the courtroom, where a right to silence is contested.

The standard, coercion-based justification treats silence as a merely negative option — the absence of a forced answer, valuable only because the alternative (compelled speech) would be objectionable. On this view, silence has no positive content of its own; it is simply what remains once an unacceptable option (compelled self-accusation) is removed. This framing explains why the right exists, but it does not explain a puzzling feature of how silence actually functions once granted: a defendant's silence is typically treated, by rule, as carrying no evidentiary weight at all — juries are instructed

not to infer guilt from it — even though, as a matter of ordinary human inference, silence in the face of an accusation is, in nearly every other context, treated as significant information, often as evidence against the silent party.

This gap suggests silence is not merely a negative absence in the legal context; it has been deliberately and artificially stripped of its ordinary evidentiary content, precisely because the legal system has decided that this particular communicative act — remaining silent under a specific kind of official accusation — needs to be available to the innocent and the guilty alike without cost, in a way that no other communicative option can be, including the option to speak. A person who testifies and lies faces perjury liability; a person who testifies and tells an unhelpful truth has, in an important sense, no further options once the words are out. Silence alone offers a genuinely cost-controlled position: it commits its holder to nothing, forecloses nothing further, and (by legal fiat) carries no inferential content that could be held against them later. This is not merely the absence of a bad option; it is the deliberate legal construction of a specific, protected communicative act — refusal to answer this specific kind of question, in this specific context — with a stipulated meaning (no inference permitted) that ordinary silence, outside this narrow legal construction, does not carry at all.

Seeing the right this way — as constructing a specific, artificially meaning-neutral communicative act, rather than merely permitting an absence — clarifies why the right’s scope is, and should be, narrower than a general “right to say nothing” might suggest. The specific legal construction only strips silence of its ordinary evidentiary content within the narrow context it was built for (formal accusation, under official compulsion, in specified proceedings); it does not and should not extend that same artificial neutrality to silence in other contexts — an employer, for instance, is not similarly

barred from drawing conclusions from an employee's silence when directly and appropriately asked about workplace conduct, because no comparable legal construction has stripped that silence of its ordinary inferential content. The right against self-incrimination, on this view, was never a general philosophical claim that silence is always meaning-neutral; it was always a specific, deliberately constructed legal artifact, valuable exactly because of how narrow and context-bound its artificial neutrality actually is.

This has a practical implication for debates about extending "right to silence" style protections into new domains — workplace investigations, school disciplinary hearings, informal community proceedings — where advocates sometimes argue the same right should simply transfer by analogy. The analysis here suggests this transfer cannot be automatic: what needs to be argued, in each new context, is not merely that compelled speech would be uncomfortable for the accused (true in nearly every accusatory context, and not sufficient on its own), but that the specific, artificial construction — silence stripped of its ordinary evidentiary meaning — is worth building afresh in this new context, which is a real, substantive institutional decision with real costs (some guilty parties will benefit from an inference that would otherwise have been drawn), not a default extension of a general right to say nothing that was actually never general in the way the popular phrase suggests.

HONORABLE MENTION · ETHICS

What the Trolley Problem Forgets

Trolley cases strip away nearly everything except the number of lives at stake. This essay argues the stripping itself is where the philosophical action actually is.

Yuki Tanaka Azabu High School, Tokyo

The trolley problem, in its many variants, asks whether it is permissible to divert or otherwise use force to redirect a runaway trolley so that it kills one person instead of five, and has generated an enormous philosophical literature comparing our differing intuitions across structurally similar cases — diverting the trolley onto a side track feels permissible to most people, while pushing a bystander off a footbridge to achieve the identical numerical outcome does not, and much of the literature is devoted to explaining this asymmetry through principles like the doctrine of double effect. This essay does not propose a new solution to the trolley problem’s central puzzle. It argues instead that the case’s enormous popularity has had a distorting effect on moral philosophy, by training generations of students to treat morally relevant information as noise to be stripped away, when in ordinary moral life, that stripped-away information is usually where the actual moral work gets done.

The trolley case is deliberately constructed to remove nearly everything a real moral situation would contain: we are told nothing about who the five people are, who the one person is, whether any of them behaved recklessly or heroically to end up on the tracks,

whether the agent has any special relationship to any of them, what their deaths would mean to anyone left behind, or whether alternative options exist that the stipulated two-choice structure has simply defined out of existence. This stripping is a deliberate methodological choice, meant to isolate the effect of a single variable (the numbers, the mechanism of harm) by holding everything else constant — exactly the logic of a controlled experiment, and a genuinely useful one for isolating whether, and how much, the doing/allowing distinction or the intended/foreseen distinction affects moral judgment independent of any confounding detail.

The trouble is not with the method as applied to its narrow, intended question, but with what happens when several generations of students and, increasingly, popular culture absorb the trolley problem as the paradigm case of moral reasoning itself — as though real moral situations were trolley cases with the details filled back in, rather than trolley cases being real moral situations with nearly all the details deliberately removed. In actual moral life, the information the trolley case strips away is rarely irrelevant noise; it is usually the primary bearer of moral significance. Whether the one person on the side track is a stranger or your child, whether the situation arose through anyone's negligence, whether the five people got themselves into danger through reckless disregard for posted warnings, whether an alternative you have not yet considered actually exists — these are not distractions from the real moral question in an ordinary situation; for most people's actual moral reasoning, they typically *are* the real moral question, and the pure headcount comparison the trolley case isolates is a comparatively rare and impoverished way of approaching real moral choices.

This produces a specific and, I think, underappreciated distortion when trolley-style reasoning is imported wholesale into real policy and ethical debates — triage protocols, autonomous vehicle

programming, pandemic resource allocation — that involve genuine trade-offs between different numbers of lives. Importing the trolley frame encourages treating the numerical comparison as though it were the whole of the relevant moral question, when the stripped-away particulars (whose negligence created the shortage, what special relationships or obligations exist, whether the categories being compared are actually as clean and mutually exclusive as the trolley’s two tracks) are often doing more moral work in the real case than the pure number comparison the trolley problem trained us to isolate and focus on. A policy debate conducted entirely in trolley-problem terms — “would you divert harm from five onto one” — has already, often silently, decided that the particulars do not matter, which was a useful stipulation for isolating one variable in a thought experiment and a substantive, contestable, and frequently false assumption when carried into an actual policy debate with real particulars still very much present and relevant.

None of this is an argument against the trolley problem as a piece of methodology within philosophy, where its narrow, isolating purpose is well understood by the specialists using it. It is an argument about pedagogy and public reasoning: a thought experiment built specifically to strip moral situations of their particulars, in order to isolate one variable for careful study, should not be mistaken for a template of what moral reasoning looks like once the particulars are put back — and a great deal of confused, particulars-blind public reasoning about real trade-offs, I suspect, comes from exactly this mistake, made by people who encountered the trolley problem as their first and sometimes only formal introduction to moral philosophy, and absorbed its deliberate impoverishment as though it were a discovery about what moral reasoning essentially consists of, rather than a methodological choice suited to a narrow, isolated question and poorly suited to almost everything else.

HONORABLE MENTION · ETHICS

Owning a Mistake You Didn’t Make

Institutional apologies are routinely issued by people who had no part in the original wrong. This essay asks what, exactly, such a person is doing when they say ‘we were wrong.’

Emilia Novák Gymnázium Jana Nerudy, Prague

A university president apologizes for admissions decisions made by administrators fifty years dead; a company’s current CEO apologizes for environmental damage caused under a previous, unrelated management team; a national leader apologizes for historical injustices committed by governments generations removed from the one currently in office. In each case, the person delivering the apology bears no personal responsibility, in any ordinary sense, for the original wrong — they did not do it, often were not born when it happened, and could not have prevented it. This essay asks what such a person is actually doing when they say, apparently sincerely, “we were wrong,” given that the most natural analysis of apology (acknowledging one’s own wrongdoing) does not seem to apply to them at all.

One response is deflationary: institutional apologies of this kind are not really apologies in the full moral sense, but a different speech act entirely — an official acknowledgment of historical fact, perhaps combined with an expression of regret that such things happened, dressed in apology’s grammar for rhetorical and diplomatic effect, without actually claiming personal moral responsibility for anything. This deflationary reading has some appeal, since it avoids the puzzle

of assigning personal guilt to someone who did nothing wrong. But it does not fully match how these apologies are actually received and evaluated — audiences do seem to hold current officeholders to a standard beyond bare historical acknowledgment, judging some institutional apologies as more adequate, more genuine, and more binding than others, in ways that a mere “yes, this happened, and it was bad” statement would not obviously earn or deserve.

A more interesting analysis treats the apologizer not as claiming personal responsibility for the original act, but as claiming a specific kind of representative responsibility, grounded in their current occupancy of a persisting institutional role that has continuity with the role occupied by the original wrongdoers. This is a real and independently recognized category of responsibility: a company’s current leadership can be genuinely on the hook, legally and morally, for correcting ongoing harms set in motion by a previous administration, precisely because the company persists as the same institutional entity across the change in personnel, in a way that makes its current officeholders the appropriate parties to address in the present, whatever their personal innocence regarding the original decision. The apology, on this view, is sincere and appropriately made by the current officeholder not because they did anything wrong, but because they currently hold the specific institutional position whose occupant is the correct party to speak for, and act on behalf of, the persisting entity that did.

This account explains a real and otherwise puzzling asymmetry in how we judge these apologies: we consider an institutional apology considerably weaker, or even suspect, when it is not accompanied by any actual, present-day action within the apologizer’s current power — correcting an ongoing harm, changing a policy that perpetuates the original wrong’s effects, providing some form of redress the institution is presently capable of providing. If the apology were

merely historical acknowledgment, this expectation would be puzzling; historical facts do not require present-day correction to have been accurately acknowledged. But if the apology is a claim of representative responsibility grounded in persisting institutional role, the expectation makes complete sense: representative responsibility, unlike personal guilt for a past act, is inherently forward-looking as much as backward-looking, because the role itself persists into the present and carries whatever powers and duties the present moment actually makes available, in a way personal guilt for a specific past act does not automatically generate any parallel present-day duty to act.

This reframing also explains why some institutional apologies ring false in a specific and recognizable way — not because the officeholder lacks genuine feeling, but because they are, whether they realize it or not, attempting to borrow the emotional register of personal apology (I did wrong, I am sorry) for an act that is actually representative apology (this institution, which I currently represent, did wrong, and here is what I am doing about it now), and the mismatch between the register borrowed and the claim actually being made is what produces the frequent, justified sense that something in the apology does not quite add up. Getting institutional apology right, on this view, requires the apologizer to be honest about which of these two very different things they are actually doing — and the more useful and more honest of the two, however less immediately satisfying it sounds, is very often the second.

MASTHEAD

The Dialectic Prize — PRIZE ’25

CHAIR, PRIZE COMMITTEE

Prof. Elena Marsh

Princeton University

JUDGE, EPISTEMOLOGY

Dr. Amit Sharma

University of Oxford

JUDGE, ETHICS

Prof. Naomi Katz

University of Toronto

JUDGE, PHILOSOPHY OF MIND

Dr. Julian Voss

National University of Singapore

JUDGE, METAPHYSICS

Dr. Helena Bruch

University of St Andrews

JUDGE, POLITICAL PHILOSOPHY

Prof. Daniel Osei

University of Melbourne

EDITOR-IN-CHIEF

Aanya Bhatt

Cathedral and John Connon School, Mumbai