Awareness and the Field of Intelligibility

Riyan Christy

2026

Awareness and the Field of Intelligibility

This book defends a thesis with deep philosophical precedent: awareness is not one item within the inventory of what appears, but the field within which any such inventory becomes intelligible. Versions or near-relatives of this thesis appear in Advaita Vedānta, in Yogācāra Buddhism, in Husserlian phenomenology, in Sartre's account of the pre-reflective *cogito*, and in contemporary work on consciousness and self-awareness by figures such as Zahavi, Thompson, Strawson, Kastrup, Goff, Shani, and Albahari. The present book does not claim to originate the thesis. Its contribution is methodological. It offers a disciplined eliminative argument against the main ways awareness is mislocated: as object, function, representation, reportable access, temporal synthesis, neural correlate, or countable subject.

The book's argumentative promise can be stated in one sentence. Attempts to derive, identify, or individuate awareness inevitably rely on the field in which the relevant derivation, identification, or individuation is made intelligible. That sentence is the work's discipline, not its conclusion. Each major chapter takes one form of the relevant attempt — derivation, identification, temporal subordination, individuation — and asks what its load-bearing terms commit it to. The chapters do not argue that awareness has some hidden positive feature. They argue that the proposed reductions and locations cannot be stated within the field they propose to subordinate.

A few things the book is therefore not. It is not idealism. It does not claim that the external world is mind-dependent, or that physical inquiry into matter is invalid, or that empirical science overreaches when it describes neural, cognitive, or computational organization. The argument does not contest first-order results in cognitive science, neurology, or psychology. It contests one philosophical reading sometimes attached to those results: that the items the sciences describe could subordinate the field in which their descriptions become intelligible. The argument is not mysticism. Awareness, on the view defended here, is not ineffable. The book consists of saying things about it within an eliminative discipline; that discipline is the work, not a refusal of work. The argument is not anti-logic. Its proof discipline is stricter than the classical inheritance it constrains; it asks that conclusions not outrun what can be shown. And the argument is not verificationism. The referential discipline it employs is local, narrow, and operative only inside ontological dependence claims, where it asks whether a term has acquired enough determinacy to make the dependence relation it figures in do argumentative work.

What the book asks for, from a reader, is patience with a particular kind of move. A reduction of awareness will be stated as forcefully as its proponent states it; its terms will be granted whatever local legitimacy they possess; the question will then be whether the reduction can bear the load its own form assigns it. When it cannot, the failure is not a rhetorical victory but a constructive exposure of an internal limit. Nothing has been imported from outside the proposal. The proposal has been followed to the point where it no longer holds. The book's arguments are arguments of that kind throughout.

What remains to be said in the chapters that follow is therefore not the announcement of a new metaphysical object. It is the demonstration that a thesis with long philosophical history — awareness as

the field of intelligibility rather than as an item within it — survives, as a disciplined eliminative result, against the contemporary attempts that have most strongly pressed against it.

Introduction — Theories of Consciousness Begin Too Late

0.1 The problem

Theories of consciousness begin too late. They take as their starting point a field of intelligible items — neurons, representations, functional states, information structures, conscious episodes, perspectival contents — and ask, within that field, what makes some of those items conscious and others not. The answers offered have grown sophisticated: higher-order representational theories of conscious states, global workspace and access-architectural theories of broadcast and availability, illusionist accounts of why our intuitions about phenomenal experience mislead us, transcendental and phenomenological accounts of the conditions under which any item appears in time and from a perspective. Each of these answers operates competently within its proper jurisdiction. Each is, in its own register, a serious philosophical contribution.

None of them addresses the question this book is concerned with. The question is one level adjacent to the questions the theories address. It is this: the field within which any of these accounts is intelligible as an account — within which higher-order representations are intelligible as representations, functional states are intelligible as functions, conscious episodes are intelligible as episodes, perspectival contents are intelligible as contents — is itself not one of the items the accounts catalogue. The accounts presuppose the field; none of what they offer supplies the field as an item within their account. To begin an account of consciousness from within an inventory of content-bearing items is to have begun at a level downstream from the question of awareness as such.

This is what "begin too late" means. The book is not claiming that contemporary theories of consciousness are wrong in their domains. It is observing that their domains are not where the question of awareness as field has its purchase. Each theory does its work; the question of the field within which the work is intelligible as work remains. The book has been written to ask that question with the discipline it requires, and to argue, against the principal attempts to dissolve or reframe it, that the question is not answered by any further item within the inventory the theories assemble.

The structural diagnosis can be stated briefly, though its development will take chapters. Whenever a contemporary theory proposes that consciousness is to be derived from, identified with, or located among items within a content-architectural account, the proposal turns on a transition from items determinate enough to play the proposed role to a feature of consciousness those items are meant to explain. The transition assumes that consciousness is the kind of feature an item-account could explain. The book argues, through chapter after chapter, that the assumption is what fails. Items, properties, structures, dependence relations, and grounding relations are intelligible to us as the items, properties, structures, dependence relations, and grounding relations they are because they sit within a field of in-

telligibility. The field is what the contemporary theories presuppose but cannot deliver. Their inability to deliver it is not a defect of their accounts; it is a feature of what they are doing. The question of the field is a different question, and the book is the discipline of asking it.

0.2 The thesis

The book defends one central thesis.

Awareness is not one item within the inventory of what appears, but the field within which any such inventory becomes intelligible.

Three connected refusals follow from this thesis, and a fourth positive observation. The book argues, first, that awareness is not ontologically derivative from anything determinate enough to play the role the derivation would assign it. Anything determinate enough to figure as the deriving-base of awareness is already an item within the field of intelligibility, and the derivation cannot reach awareness as field without already presupposing it. Second, awareness is not identical to any content, content-relation, function, representation, or access architecture. Anything determinate enough to play the identifying role is itself content within the field rather than the field within which content has its standing. Third, awareness is not subject to the kinds of structural conditioning — temporal, perspectival, modal — that operate on items within the field. Temporal structure is a feature of how content unfolds; perspectival individuation is a feature of how items are encountered from a body or a history; neither extends to a condition governing the standing of the field within which such conditioning is intelligible.

The positive observation, which Chapter 8 develops directly, is that awareness as field is the standing within which any cataloguing of items, properties, structures, and relations is itself intelligible as cataloguing. The discipline that does the cataloguing — ontology, broadly construed — has a jurisdiction within which its methods operate. The field is not within that jurisdiction as an item to be catalogued. It is the standing within which the discipline's own operations are intelligible. The book's argument is jurisdictional rather than itself ontological: a structural observation about what kinds of question the inventory-discipline can answer, and where it reaches the limit of its own jurisdiction.

The thesis has deep philosophical precedent. Versions or relatives of it appear in Śaṅkara's Advaita Vedānta as the *sākṣin* (witness) doctrine, in the phenomenological tradition's accounts of intentionality and pre-reflective self-awareness, in parts of contemporary philosophy of mind that resist deriving consciousness from non-conscious matter, and in traditions that may warrant later comparison, including apophatic theology. Chapter 9 engages Śaṅkara as the closest classical precedent. The book's contribution is not the bare thesis. It is a methodological one.

0.3 The contribution: methodological, not first-doctrinal

The book does not claim to originate the thesis. Versions of it have been articulated for over a millennium in the classical Indian tradition, with comparable rigour and at comparable scale. Contemporary work in the phenomenological tradition and parts of analytic philosophy of mind have articulated structurally similar concerns; other possible comparisons, including apophatic theology, are held back for source-disciplined later work. Acknowledging this is part of the book's intellectual honesty: it is not in a position to claim that the structural shape of its central thesis is unprecedented, and the chapter on philosophical antecedents (Chapter 9) makes the closest classical precedent explicit.

The book's contribution is methodological rather than first-doctrinal. What is offered is a disciplined eliminative argument, run under a narrowly specified referential constraint, against six structurally

distinguishable ways in which awareness can be mislocated within content-architectural features of conscious life. Each of the six receives a chapter, with the contemporary opponent who has articulated the relevant mislocation most carefully serving as the chapter's worked case.

The first move the book negates is *ontological derivation*: the claim that awareness can be derived from items more basic in the ontological order. Chapter 3 develops this against the family of physicalist proposals that attempt such derivation, with care given to the Stroud-style transcendental objection. The second move is *identification*: the claim that awareness is identical to a specific content-architectural feature — a higher-order representation, an access function, a representational state at some particular cognitive level. Chapter 4 develops this against Carruthers and Block. The third move is *reframing*: the claim that the problem of awareness is the problem of explaining our intuitions and reports about consciousness, with the substantive question dissolving or relocating in the explanation. Chapter 5 develops this against Frankish and Chalmers. The fourth move is *temporal subordination*: the claim that awareness, being temporally structured, is subordinate to the conditions governing temporal structure. Chapter 6 develops this against Kant's doctrine of time as the form of inner sense. The fifth move is *perspectival multiplication*: the claim that the manifest plurality of perspectives entails the plurality of awareness-fields. Chapter 7 develops this against the phenomenological articulation in Husserl, Sartre, and Zahavi. The sixth move is *ontological placement*: the constructive claim that awareness is the deep nature or fundamental property of some metaphysical absolute. Chapter 8 develops this against Goff's constitutive cosmopsychism.

What unifies the six chapters is procedural rather than doctrinal. Each chapter grants the opponent's argument within its proper jurisdiction. None of them is anti-Carruthers, anti-Block, anti-Frankish, anti-Chalmers, anti-Kant, anti-Husserl, anti-Sartre, anti-Zahavi, or anti-Goff. Each opponent has done careful work; each is granted whatever it establishes within its own programme; the book takes none of them as failed within their own programmes. What each chapter does is locate a structural pressure-point adjacent to the opponent's apparatus — the point at which the opponent's account either reaches or does not reach the field within which its operations are intelligible. The pressure-point is structural in every case; the diagnosis is the same in form across all six chapters; the differences are in the specific content-architectural features each opponent has proposed.

The six chapters together form one cumulative argument rather than six independent engagements. Each operates on a structurally distinct way of mislocating awareness — derivation, identification, reframing, temporal subordination, perspectival multiplication, ontological placement — and each runs the same diagnostic on its specific case. The diagnostic is invariant; the cases are not. No chapter is structurally redundant, and no chapter functions merely as illustration of another. What the arc as a whole establishes is what no single chapter could establish on its own: that the family of attempts to subordinate awareness to items within the field of intelligibility shares a structural exhaustion-point, and that the exhaustion-point is not an artefact of any one opponent's apparatus but a feature of the kind of claim each opponent has tried to make.

This is what the book means by "disciplined eliminative argument." The book proceeds by ruling out each candidate way of locating awareness within an inventory of items, properties, structures, or relations, without proposing a positive ontological location of awareness in their place. The cumulative effect of six chapters of such argument is not the construction of a positive alternative ontology — the book has none — but the marking of a jurisdictional limit. Ontology has its scope; the question of awareness as field is at the jurisdictional edge of that scope. The book's discipline is the discipline of articulating where the edge falls and why.

The discipline matters because the result it produces is positive without being ontological. The book is

not arguing that something is missing from contemporary ontology and then declining to supply it; the book is establishing that the question of awareness as field is the kind of question ontology's apparatus cannot answer, and that this jurisdictional result is itself the contribution. The eliminative arc does not produce a hole that a positive ontology of awareness would later fill. It produces a structural observation about where ontology's reach ends with respect to the field-question. The restraint is therefore not a holding back; it is the form the result takes.

What this means for the dialectic with each opponent is that the chapters are not partial demolitions awaiting a final reconstruction. Each is a complete piece of work at the level it operates. The arc's cumulative result is not a synthesis of the chapters into a single positive thesis about what awareness is; it is the marking of the structural feature each chapter has independently shown: that awareness, considered as the field of intelligibility, is not subject to ontological articulation as one of the items ontological articulation tracks. The contribution is methodological in this sense — it produces a disciplined philosophical result about the relation of awareness to ontology — and it is positive in this sense — it states something definite about that relation, not merely that one or another candidate location fails.

0.4 The shape of the argument

The book has eleven units, including this Introduction and the closing Conclusion. The nine chapters between them follow a sustained argumentative arc.

Chapters 1 and 2 lay the methodological ground. Chapter 1 introduces the basic distinction between awareness and determinate content, and develops the mislocation diagnosis at a level abstract enough to apply across the contemporary debates the later chapters will engage. Chapter 2 develops the local referent constraint — the disciplined diagnostic device the book uses to test whether the load-bearing terms in any ontological dependence claim about awareness have enough determinacy to support the claim. The constraint is not a general theory of meaning, not a verificationist criterion, not a semantic replacement for any contemporary work in philosophy of language; it is a local diagnostic, narrowly specified, used for the specific job of testing dependence claims about awareness. Chapter 2 makes that scope explicit.

Chapter 3 is the central structural argument, the *Bridge against ontological derivation*. The chapter states the rival thesis — that awareness is ontologically derivative from some prior item — and presses on the move from determinate items to the dependence relation the rival proposes. The pressure-point: any specification of the items would already operate within the field of intelligibility within which the items, the dependence, and the awareness being derived are themselves intelligible. The chapter includes a sustained engagement with the Stroud-style transcendental objection, which is the most serious worry the book's argument faces in this neighbourhood. The reply distinguishes the book's argument from a standard Kantian transcendental argument from conditions of knowledge to conditions of being; the book's argument concerns the intelligibility of ontological dependence claims themselves, not the general move from epistemic to ontological priority.

Chapters 4 and 5 take up the two principal contemporary reframings of the problem of consciousness — identification with content-architectural relations, and illusionist or meta-problem reframing of phenomenal consciousness. Chapter 4 engages Carruthers's higher-order theory and Block's distinction between access and phenomenal consciousness, arguing that identification claims of either kind locate awareness within a determinate content-architectural feature, and that the feature, whatever its content-architectural detail, cannot be where the field has its standing. Chapter 5 engages Frankish's illusionism and Chalmers's meta-problem of consciousness, arguing that an account of why we have the

intuitions and reports we have about consciousness leaves untouched the question of the field within which the intuitions, the reports, and the explanations are themselves intelligible.

Chapter 6 develops a different kind of subordination claim. Time, on Kant's celebrated doctrine, is the form of inner sense — the transcendental condition under which the manifold of inner intuition is necessarily ordered. The chapter grants Kant's account within its proper jurisdiction (the form of organised inner appearance) and presses on whether this jurisdiction extends to a condition governing the field within which any organisation of inner appearance is itself intelligible. The argument is structural rather than anti-Kantian: time conditions content-experience; the field is not what such conditioning specifies.

Chapter 7 takes up perspectival plurality. Other people are real. Their perspectives are not modifications of mine. Embodiment, sensory access, memory, expressive resources, and intersubjective correction richly individuate perspective-loci within experience. The chapter grants all of this. What it resists is the inference from the plurality of perspective-loci to the plurality of awareness-fields. The most articulate phenomenological accounts of perspective and self-awareness — Husserl on intentionality, Sartre on the pre-reflective cogito, Zahavi on the minimal self — each articulate perspectival plurality at the level of perspective-loci within a field, not at the level of fields themselves. The chapter draws on each, including Husserl's monadological intersubjectivity in the Fifth Meditation and Zahavi's multidimensional account of intersubjectivity, to mark that no constructible principle individuates fields by counting perspective-loci.

Chapter 8 consolidates the book's positive thesis about the relation of awareness to ontology. The chapter is structured around the engagement with Philip Goff's *Consciousness and Fundamental Reality* (2017), which is the most articulate contemporary defence of an ontology that takes consciousness as fundamental. Goff is the closest contemporary ally on the negative phase of the book's argument: he agrees that consciousness cannot be derived from non-conscious matter, that the standard physicalist identifications fail, and that illusionism does not deliver what an account of consciousness needs to deliver. The book and Goff diverge at the constructive step. Goff places consciousness as the intrinsic nature of the cosmos within an ontological architecture; the book argues that the field-question concerns the standing within which any such architecture is intelligible as architecture, and that the field is not within ontology's jurisdiction as an item to be catalogued. The chapter develops the jurisdictional thesis directly.

Chapter 9 is the closest classical precedent. Śaṅkara's Advaita Vedānta has articulated the structural position the book defends, in a different idiom and within different doctrinal commitments, for over a millennium. The chapter marks four specific points of structural kinship — the negation method (*neti neti*), the witness-doctrine (*sākṣin*), the resistance to identifying consciousness with the inner organ (*antaḥkaraṇa*), and the self-luminosity of consciousness (*svayaṃprakāśa*) — and six specific points of doctrinal distance: Brahman as metaphysical absolute, sat-cit-ānanda as positive characterisation, Sruti as scriptural warrant, mokṣa as soteriological purpose, māyā as the ontological status of the empirical world, and saṃsāra as framing cosmology. The book does not take on any of the six points of distance. The kinship is structural; the doctrines remain Advaita's.

The Conclusion restates the thesis in modest form, identifies open problems, and points to future work the eliminative arc suggests. It is deliberately brief.

0.5 What the book is not

The book's argument is structural and procedural, not constructive, and the reader is owed an explicit statement of what the book does not claim. The negations matter because the cumulative effect of the eliminative arc may otherwise be misread as a series of disguised positive metaphysical commitments.

The book is not idealism. It does not claim that the world is constituted by mind, that the items and structures the world contains are constructions of a mind or community of minds, or that the existence of stars, mountains, organisms, or the events in which they participate depends on any subject's awareness of them. Idealism is one specific metaphysical thesis within ontology; the book's argument is jurisdictional, and the argument that the field is not one of the items ontology catalogues is not a claim about the items' dependence on the field. The point can be made sharper. The book is not claiming that the items in any catalogue depend on the field for their existence; it is claiming that any catalogue of items is articulated within the standing the book has been calling the field. To say that a catalogue of stars is articulated within the field is not to say that whatever reality the catalogue concerns depends on a subject's awareness for its existence. The two claims are at different levels, and only the first is what idealism claims.

The book is not dualism. It does not claim that mind and matter are two fundamental substances, that consciousness is a non-physical thing standing alongside the physical, or that any mechanism of interaction needs to be specified between two kinds of being. Dualism makes a substantive count of fundamental kinds; the book makes no such count.

The book is not monism in the positive sense. It does not claim that all reality is one substance, that all distinctions are merely apparent, or that the manifest plurality of items is illusion overlying a deeper unity. The arguments against pluralising awareness in Chapter 7 may sound like an assertion of unity at the level of awareness; they are not. Both pluralisation and counting-to-one operate within the inventory the book's structural argument is about, and the book takes neither side within the inventory.

The book is not panpsychism. It does not claim that consciousness is ubiquitous in matter, that fundamental physical entities have phenomenal properties, or that the deep nature of the physical is experiential. Panpsychism makes a specific positive thesis about the distribution of conscious properties across an ontology of fundamental items; the book takes no position on that distribution.

The book is not cosmopsychism. It does not claim that the cosmos is a conscious subject, that the universe has a unified mental life of which our individual minds are aspects, or that there is one cosmic awareness within which all individual awarenesses are subsumed. Chapter 8 makes this negation explicit in its engagement with Goff. The book's argument is not a cosmopsychist argument with the cosmic subject given a different name.

The book is not mysticism. It does not claim that awareness is ineffable, that the field is beyond reasoned discourse, or that what the book has been pointing toward is something the methods of analytic philosophy cannot reach. The argument throughout is analytic, structural, and inferential. The book's restraint about positive ontological characterisation of awareness is not a mystical reticence; it is the structural observation that ontological characterisation is not the level at which the question of awareness as field is to be answered. The difference matters. In the relevant contrast, mysticism would take the object of contemplation to exceed reasoned discourse; the book takes the field-question to be answerable by reasoned discourse of a particular structural-jurisdictional kind. The chapters that follow are instances of that discipline at work. The book's restraint about positive ontological characterisation is not a refusal to engage; it is the form the engagement takes.

The book is not anti-science. It does not deny that the sciences deliver genuine knowledge about reality, that empirical methods are reliable within their domains, or that the items the sciences catalogue exist. The book's argument operates at a level adjacent to all such claims. The sciences succeed within their proper jurisdictions; the book's argument concerns a structural feature of the field within which any such inquiry is itself intelligible as inquiry.

Nor is the book realism as a positive doctrine. It is compatible with realism about ordinary and scientific objects, but it does not make realism its thesis. Realism is a position within ontology about the status of items in the inventory; the book's claim concerns the standing within which any such inventory is intelligible.

The book is not anti-logic. It does not deny the validity of inference, the canons of analytic argumentation, or the procedural norms of careful philosophical reasoning. The book's argument is itself an exercise in those norms throughout.

The book is not verificationism. The local referent constraint developed in Chapter 2 is not a general criterion of meaningfulness, not a denial that terms without inspectable referents can be meaningful, and not a semantic replacement for any contemporary work on reference or meaning. It is a local diagnostic, narrowly specified, used only to test whether ontological dependence claims about awareness have enough referential determinacy to do the work the claims assign their terms. Chapter 2 makes the scope of the constraint explicit.

The negations are not, taken together, a positive metaphysical position. They are the negative side of a jurisdictional argument: the book has been arguing, throughout, that ontology has a jurisdiction within which its first-order disputes about idealism, dualism, monism, panpsychism, and the rest are proper. The field-question is at the jurisdictional edge. Refusing to position the book within any of the standard first-order disputes is part of the discipline; it is not evasiveness.

0.6 How to read the chapters

Each chapter has the same dialectical shape. The opponent's argument is stated fairly, in its strongest form, from the primary text. The load-bearing premise of the argument is identified and isolated. The structural pressure-point is located at a level adjacent to the opponent's apparatus, not within the opponent's terms. The chapter then states, in a parallel closing section, what the structural argument establishes and what it does not.

The dialectical posture is consistent across the chapters. The opponent is granted what the opponent achieves within their own programme. None of the chapters argues that Carruthers, Block, Frankish, Chalmers, Kant, Husserl, Sartre, Zahavi, or Goff has failed at what they were trying to do; each is granted whatever it establishes within its own programme. The opponent's apparatus is granted, sometimes generously; the chapter's pressure-point operates at the level the opponent's apparatus does not address.

This is what the reader should expect: repeated narrow arguments, not a revealed system. Each chapter does its own work. Each grants its opponent's local legitimacy, tests whether the opponent's load-bearing terms can do the work the opponent assigns them, and states the limit precisely. The cumulative effect of the chapters — the eliminative arc — is what matters; no individual chapter delivers more than its narrow procedural result. The book's positive thesis emerges from the arc rather than from any individual chapter.

A reader looking for the book's positive ontological account will not find one. There is no chapter in which the book proposes what awareness positively is in metaphysical terms — no chapter in which

the field is given a substantive characterisation, identified with a metaphysical absolute, located within the inventory of fundamental individuals, or attributed the kind of property that would allow it to enter the standard first-order ontological disputes. The book's restraint about positive ontological characterisation of awareness is itself substantive: it is the chapter-by-chapter consequence of the structural arguments the eliminative arc develops. Chapter 8 makes the restraint explicit; the other chapters develop it.

The classical and phenomenological precedents are engaged in Chapter 7 (Husserl, Sartre, Zahavi) and Chapter 9 (Śaṅkara). Neither chapter claims that the precedent contains the book's argument in disguise. Each chapter marks the structural kinship between the precedent and the book's position while preserving the precedent's own commitments. The book does not extend Husserl. The book does not recover what Advaita meant. The book does not secularise any classical tradition. The chapters use the precedents to articulate the structural shape of the position the book defends, without taking on the precedents' doctrinal or scriptural commitments.

Two practical notes for the reader. First, technical terms are introduced where they earn their place. The book does not run a public-facing system, does not announce its operations through labels, and does not develop a private vocabulary. The terms that recur — "field," "awareness as field," "content-architecture," "the local referent constraint," "mislocation" — are introduced once and operated with thereafter. Where a chapter engages a specific philosopher's apparatus (Carruthers's higher-order theory, Husserl's noetic-noematic correlation, Goff's grounding by subsumption), the apparatus is reconstructed from the primary text and used within the chapter; it is not adopted as part of the book's own vocabulary.

Second, the citations are deliberately narrow. Each contemporary engagement rests on a specific primary text or texts, consulted directly, with the reconstruction memos that preceded the chapter drafts available as a record of the source-discipline. The book does not cite ornamentally. Where a figure is named, it is because the figure does specific argumentative work in the chapter at hand; the book does not draw on figures it has not consulted, and it marks explicitly where its engagement with a tradition is structural rather than primary-text.

The argument begins in Chapter 1, with the basic distinction between awareness and determinate content. The disciplined eliminative arc develops from there.

Chapter 1 — Awareness and the Mislocation Problem

1.1 The problem

The problem is not that theories of consciousness lack sophistication.

The complaint is not about historical timing or the slowness of the field. It is about logical placement. By the time a theory of consciousness has set out what it is to explain — pinned down its explanandum, decided what kind of thing it is, fixed the form of the question — the explanandum has been described in a vocabulary that takes its own terms for granted. Awareness, on which the description's intelligibility depends, has been quietly stationed inside the inventory of items the description is meant to handle. The work that follows is then earnest, often careful, sometimes spectacular. What it cannot reach, from where it has already begun, is the question of what was placed where.

The pattern recurs across otherwise different theories. A functionalist begins with mental states characterised by causal role. A higher-order theorist begins with first-order states made conscious by being represented from above. A workspace theorist begins with information made globally available across cognitive subsystems. An integrated-information theorist begins with the intrinsic causal structure of a physical system. A representationalist begins with content directed at objects in the world. A neuroscientist studying the correlates of consciousness begins with reportable states and their neural signatures. In each case, the theory takes itself to be saying what consciousness is by saying what it is made of, what it is constituted by, or what its structural signature is. In each case, the specification proceeds by picking consciousness out as one item among the items the inquiry already knows how to describe.

This is not, considered in itself, a defect. Inquiry has to begin somewhere; describing items is what inquiry does. The pressure-point lies one step earlier. Before a theory can place consciousness among the items it describes, consciousness has to have been picked out as the kind of thing that can be so placed. That step is rarely examined. It is performed at the threshold of inquiry and, once performed, ceases to be visible. The argument of this book is concerned with that threshold step. The argument is that the threshold step, in the dominant theoretical traditions, picks out something different from what the inquiry takes itself to be addressing — that the picking-out, however carefully done, places what conditions theoretical specification inside the inventory the specification means to handle.

The word "consciousness" carries a great deal of theoretical history. To avoid prejudging that history, this book prefers the word "awareness." The preference is not a definitional manoeuvre. It is a request for a slower beginning. "Awareness," in what follows, names what is operative when anything appears as anything — when a word appears as a word, when a sound appears as a sound, when a thought appears as a thought. The name is not a theory of what is so operative. It is a placeholder for the dimension that the book will argue cannot be located among the items that appear within it. Other vocabularies

have done analogous work. Husserl spoke of intentional consciousness in a specific technical sense. Phenomenologists working in his wake have spoken of the field of appearance, of disclosedness, of the horizon of givenness. Classical Indian traditions have spoken of witness-consciousness, of pure cognizance, of the *sākṣin*. Contemporary writers in philosophy of mind sometimes speak of phenomenal character, of *what-it-is-likeness*, of the experiential. Each vocabulary brings its own resources and its own risks. The book uses "awareness" and tries to mean by it nothing more than this: the dimension within which the items of an inventory are present as items.

The chapter that follows is concerned to show, in modest steps, what is meant by saying that this dimension cannot be placed inside the inventory it conditions, and what kind of philosophical claim the book is and is not making in saying so. The first step is to make the relevant distinction sufficiently concrete that subsequent arguments can be evaluated against actual cases rather than against rhetorical gestures.

1.2 What is at the threshold

Take a simple case. A reader, encountering a printed page, picks out the letters, the words, the sentences. Each is something one can point at, comment on, miscopy, lose track of, return to. Each is determinate: it has features one can describe, it occupies a position, it stands in identifiable relations to its neighbours. There is no difficulty in saying what these things are. The difficulty, if any, lies in saying it better.

Now consider, alongside the words on the page, the reading of them. Not a particular sentence as read, but the activity of reading taken as the present occurrence in which the words are appearing. The reader who looks for this finds something curious. No additional item enters the inventory. No further word, image, or object presents itself alongside the ones already there. Yet it is plainly not the case that nothing is present beyond the words. The words' being read is precisely what makes them more than ink. Their appearing as words rather than as marks is a feature of the present situation. It is just not a feature catalogued the way the words are catalogued.

The same structure appears in other modalities. A listener follows a piece of music. Notes, phrases, rhythms, harmonic progressions — each can be picked out, described, contrasted with neighbours. The listener can attend to the timbre of the cello, the rest before the entrance of the strings, the modulation in the second movement. Each of these is an item in the inventory of the listening. But the listening itself — the standing in which any of these is heard as something rather than as nothing — does not show up as one more note, one more phrase, one more harmonic feature. It is not catalogued. It is what makes the cataloguing a cataloguing of music rather than a registration of acoustic data.

A third case. A reader feels an ache in the shoulder, recognises a familiar tension, notices the pressure as a kind of fatigue. The ache is locatable, characterisable, comparable to past instances; it is an item in the inventory of what is presently going on. The recognising and the noticing, however, do not show up alongside the ache as further items. The reader who looks for them finds the ache, perhaps more vivid for being attended to, but no second occurrence that could be picked out as the having of it.

These three cases share a structure. In each, there is what is present — words on the page, notes in the air, sensation in the shoulder — and there is the standing in which the presence takes the form of presence. The first is determinate; the second is not determinate in the same way. The book calls the first contents and the second awareness, but the choice of vocabulary is provisional. What matters is the structural distinction. Contents are what can be placed in an inventory: pointed at, described, contrasted, theorised. Awareness is what makes the placing possible. It is not absent when contents are absent, and it is not added when contents are added; it is the dimension within which contents are

determinate as the contents they are.

This is not yet a metaphysical thesis. To distinguish contents from the standing in which they are present is not to say that contents are unreal, that they exist only as ideas, that the world reduces to mind, or that awareness has some hidden positive nature waiting to be revealed by further inquiry. The distinction is more modest. It is a difference in role. Contents play the role of items within investigation; awareness plays the role within which items can be picked out as items. The chapter will argue that confusing these roles — treating awareness as one more content — is what produces a recurring family of difficulties in the philosophy of mind.

The distinction is not new. Family resemblances appear in Husserl's account of act and content, Sartre's pre-reflective *cogito*, Zahavi's work on minimal selfhood, Strawson's treatment of the experiential, Advaita discussions of the *sākṣin*, and Yogācāra analyses of cognition. These are not engagements yet; the relevant historical and textual work is deferred to the later chapters. The book does not claim to originate the distinction. It claims that the distinction can be defended, against the contemporary moves that press most strongly against it, by a disciplined eliminative argument that does not depend on the broader metaphysical, soteriological, or contemplative commitments of those traditions. That is the book's contribution.

1.3 The shape of the mislocation

When a theory of consciousness picks out its explanandum, it picks it out as a content. The specific content varies. A higher-order theory picks out the relation between a mental state and a representation of that state. A workspace theory picks out the configuration of information made available across cognitive subsystems. A representationalist theory picks out the structure of content directed at objects. An integrated-information theory picks out the intrinsic causal architecture of certain systems. A biological-naturalist account picks out the higher-level features of neurobiological organisation. In each case, the picked-out item is determinate: it has features one can describe, an internal structure one can characterise, identification conditions one can in principle apply. The theory then offers an account of how the item arises, what it correlates with, why it has the features it has.

The trouble is not with this procedure considered in itself. Picking out an item is what makes investigation possible, and many of the descriptions produced by these theories are substantial empirical and theoretical achievements. The trouble is with the assumption that the procedure can be applied to awareness without remainder — that what shows up when one picks out awareness as the explanandum is what awareness is. The argument of this book is that this assumption mislocates awareness from the start. What gets picked out is some determinate item the theory takes consciousness to consist in. What gives the picking-out its character as a picking-out — the standing in which contents are determinate as contents — is not itself picked out. It does the picking-out.

The mislocation is not the result of theoretical sloppiness. It is the predictable outcome of treating a structural distinction as though it were a distinction between two items at the same level. To pick something out is to distinguish it from other items. For most things, this distinguishing is unproblematic. The reader can distinguish the page from the desk, the desk from the floor, the floor from the room, and none of these acts of distinguishing requires counting itself among the items distinguished. Awareness is not one of the items the distinguishing distinguishes itself from. It is operative in the distinguishing. To pick it out as an item requires giving it features that distinguish it from other items, and any such features will be features at the level of items — that is, at the level of contents. The result is that the picking-out has produced a content, and the dimension that gave the picking-out its character has not been picked out.

A specific instance will help. Consider, schematically, the higher-order family of views often associated with Rosenthal, Carruthers, and others. The detailed treatment of these views is deferred to Chapter 4; the present point concerns only the form of the proposal. The view holds, roughly, that a mental state is conscious in virtue of being represented by a further mental state of an appropriate kind. A first-order perceptual state is conscious because there is a higher-order thought, or a disposition to form such a thought, that takes the first-order state as its object. The relation between the first-order state and the higher-order state is what makes the first-order state conscious. Set aside, for the moment, internal debates between actualists and dispositionalists, between thought-based and perception-based higher-order accounts. Consider only the form of the proposal.

The proposal picks out consciousness as a relational property: a relation between two mental states. Both relata are determinate items; the relation is a determinate structural feature. Each side of the relation can be described, contrasted with neighbours, modelled. Once picked out in this way, consciousness becomes the kind of thing one can have a theory of: one can ask how the higher-order state arises, what its content is, how it is implemented, why it confers consciousness on the first-order state. The theory is then in a position to make claims, test them against cases, develop refinements.

What the theory has not done — and what the form of the proposal does not allow it to do — is address the dimension within which the relation between first-order and higher-order states is present at all. The relation is a content-relation. It has features describable from outside it. It is the kind of structure one can think about. Its being thinkable, its being available to inquiry, its being intelligible as the relation it is, is not itself a content-relation. It is the standing in which the relation is determinate as a relation. Higher-order theory is not equipped to address this standing because its categories — first-order states, higher-order states, the representational relation between them — are all categories at the level of contents. Whatever theory it produces will be a theory of contents and their relations. It will not be a theory of the dimension within which contents and their relations have the standing they have.

The same observation applies, with the appropriate adjustments, to the other theories named earlier. Workspace theory describes a configuration of information-availability across subsystems; the configuration is a content. Representationalism describes the structure of intentional directedness; the structure is a content. Integrated-information theory describes the intrinsic causal architecture of certain systems; the architecture is a content. None of these descriptions, however refined, addresses the dimension within which the descriptions are intelligible as descriptions.

It is important to be clear about what this observation is and is not. It is not the claim that the theories in question have failed by their own standards. By their own standards, they may succeed magnificently. Higher-order theory, in the hands of its more careful proponents, has produced rich and detailed accounts of self-monitoring, of metacognition, of the structure of reflective access. Workspace theory has produced detailed accounts of the dynamics of information broadcasting and its neural correlates. Integrated-information theory has produced a sophisticated mathematical formalism. None of these achievements is contested by the present argument. What the present argument contests is the further claim — sometimes made by the theories themselves, sometimes attached to them by philosophical commentators — that the descriptions on offer specify what awareness is. That further claim, the argument will press, depends on a mislocation of awareness from the start.

1.4 Why the diagnosis is structural

A natural response to the foregoing is that the mislocation problem, as described, is a placeholder for a more familiar worry. Perhaps the existing theories are simply incomplete. Perhaps the right level of description has not yet been found. Perhaps some empirical or theoretical advance is needed before the

question of what awareness is can be properly addressed. If so, the response is patience: the sciences move slowly, the right account will eventually emerge, and the difficulty will dissolve.

The chapter wants to mark clearly that this is not the worry being raised. The argument is not a complaint about the state of the empirical art. It is a structural observation about how the explanandum has been set up. No empirical advance, by itself, alters the setup. A more refined theory of higher-order representation, a more precise account of global broadcasting, a richer formalism for integrated information — each may produce real and useful results, but none of them addresses the question of whether what was being picked out at the threshold was the right kind of thing.

The distinction matters because it determines the appropriate response. If the worry were a routine demand for theoretical refinement, the response would be theoretical refinement. The argument would have nothing to add to ongoing scientific work; it would merely be a slow voice on the sidelines, hoping for progress. If the worry is structural, the response is different. The argument can engage with current theory, not by competing with it on its own terrain, but by identifying the threshold move whose consequences propagate into the difficulties the field has come to call hard.

A useful comparison is with philosophical critiques that target specific theoretical claims. A critic might argue that higher-order theory is incorrect because it fails to handle a particular kind of case — non-introspective consciousness, for example, or consciousness in non-human animals. Such a critique operates at the level of the theory: it accepts the theory's terms and presses against its application. A different critic might argue that higher-order theory is unfalsifiable in some respect, or that its formalism is underdetermined. These too operate at the level of the theory. The argument of this book operates at a different level. It does not deny that higher-order theory handles its target cases. It denies that the target cases are correctly characterised as cases of awareness in the relevant sense. The denial is not a claim about whether the theory's internal moves are correct. It is a claim about whether the theory's explanandum is what the theory takes it to be.

This is what makes the chapter's argument compatible with empirical work. Neuroscience can continue to identify which brain structures correlate with which reportable conscious states, and the argument does not constrain this. Cognitive science can continue to model the dynamics of attention, working memory, perceptual integration, and the broadcasting of information across cognitive subsystems, and the argument does not constrain this. Psychology can continue to study reports, judgements, intuitions, and behavioural correlates of consciousness, and the argument does not constrain this. The argument concerns one specific philosophical reading sometimes attached to those programmes: the reading that takes their items to specify what consciousness is, or to derive consciousness from something more basic, or to identify consciousness with a structural feature within the items they describe. That reading, the argument will press, depends on a mislocation. The reading is contested; the empirical work is not.

The compatibility is not a diplomatic concession. It is a consequence of the structural character of the argument. A structural argument about the form of the explanandum has no quarrel with first-order results. It has a quarrel with the philosophical interpretation under which those results are taken to bear on the question they are taken to bear on. Once the interpretation is corrected — once the items that the empirical and cognitive sciences describe are recognised as items within awareness rather than as candidates for what awareness consists in — the empirical programme continues unaltered. What changes is what the programme is taken to have established. It has established a great deal about content-organisation, about correlation, about the dynamics of cognitive processing. It has not established a specification of awareness as a content-level item, because the form of its inquiry was not equipped to do so.

1.5 What follows

The mislocation problem, named in this chapter as a recurring pattern in theoretical inquiry, does not yet do philosophical work on its own. Naming a pattern is not the same as showing that any particular claim instantiates it, and showing that a particular claim instantiates it requires a discipline that the next chapter is concerned to introduce. The discipline is referential and local. It asks of an ontological dependence claim whether its load-bearing terms have the determinacy required for the relation they assert to do argumentative work. That is the entire content of the discipline. It is not a general semantics, not a verification principle, not a replacement for any of the going theories of reference or meaning. It is a constraint that any philosophical claim involving an ontological dependence relation must already meet if it is to function as such a claim. The next chapter introduces it carefully, distinguishes it from neighbouring positions with which it is liable to be confused, and shows what it is and is not entitled to do.

With that discipline in hand, the third chapter applies it to the most direct way of mislocating awareness: by claiming that awareness is ontologically derivative from some prior item. The result, argued there, is narrow and specific. It is not that no reality exists independently of awareness, and not that the empirical sciences have overreached. It is that no ontological derivation of awareness can be stated without presupposing the field within which the derivation's terms become intelligible. The same form of argument extends, in the chapters that follow, to identification claims (Chapter 4), to illusionist and meta-problem responses (Chapter 5), to temporal subordination (Chapter 6), and to the multiplication of awareness into countable subjects (Chapter 7). Each of those chapters takes a specific contemporary proposal, states it as forcefully as its proponents state it, and asks whether its load-bearing terms can do the work the proposal requires. Each is a separate argument.

What the book aims to leave standing, at the end of these arguments, is not a new metaphysical posit but a clarified order of dependence. Ontological description remains possible. Scientific and philosophical inquiry into the items of the world remain possible. The standing of the field's empirical results is not in question. What is no longer possible, if the arguments succeed, is to treat awareness as one of the items that such description and inquiry place within their inventories. That is the structural conclusion the book is concerned to defend, and the chapters to come are the defence.

Chapter 2 — The Local Referent Constraint

2.1 What the next chapters need

The arguments that follow this one require a discipline for evaluating one specific kind of philosophical claim: an ontological dependence claim about awareness. The kind of claim is familiar. Awareness depends on neural activity. Awareness is constituted by a higher-order representational relation. Awareness is identical with integrated information. Awareness is grounded in functional organisation. Awareness is what biological systems of a certain kind produce. Each of these claims places awareness on the dependent side of an asymmetric relation; each takes some other item — neural activity, a representational relation, integrated information, functional organisation, biological systems — to be the prior side. The chapters to come press on the dependence relation in each case and ask whether the proposed asymmetry can be stated coherently.

To press on the dependence relation, the chapters need a constraint. Without one, an ontological dependence claim can be made by putting words in the right grammatical order: *awareness depends on X*. The grammatical operation produces a sentence; it does not, by itself, produce a claim that has done what an ontological dependence claim has to do. Such a claim has to mark an asymmetric relation between two relata in a way that lets the relation play its argumentative role. The relata have to be determinate enough for the relation between them to obtain or fail to obtain. The relation has to be determinate enough for the asymmetry it asserts to be marked rather than merely declared. If the load-bearing terms are not specified with enough determinacy for the relation to do its work, the claim has the appearance of being argued without yet being argued. The work of evaluation is then blocked, not because the claim is false, but because there is not yet anything determinate enough to be evaluated as true or false.

The constraint this chapter introduces does no more than register what such a claim has to do in order to function as a claim of the kind it advertises. It is not imposed from outside the relevant practice. It is the constraint that the practice itself requires for its terms to refer at the level of determinacy that an ontological dependence claim presupposes. Once stated explicitly, the constraint will look like a familiar feature of how dependence claims work in any domain. The chapter's task is to state it carefully, distinguish it from neighbouring positions with which it is liable to be confused, and show that it is doing only the local work the subsequent chapters need.

2.2 The constraint

The constraint can be stated in one sentence.

> A term used in an ontological dependence claim must have, in its load-bearing positions, enough determinacy of reference for the dependence relation it figures in to do argumentative work.

Each element of the formulation does specific work. "Load-bearing positions" are the positions on which the dependence claim's force depends — typically the names of the relata and the description of the relation. "Determinacy of reference" means that the term picks out something specific enough that questions about the relation it figures in have answers, even where those answers are difficult to establish. "Argumentative work" means that the relation can be marked in a way that permits the asymmetry being asserted to be recognised, contested, or defended.

Three features of the constraint are worth highlighting at the outset.

First, the constraint applies only to the load-bearing positions in an ontological dependence claim. It does not apply to all uses of all terms. A term may figure in many kinds of discourse — narrative, hypothetical, fictional, evocative — without facing this constraint. The constraint is local in this precise sense: it kicks in where a term is asked to bear the weight of an ontological dependence relation, and it does not kick in elsewhere.

Second, the constraint concerns determinacy, not method. It does not specify how the determinacy is to be achieved. It does not require that referents be directly observable, or theoretically tractable, or experimentally accessible, or anything else of that kind. The mode by which a term acquires its determinacy is not the constraint's concern. The constraint cares only that, in the relevant position in the relevant kind of claim, the determinacy is present.

Third, the constraint is internal to the practice of making ontological dependence claims. It is not imported from a theory of meaning or a theory of reference. The chapter does not commit to any particular such theory. What the constraint registers is that, whatever theory of reference one prefers, a dependence claim whose load-bearing terms fail to refer with the relevant determinacy does not yet function as a dependence claim. It is a sentence that uses dependence-claim grammar; it is not yet doing dependence-claim work.

2.3 Varieties of dependence

The chapters to come use "ontological dependence" as a cover term for several distinct relations that have been treated separately in contemporary metaphysics. It will help to mark the differences here, without taking a position on which of them is the correct account of any specific case.

Causal dependence is the relation a state stands in when it is brought about by some prior process or event. Brain activity causes a conscious experience; the experience is causally dependent on the activity. The relation is asymmetric in the standard direction: causes are prior to their effects, and the effect cannot be specified except as something the cause produces or sustains. Many proposals in the philosophy of mind treat the relation between brain and consciousness as causal in this sense, sometimes with further metaphysical commitments attached.

Constitutive dependence is the relation a composite stands in when its existence requires the existence and configuration of its constituents. A statue is constituted by the marble that makes it up. The relation is asymmetric in a different way from causal dependence: the constituents do not bring about the composite, but the composite cannot exist independently of them. Some philosophers of mind treat consciousness as constitutively dependent on the relevant brain states, sometimes invoking a "realisation" relation that is neither identity nor mere causation.

Grounding, in the contemporary metaphysical sense developed by Fine, Schaffer, Rosen, Audi, and others, is a relation in which one fact obtains in virtue of another. The phrase "in virtue of" marks the relation: the lower fact does not merely cause or constitute the higher; it is what the higher holds in

virtue of. Grounding has been deployed in recent work on the metaphysics of consciousness as a way of relating the physical and the mental that is neither causal nor identitarian.

Reductive dependence asserts that what is reducible is, in some specifiable sense, identical with or fully accounted for by what it reduces to. Reduction comes in several forms: type-type identity, token-token identity, theoretical reduction, functional reduction, structural reduction. In each, the reducing item is taken to fix the reduced item completely, leaving nothing further to explain.

Identity is the limiting case. To say that X is identical with Y is to assert that there is only one item, not two related items. Identity claims about consciousness — that consciousness is integrated information, that consciousness is a certain pattern of cortical activity, that consciousness is a higher-order representational relation — are not strictly dependence claims, but they raise the same question for the present constraint. An identity claim, no less than a dependence claim, requires its terms to be specified with enough determinacy for the identity relation to hold or fail.

These relations are not equivalent, and the present book does not need them to be. What it needs is that, in each case, the constraint of §2.2 applies. A causal dependence claim requires the cause and the effect to be specified determinately enough for the causal relation between them to be marked. A constitutive dependence claim requires the constituents and the composite to be specified determinately enough for the constitution relation to be marked. A grounding claim requires the lower and higher facts to be specified determinately enough for the in-virtue-of relation to be marked. A reductive claim requires the reducing and reduced terms to be specified determinately enough for the reduction to be evaluable. An identity claim requires both terms to be specified determinately enough for identity to obtain or fail.

The arguments of the chapters that follow do not turn on choosing between these forms of dependence. The argument is structural: whichever form the proponent of an awareness-derivation claim selects, the form requires determinate specification of its load-bearing terms, and the chapters press on whether such specification is available without already drawing on the field whose subordination is at issue. The constraint applies to all five forms in the same way. Distinguishing them at this stage clarifies what the arguments to come are not committed to: they are not committed to any particular theory of grounding, any particular account of constitution, any particular formal apparatus for reduction. They are committed only to the local constraint, applied across whichever form of dependence a specific opponent prefers.

2.4 What the constraint does not say

Because the constraint is liable to be confused with several stronger and more familiar positions, it is worth saying explicitly what it does not say. The list is selective; it covers the confusions most likely to arise.

The constraint does not say that a term is meaningless unless it has an observable referent. A term may carry meaning, be intelligibly used, and play a role in inquiry without facing the constraint at all. The constraint applies only to the load-bearing positions in an ontological dependence claim. Outside those positions, terms function under whatever conditions semantic theory and ordinary practice license. The constraint has no jurisdiction over them.

The constraint does not deny theoretical reference. A theoretical posit — an unobservable particle, a postulated mechanism, an abstract structure — may serve as the prior term in an ontological dependence claim, provided that its theoretical specification is determinate enough for the dependence relation to do its work. The constraint asks whether the specification is sufficient for the asymmetry being asserted; it does not require that the posit be observable, accessible, or otherwise epistemically

privileged. The theoretical practices of physics, biology, and the special sciences continue under their own conditions, and the constraint does not contest any of them.

The constraint does not deny modal discourse, mathematical reference, memory-based reference, inferential reference, or testimony-based reference. These remain available. What the constraint asks is whether, in the specific case of an ontological dependence claim about awareness, the relata's specifications support the asymmetric relation that is being asserted. If they do, the constraint is satisfied. If they do not, what has been said is not yet an ontological dependence claim in the operative sense; it is a sentence that uses the grammar of such a claim without yet doing what such a claim has to do.

The constraint is not verificationism. Verificationism, in its historically influential form, made a stronger and more general claim: that a sentence is cognitively meaningful only if it can be empirically verified. That claim faced two well-known difficulties. First, the claim itself is not empirically verifiable, which led to standard charges of self-defeat. Second, the claim is too restrictive to allow standard scientific theorising about unobservables, which led to a long line of internal revisions by logical empiricists and ultimately to the position's collapse. The constraint of this chapter does neither of these things. It does not make a general claim about cognitive meaning, and so it is not subject to self-application in the way verificationism was. It does not restrict scientific theorising about unobservables, and so it does not face verificationism's empirical difficulties. It is a local constraint internal to one specific kind of philosophical claim. The lessons of the failure of verificationism do not carry over to it.

The constraint is not a substitute for any contemporary theory of reference. Kripke's account of rigid designation, Putnam's account of natural-kind terms, Stalnaker's account of two-dimensional semantics, Williamson's account of vagueness and reference — each is a substantial theoretical contribution to the philosophy of language, and the chapter takes no position on which of them is correct. The constraint is compatible with several of them; it does not require any of them; and it does not propose itself as an alternative to them. The constraint operates at a level too local to compete with general theories of reference. It registers only what a dependence claim has to deliver in its load-bearing positions if the dependence relation is to do its work. Theories of reference may explain how such delivery is possible; the constraint asks only that, in the case at hand, it has occurred.

The constraint is not a hidden semantic thesis. It does not claim that being is to be intelligibly assertible, or that ontological status is to be referentially marked, or that to exist is to figure in successful reference. None of these is among the constraint's premises or its consequences. The constraint says only that an ontological dependence claim, considered as a dependence claim, requires determinate load-bearing reference. Whether the world contains items that are not so referred to, and whether such items have ontological standing of various kinds, are questions the constraint does not address.

2.5 Why the constraint is needed only locally

The chapter has emphasised that the constraint is local. It is worth saying why locality is enough.

The arguments of the chapters that follow are not arguments about awareness in general. They are not arguments about how the term "awareness" should be used in ordinary discourse, in literature, in introspective report, in clinical assessment, in contemplative writing, or in the casual exchanges of everyday life. They are arguments about a specific family of claims: claims that take awareness to stand in an ontological dependence relation with some other item, with awareness on the dependent side. To press on such claims requires nothing more than the constraint that such claims have to meet to be the kind of claim they are advertised as being.

It might seem, in light of this, that the constraint is suspiciously tailored to its target. A constraint introduced in order to defeat a particular family of claims may look like a piece of philosophical equipment manufactured for a purpose. The appearance is misleading. The constraint is not designed to defeat anything. It is the constraint any ontological dependence claim has to meet — whether the claim is about awareness, about persons, about social entities, about mathematical objects, about anything for which an asymmetric ontological relation is being asserted. To say that awareness depends on X is to assert an asymmetric ontological relation between awareness and X. To assert such a relation is to assert it under whatever conditions any such relation has to meet to be assertible. Whatever those conditions are, they apply equally to claims about awareness and to claims about anything else. The constraint is general in this respect.

Locality, then, is not narrowness of application but narrowness of ambition. The constraint does only what it does. It says nothing about meaning generally, about reference generally, about cognition generally, about being generally. It says one thing about one kind of claim. The arguments that follow take the constraint and apply it to one specific case after another: ontological derivation (Chapter 3), identification (Chapter 4), temporal subordination (Chapter 6), perspectival multiplication (Chapter 7). Each of these cases is a case in which an ontological dependence relation, or a structurally adjacent kind of subordination relation, is being asserted with awareness on the dependent side. In each case, the constraint asks whether the load-bearing terms support the relation being asserted. In each case, the chapter argues, they do not — not because of any general semantic difficulty, but because of a specific structural feature of awareness that makes the relevant determinacy unachievable for awareness in the role the subordination claim assigns it.

That argument cannot be made until the constraint is in hand. With the constraint stated, with its scope clarified, and with the most likely confusions removed from view, the first specific application is now possible. The next chapter takes up that application: the argument against ontological derivation.

2.6 How the constraint is met

To say that a dependence claim's load-bearing terms must have enough determinacy for the dependence relation to do its work is not to say that the determinacy must take a single form. Different kinds of terms achieve determinacy in different ways. An empirical posit may be specified by its measurable features, its experimental signatures, its predicted behaviour under specified conditions. A theoretical posit may be specified by its role in a theory, its inferential connections to other posits, its place in the explanatory structure of the relevant domain. A structural posit may be specified by its formal properties, its relations to other structural items, its representability in a formalism. Each route to determinacy is legitimate. The constraint does not privilege one over another. It asks only that, by some legitimate route, the relevant determinacy is in place.

The constraint also recognises that determinacy is a matter of degree. A term may be partially specified in some respects and not in others; a theory may have refined some of its central terms more than others; a research programme may be in mid-flight, with some of its commitments well-articulated and others still gestural. The constraint does not require complete determinacy in every respect. It requires determinacy sufficient for the specific argumentative use to which the term is being put. A term in the load-bearing position of an ontological dependence claim must be determinate enough for the dependence relation it figures in to be marked. It need not be determinate in respects that are irrelevant to that relation.

In practice, the constraint is met when one can answer, with reasonable specificity, two kinds of question about a dependence claim. First, what is being claimed to depend on what? That is, can one say,

with enough precision for the relation between them to be characterised, what each relatum is? Second, in what sense is the dependence asymmetric? That is, can one say, with enough precision for the asymmetry to be recognised, why the prior term is prior and the dependent term is dependent?

When both questions have specific answers, the claim has met the constraint and the argumentative work of evaluating its truth or falsity can begin. When one or both questions lacks a specific answer in the relevant respects, the claim has not yet met the constraint, and the argumentative work is blocked at the threshold. The blockage is not a verdict that the claim is false. It is a verdict that the claim has not yet been brought to the form in which truth-evaluation can proceed.

The next chapter applies this procedure to a specific case: the claim that awareness is ontologically derivative from some prior item. The application will not turn on any general semantic principle or on any verificationist commitment. It will turn on a specific question about whether, when a proponent of derivation specifies the prior item, the dependence relation, and the asymmetry between them, the specifications can support the work the claim is asking them to do. The chapter will argue that they cannot, but the argument will be made case by case, with the relevant specifications taken in their strongest form. That is what the local constraint, once in hand, makes possible.

Chapter 3 — The Bridge: Against Ontological Derivation

3.1 The rival thesis

The chapters preceding this one have done two things. They have introduced the working distinction between awareness as a field and what appears within it; and they have isolated a local referential discipline for evaluating ontological dependence claims. With that, the first central argument of the book can be stated. Its target is the most direct way of placing awareness inside the inventory of the world. The target view is the following.

> Awareness is ontologically derivative if there exists some X such that X is prior to awareness and awareness depends on X.

The view need not specify what X is. It may be neural activity, computation, biological organization, information, physical structure, or some other proposed prior. The argument that follows holds regardless of which X the proponent selects, because what it presses on is not the choice of X but the form of the dependence claim.

The view also need not specify what notion of priority is in play. Causal, constitutive, metaphysical, or grounding priority will each do, as long as the relation between X and awareness is asymmetric in the relevant way: X can be specified without invoking awareness, awareness cannot be specified without invoking X, and the disappearance of X would carry the disappearance of awareness while the converse does not hold. This is the ordinary shape of an ontological dependence claim, and it is the only structural commitment the argument needs the rival to share.

The argument then asks what such a claim must do in order to be intelligible — not in order to be known to be true, but in order to have determinate content at all. The conclusion will be narrow. It will not be that no reality exists independently of awareness. It will not be that everything is mental. It will be that no ontological derivation of awareness can be stated without presupposing the field in which the derivation's key terms become intelligible. The work of this chapter is to show that, and to show why the narrow conclusion is not a sleight-of-hand version of a stronger one.

3.2 The pressure on the load-bearing terms

To state the rival view at all, the proponent must specify three things determinately enough for the dependence relation to do argumentative work: the prior item X; the relation of priority or dependence; and the sense in which X is independent of awareness in a way awareness is not independent of X.

Consider each in turn.

The prior item X must be specifiable. If the proponent cannot say what X is — what kind of item, with what features, with what individuation conditions, in what relation to the items already accepted as real — the view is not yet a view. It is a slot waiting for content. So the proponent will say: X is neural activity in such-and-such cortical regions; or X is the integrated information of a system of such-and-such structure; or X is the computational implementation of such-and-such functional architecture; or X is the biological organization that brain systems of vertebrates exhibit. Each of these is a substantive proposal. Each delivers an X that has features one can describe, structure one can investigate, and identification conditions one can in principle apply.

The dependence relation must also be specifiable. The proponent must say what it is for awareness to depend on X rather than merely to be correlated with X or co-occur with X. The available options are familiar. Some take the relation to be causal: X causes the occurrence of awareness, in a sense made precise by some theory of causation. Some take it to be constitutive: awareness is constituted by X, in the way a statue is constituted by the marble that makes it up. Some take it to be grounding: X grounds awareness, in the contemporary metaphysical sense in which one fact obtains in virtue of another. Some take it to be reductive: awareness is identical with or reducible to a feature of X. In each case the relation is a specific philosophical posit with its own commitments, its own structural form, and its own consequences for how the dependent term inherits its features from the prior term.

The independence claim must also be specifiable. The proponent does not merely say that awareness depends on X; the proponent says that X is prior in a way that awareness is not. This asymmetry is the heart of derivation. Without it, the view is at best one of mutual ontological co-determination, which is not the thesis under examination. So the proponent owes some account of what the asymmetry consists in: that X could obtain without awareness obtaining, that X is specifiable without reference to awareness, that the laws or principles governing X make no essential reference to the awareness it produces or constitutes or grounds. Each of these is a substantive commitment.

The three commitments are distinguishable, and the rival typically needs all three. *Modal independence* is the claim that X could obtain without awareness obtaining: the counterfactual relation runs in one direction only. *Specification independence* is the claim that X is specifiable in terms that do not invoke awareness: the descriptive vocabulary stands on its own. *Explanatory independence* is the claim that the laws or principles governing X make no essential reference to awareness: the explanatory order in which X figures does not need to mention awareness to be complete.

These three labels are descriptive, not apparatus-introducing. The chapter does not use them as recurring terms; they are placeholders for what the rival has to defend if the asymmetry the derivation requires is to be marked. After this section the chapter returns to speaking of "the asymmetry," "the independence claim," and "what the rival owes," without rehearsing the labels.

Each is a different commitment, and each can in principle fail without the others failing. A view might hold that X is modally independent of awareness while quietly conceding that X is specifiable only in terms that already invoke awareness — in which case the modal claim survives but the specification claim does not. A view might hold that X is specifiable without reference to awareness while conceding that the laws governing X mention awareness at some level — in which case the specification claim survives but the explanatory claim does not. The chapter's structural pressure applies wherever any one of the three is offered as part of the asymmetry the derivation requires. It is enough for the chapter's argument that the rival's specification of any one of these independence claims be examined; the rival cannot retreat from a failed specification of one to an unspecified version of another without conceding that the unspecified version no longer marks the asymmetry the derivation needs.

The structural argument does not depend on the proponent's choice of dependence relation, either.

Whether the relation is causal, constitutive, grounding, or reductive, the chapter's pressure operates on the same point: the specification of X, the relation, and the asymmetry has to be done by someone, in some vocabulary. The four kinds of dependence relation differ in their internal structure and in the philosophical commitments they bring; they do not differ in whether their load-bearing terms must be specifiable. They must.

Together, the three commitments make up what the rival has to deliver. The proponent can defend each, can concede some while pressing the others, can offer combinations weighted in different ways. What the proponent cannot do is offer none of them: a dependence claim with no specifiable asymmetry between its terms is not yet an asymmetric ontological dependence claim. The rival's position requires the asymmetry; the asymmetry requires the specifications; the specifications are what the chapter's diagnosis presses on.

So far, the rival view has been stated charitably. The proponent has been credited with the resources philosophy of mind ordinarily uses. Nothing in the argument that follows depends on construing the view weakly.

3.3 The diagnosis

Now consider what the proponent has supplied. Each of the three load-bearing items — the prior X, the dependence relation, the independence claim — is articulated within a fabric of intelligible meaning. To specify X as neural activity is to draw on the conceptual resources of neuroscience: a vocabulary of cells and signals, of regions and dynamics, of measurements and models. To specify the dependence relation as causal or constitutive or grounding is to draw on the conceptual resources of metaphysics and the philosophy of science: a vocabulary of causes, of constituents, of in-virtue-of relations, of laws and counterfactuals. To specify the independence claim is to draw on the conceptual resources of ontological discourse: of priority, of asymmetry, of what holds without what.

These are not idle observations. The point is structural. Each of those vocabularies has the features it has because each is intelligibly deployed by a speaker, asserter, or thinker who is doing something with it. A neural specification is intelligible because the term "neural activity" has a determinate use within a practice of inquiry. A dependence relation is intelligible because the relation has features one can think about, distinguish from neighbouring relations, and apply in particular cases. An independence claim is intelligible because the asymmetry it asserts can be marked, contrasted with mutuality, and used to draw consequences. In each case, intelligibility is at work.

The argument now asks what is required for that intelligibility to be operative. The answer is not metaphysically tendentious. Intelligibility is not stored in marks on a page or in soundwaves in the air. It is a feature of meaning-bearing activity — of assertion, inference, identification, comparison, denial. The present point does not require deciding whether such activity can also be physically described; considered as meaning-bearing activity, it occurs within the field in which items, vocabularies, dependence relations, meanings, distinctions, and claims have whatever standing they have. To specify X is to be in that field. To specify the dependence relation is to be in that field. To specify the independence claim is to be in that field.

The rival's claim, then, has the following shape. Within the field of intelligibility — that is, within awareness, in the present book's sense — the proponent assembles three determinate items: a prior X, a dependence relation, and an independence claim. The proponent then asserts that these items, taken together, place awareness ontologically below X. But the field within which the items are determinate, the relation is intelligible, and the independence is assertible is the very field the assertion proposes to

subordinate.

This is the diagnosis. It is not the claim that the proponent has overlooked something obvious. It is not the claim that intelligibility entails being. It is the narrower observation that an ontological derivation of awareness, when it is stated with the determinacy required for the dependence relation to do work, articulates its terms inside the field whose subordination it is meant to demonstrate. The derivation does not reach over the field. It uses the field. Its operation presupposes what its content denies.

Two cases are worth distinguishing here, because they exhaust the proponent's options.

In the first case, the proponent supplies determinate content for X, the dependence relation, and the independence claim. The local referential discipline of the previous chapter is then satisfied. The terms have referents, even if those referents are theoretical, inferential, or model-mediated. What the proponent has produced is a within-field account of one part of the world: an account of how neural systems work, or how integrated information is structured, or how computational organization underlies certain phenomena. This is a real contribution to the description of contents. It is not yet a derivation of awareness, because none of its specifications has stepped outside the field whose subordination is at issue. The proponent has supplied an empirical or theoretical articulation of content. The proponent has not supplied an awareness-independent ranking.

What within-field articulation does accomplish is worth saying explicitly, since the chapter's argument is sometimes mistaken for a complaint against empirical or theoretical work. Within-field articulation can do real philosophical and scientific work. It can identify causal regularities among contents, support predictions that further inquiry can confirm or disconfirm, mark constitutive relations between higher-level and lower-level features of a given system, and trace grounding relations among facts. Each of these is genuine knowledge of the items the articulation picks out. The argument neither contests nor diminishes any of it. What the argument contests is the further claim that work of this kind, however refined, can mark the priority of any such item over awareness. The within-field standing of the articulation is also what makes it work; the intelligibility that lets it be confirmed, contested, and refined is what the further priority claim has to step outside in order to deliver what the priority claim is supposed to deliver. The articulation cannot do both at once.

In the second case, the proponent acknowledges this and replies that the dependence relation is meant to hold outside the field. It is meant to obtain whether or not anyone is intelligibly asserting it. It is meant to be an ontological fact independent of meaning-bearing activity. But then the local referential discipline is no longer satisfied. The terms no longer have the determinacy required for the dependence relation to do argumentative work. They remain words within the sentence, but the asymmetry the rival's view requires — between something prior and awareness as derivative — can no longer be marked in the form that view needs.

What withdrawal from intelligible specification looks like in practice is worth pausing on, because the move is structurally subtle. The proponent does not retract the claim. The grammar of the sentence is preserved. What changes is the proponent's account of what the sentence is doing. Some versions of the move appeal to future inquiry: the prior item, on this version, is whatever inquiry will eventually settle on as the relevant physical or computational base, and the claim is that awareness will turn out to depend on it. Other versions appeal to an unspecifiable ontological substratum: the dependence relation is real, even if we cannot articulate its terms with current resources. Other versions appeal to a stipulation: the proponent declares that the asymmetry holds, leaving the burden of articulation to others. In each version, the asymmetry the rival's view requires is no longer being marked. It is being promised — by future inquiry, by hypothetical resources, or by stipulation — but a promise is not a marked asymmetry. The dependence relation is meant to do argumentative work; a promise that

someone will eventually be able to do that work does not yet do it.

The two-case diagnosis is the chapter's structural result. It does not depend on the rival's having made a mistake. It depends only on the form a coherent ontological derivation has to take and on the recognition that the form, taken at the required level of determinacy, articulates within the field whose subordination it proposes.

Either the proponent stays within the field of intelligibility, in which case the proposed derivation has not achieved the external priority it claims; or the proponent withdraws the claim from intelligible specification, in which case it no longer functions as an argued derivation. This is the structure of the result.

3.4 The modest conclusion

The conclusion is best stated without inflation.

> The argument does not show that no reality exists independently of awareness. It shows that no ontological derivation of awareness can be stated without presupposing the field in which the derivation's terms become intelligible.

That is the whole of what the chapter establishes. It is consistent with the existence of an external world. It is consistent with the empirical success of neuroscience, cognitive science, and physics. It is consistent with the existence of mind-independent objects, properties, and processes. It is consistent with realism about a great many things. What it is not consistent with is the further claim that the items those disciplines describe could play the role of priors in an ontological derivation of awareness without their description's availability already drawing on the field they are proposed to subordinate.

It is worth saying explicitly what this leaves untouched.

The argument leaves untouched the empirical questions about which physical and biological structures correlate with which conscious states. Those are first-order scientific questions on which the argument has no jurisdiction. The argument leaves untouched the explanatory project of cognitive science, considered as the project of describing how contents are organised, processed, integrated, and used. The argument leaves untouched the metaphysical project of describing how items in the world stand in relations of causation, constitution, and grounding. What the argument restricts is one specific extension of those projects: the extension that takes their items to play a derivational role with respect to awareness itself. That extension is what the chapter has argued cannot be coherently stated.

3.5 The Stroud-style objection

A serious objection now arises. It is associated, in contemporary form, with Barry Stroud's work on transcendental arguments, and it has been developed in different directions by Robert Stern, Anil Gomes, Lucy Allais, and others who write on Kant and on the transcendental-argument tradition more broadly. The argument of this chapter must engage it directly, because if the objection is correct as stated, the chapter's result collapses into a much weaker one.

The objection runs as follows.

> Even if every claim-act presupposes awareness — even if every assertion of an ontological dependence claim occurs within a field of intelligibility — it does not follow that awareness is ontologically non-derivative. To draw that conclusion is to move illicitly from a condition of our asserting or knowing to a condition of being. Transcendental arguments of this form

> have, on standard analysis, the resources only to license claims about how we must think about the world, not claims about how the world is.

The objection is not trivial. It is the central worry that has shaped the analytic reception of transcendental argumentation since the mid-twentieth century, and the work of Stroud and his successors has rendered the worry sharp. The most that a transcendental argument of this form is supposed to license, on Stroud's analysis,[1] is the conclusion that we must believe or assert that p. It does not license the further conclusion that p obtains independently of our believing or asserting it. To take a familiar example: even if we must believe in the existence of an external world in order to engage in the practices we engage in, it does not follow that an external world exists. The strongest conclusion the argument-form supports is that we cannot help but believe in one.

If the argument of §§3.2–3.4 is a transcendental argument of that form, then it has the same limitation. The strongest conclusion it can support is that we cannot help but presuppose awareness whenever we attempt to derive it. It cannot support the further conclusion that awareness is, as a matter of how things are, ontologically non-derivative. Awareness might still be derivative, in the way the rival view claims; what the argument would have shown is only that we cannot articulate that derivation without using awareness. That is a less interesting result. It would amount to a constraint on philosophical discourse rather than to a constraint on what there is.

The objection is the most exposed pressure point in the chapter. The reply must take it seriously.

3.6 Reply

The reply is that the argument of this chapter is not, in fact, a transcendental argument of the form Stroud's analysis targets. It is structured differently, and the difference matters for what the conclusion can support.

A standard transcendental argument, on the analysis Stroud gives, has the following shape. Some practice — asserting, knowing, doubting, engaging in inquiry — is identified as one we cannot give up. The argument then claims that the practice has certain necessary conditions. The conclusion drawn is that these conditions must obtain. Stroud's well-known objection is that even granting the first two steps, the most the argument licenses is a claim about what we must take to obtain, not a claim about what obtains independently of our taking it to obtain. The illicit move is from the conditions of a practice to the conditions of the world.

The argument of this chapter is not of that form. It does not say: we must engage in the practice of asserting ontological dependence claims; therefore the conditions necessary for such assertion obtain. It says something narrower. It says: the ontological dependence claim under examination is the claim

[1] The objection has its original source in Barry Stroud, "Transcendental Arguments," *The Journal of Philosophy* 65 (1968): 241–256. For any candidate proposition S whose truth is offered as a necessary condition of an indispensable practice (language, thought, or experience), the sceptic "can always very plausibly insist that it is enough to make language possible if we believe that S is true, or if it looks for all the world as if it is, but that S needn't actually be true" (p. 255). The most a transcendental argument of this form can establish is that "we must believe that there are material objects and other minds if we are to be able to speak meaningfully at all" (p. 256). Closing the gap to a conclusion about how things are independently of our believing or appearing requires a verificationist or idealist supplement; if either supplement is granted, the sceptic is refuted directly and the transcendental argument is superfluous. The contemporary reconstruction is given by Robert Stern, *Transcendental Arguments and Scepticism* (Oxford: Clarendon, 2000), §2.3 "The Modal Objection," pp. 57–64. Stroud's own later treatment in *The Significance of Philosophical Scepticism* (Oxford: Clarendon, 1984), Chapter IV, pp. 128–169, presses a structurally adjacent worry on Kant's refutation of idealism — that the only "refuge" left open for the Kantian project is transcendental idealism, which at its own level is difficult to distinguish from the scepticism it was supposed to overcome (pp. 161–168). The Kantian refuge-in-idealism worry is not directly engaged in this section.

that awareness is derivative; that claim has a content; its content requires the determinate specification of an X, a dependence relation, and an independence asymmetry; the determinate specification of these items places the claim inside the field of intelligibility; therefore the claim, as a coherent ontological dependence claim, cannot exhibit awareness as ontologically prior in the relevant sense.

The argument is not about the conditions of our asserting things in general. It is about the internal structure of one specific kind of claim: an ontological derivation of awareness. The conclusion is not that we cannot help but think awareness is fundamental. The conclusion is that no proposed derivation of awareness can articulate its terms outside the field whose subordination is at issue. That is a claim about the proposed derivation, not a claim about the inevitability of a belief.

The contrast is worth taking slowly. A standard transcendental argument identifies a practice we cannot give up — thinking, knowing, having experience — and proceeds from the necessity of the practice to a claim about what its conditions imply about the world. The chapter's argument does not begin from the necessity of a practice. It begins from one specific kind of philosophical claim: the claim that awareness is ontologically derivative from some prior X. The chapter asks what such a claim must do to function as the claim it advertises itself as being. The conclusion — that the claim cannot be stated in a way that achieves the priority asymmetry it asserts — is a result about what is internal to the claim, not a result about what is external to all claim-making.

The two argument-forms can be displayed side by side without invoking any apparatus the chapter otherwise avoids. The standard transcendental argument moves from "we cannot give up practice P" to "P has necessary conditions C" to "therefore C obtains." The chapter's argument moves from "the rival's view is an ontological derivation of awareness from prior X" to "the derivation's content requires the determinate specification of X, the dependence relation, and the independence asymmetry" to "the specification places the claim inside the field of intelligibility" to "the claim cannot exhibit awareness as ontologically prior in the relevant sense." The first form moves from a practice's conditions to a conclusion about the world. The second moves from a claim's internal requirements to a conclusion about what the claim can establish. The difference is the difference between asking what the world must be like for our practices to be intelligible and asking what one specific philosophical thesis is internally committed to.

The argument does not infer from the necessity of awareness for our cognition to the metaphysical fundamentality of awareness. It argues that the *derivation claim itself* cannot meet its own burden without presupposing the field whose derivation it asserts. That is the exact inference the chapter declines to make, and it is the inference Stroud's analysis rightly targets when it appears in transcendental arguments.

The Stroud-style objection presupposes that the argument moves from the conditions of thinking awareness to the metaphysical status of awareness. It does not. It moves from the internal incoherence of stating an ontological derivation under a referential discipline to the inability of any such derivation to deliver the subordination its proponent intends. That is a structural result about derivation claims, not an epistemic-to-ontological inference.

To put the same point another way: the rival's claim that awareness is ontologically derivative is not a feature of the world available for inspection independently of its articulation. It is, considered as a philosophical thesis, an articulation. The argument addresses the articulation. The articulation either supplies determinate terms — in which case it stays within the field — or does not — in which case it fails the local referential discipline. Either disjunct blocks the derivation as a derivation. Neither disjunct rests on the move Stroud rightly criticises.

There is a related worry that should be addressed in the same neighbourhood. One might press: are you not at least relying on a premise of the form *to be is to be intelligibly assertible*, or *to be ontologically fundamental is to be presupposed by all intelligible assertion*? Either of those premises would be a substantive metaphysical thesis, and the chapter has not argued for it. The answer is that the chapter does not need either. The chapter is not arguing that awareness is fundamental because all assertion presupposes it. The chapter is arguing that any *proposed derivation* of awareness — any attempt to rank awareness ontologically below something more basic — cannot be stated in a way that achieves what it intends. That argument requires only the local discipline introduced in the previous chapter and the observation that ontological dependence claims have load-bearing terms whose determinacy is required for the relation to do work.

What the chapter does need are two things. It needs the local referential discipline introduced in the previous chapter: that the load-bearing terms in an ontological dependence claim must have enough determinacy of reference for the dependence relation they figure in to do argumentative work. And it needs the observation that the dependence claim against awareness, like any dependence claim, has load-bearing terms — a prior, a relation, an asymmetry. Both premises are weak. The local referential discipline is internal to what ontological dependence claims have to do to function as ontological dependence claims, not a substantive thesis about meaning or being. The observation about load-bearing terms is structural, not metaphysical. Neither is the verificationist premise the standard transcendental-argument literature has been criticised for requiring; neither is the idealist premise; neither is the premise that being is intelligible assertibility. The chapter rests on what dependence claims need to do, considered as dependence claims, and asks whether the rival's specific dependence claim can do it.

The reply does not pretend that the objection has no remaining force. There will be readers who insist that any argument operating in the neighbourhood of transcendental reasoning inherits the limitations Stroud identified. The right response to those readers is to invite them to identify the exact step at which the argument of §§3.2–3.4 moves from a condition of assertion to a condition of being. The chapter contends that no such step is taken. The conclusion is about the structure of a class of claims, not about the metaphysical implications of conditions of cognition. If the chapter is right that no such step is taken, the Stroud-style objection bears on a different argument than the one being made here.[2]

A reader pressing the objection might still say: the chapter has argued that the derivation cannot be stated without presupposing the field — that is itself a claim about conditions of stating, and a claim about conditions of stating is just the kind of claim Stroud's analysis limits to what we must take to be the case. The reply is precise. The chapter's claim is not about what the proponent must do in order to assert the derivation; it is about what the derivation, considered as a content-bearing claim, must

[2]Stern (2000, Introduction §0.2, pp. 10–11) classifies transcendental arguments into four kinds — truth-directed, belief-directed, experience-directed, and concept-directed — according to whether the necessary condition the argument identifies is a non-psychological fact, a belief, a feature of experience, or a context of concept-acquisition. The chapter's argument is not of any of these four kinds; it is not, in Stern's sense, a transcendental argument at all. The chapter does not identify a necessary condition of an indispensable practice and infer from that condition's necessity to a conclusion about how things are; it identifies the internal structural requirements of one specific kind of philosophical claim — an ontological derivation of awareness — and presses on whether such a claim can meet those requirements without already operating within the field whose subordination it proposes. The contemporary Kantian transcendental-argument programme, as developed by Stern (Chapters 3–6) and by Anil Gomes, "Perception and Reflection," *Philosophical Perspectives* 31, no. 1 (2017): 131–152, at §§3–4, has accepted the force of the modal objection against truth-directed arguments and has retreated to more modest forms; the chapter's structural argument is not in that programme and does not stand or fall with its prospects. Lucy Allais, *Manifest Reality: Kant's Idealism and his Realism* (Oxford: Oxford University Press, 2015), Chapter 1, defends a moderate metaphysical interpretation of Kant's transcendental idealism on which the empirical realism Kant secures is genuinely realist; this is noted here only to mark that the related Kantian refuge-in-idealism worry presupposes a reading of transcendental idealism that contemporary Kant scholarship contests, and not as a commitment of the present chapter.

internally satisfy. The move is from the internal content-requirements of the rival's claim to what that claim can and cannot establish. It is not from the external conditions of assertion to what is the case independently of assertion. The proponent who can produce a determinate ontological derivation that exhibits awareness as ontologically prior to its specifications has not been refuted by anything in this chapter. The chapter argues only that no such derivation has been produced, and that the structural reason for this is internal to the claim itself.

The reply also addresses a particular variant of the Stroud-style objection that is sometimes formulated separately. The variant says that the chapter has merely shown what the rival's claim cannot do for us — what we cannot infer from it, what we cannot articulate using it — rather than what is the case independently of us. The reply has the same structure. The chapter has not been showing what we cannot infer; it has been showing what the rival's claim, considered as the philosophical thesis it advertises itself as being, cannot internally accomplish. The asymmetry the claim asserts is not asserted by the inferrer; it is asserted by the claim. The question is whether the claim's content delivers the asymmetry. The chapter argues that it does not, for reasons internal to what the claim requires of its terms.

It is worth noting briefly what the chapter is not attempting. The chapter is not trying to rescue the truth-directed form of transcendental argument from the modal objection. It is not arguing that the standard transcendental argument really does establish what it appears to establish, properly understood. It is not proposing a modified form of transcendental argument that escapes the Stroud-style worry by retreating to belief-directed or experience-directed claims. The chapter's argument is structurally distinct from the transcendental-argument programme altogether. The standard programme starts from a practice; the chapter starts from a claim. The standard programme asks what the practice's conditions imply; the chapter asks what the claim's content requires. The differences in starting-point and in question carry through to differences in what the result can support.

The chapter's argument can be summarised constructively, as the challenge it leaves open for any proponent. The proponent of an ontological derivation of awareness can show that the chapter's argument fails by producing a derivation that meets the local referential discipline of Chapter 2, articulates determinate specifications of X, the dependence relation, and the asymmetry, and exhibits awareness as ontologically prior in the sense the derivation needs — and that does all of this without articulating its terms within the field whose subordination it asserts. The chapter has argued that no such derivation has been produced and that the structural reason for this is internal to what derivations of this kind have to do. The reader who thinks the chapter has overreached is therefore invited to take the constructive route — to produce the derivation the chapter says cannot be produced. Until then, the chapter's structural result stands as a claim about what is possible for a class of claims, not as a thesis about what cannot be coherently thought.

3.7 The Bridge in operation: a worked illustration

It will help to show the argument in operation against a specific contemporary proposal. The proposal is John Searle's biological naturalism, taken in the version Searle defends in *The Rediscovery of the Mind* (Cambridge, MA: MIT Press, 1992) and in subsequent work. The choice is not arbitrary. Searle has been at pains to distance his position from both functionalism and reductive materialism, while insisting on the ontological dependence of consciousness on the brain. His view is therefore a clean case for present purposes: an ontological derivation claim made by a philosopher who is otherwise alive to many of the worries that motivate the present book.

Searle's claim, taken in its load-bearing form, is that consciousness is caused by and realised in the neurobiological processes of the brain, while remaining an irreducible feature of the world. The brain

is prior to consciousness in a determinate sense: brain processes can be specified as if no consciousness were involved; consciousness cannot be specified except as something the brain produces and sustains. The dependence is asymmetric.

Searle's position is not a reduction of consciousness to third-person behavioural or computational description; indeed, part of its force lies in rejecting such reductions while retaining biological dependence. The argument that follows does not contest this. What it presses is the form of the dependence claim itself — the priority asymmetry — and not the wider position from which Searle deploys it.

The argument of this chapter applies to this proposal as follows.

The prior item X is determinate: neurobiological process at some level of description. This is a serious empirical and theoretical posit, and the argument has no quarrel with the underlying neuroscience. The dependence relation is determinate: causal in some specified sense, with the brain producing consciousness as one of its higher-level features. The independence claim is determinate: the brain's processes are specifiable without reference to the consciousness they produce.

Now apply the diagnosis. The specification of brain processes draws on the conceptual resources of neuroscience. Those resources are intelligibly deployed in inquiry, in publication, in laboratory practice, in argumentative philosophy. The specification of the causal relation between brain process and consciousness draws on the conceptual resources of the philosophy of causation. Those resources are also intelligibly deployed. The independence claim — that the brain can be specified without invoking consciousness — is intelligible as a claim, with consequences, and with arguments that have been mounted in its favour.

Each of these specifications is articulated within the field whose subordination Searle's view proposes. The view's content is that the brain is ontologically prior to consciousness, and that consciousness is what the brain produces. The view's articulation is an exercise of meaning-bearing intelligibility — that is, of the field within which the brain is identifiable as a brain, causation is identifiable as causation, and priority is identifiable as priority. The result is not that Searle has overlooked something. It is that his proposal, taken as an ontological derivation of awareness, has the structure the argument of §§3.2–3.4 diagnoses: its load-bearing terms can do their work only inside the field whose subordination is at issue.

This result has narrow consequences. It does not impeach Searle's objections to functionalism. It does not impeach his arguments against the Strong AI programme. It does not contest the neuroscience he draws on. What it shows is that the *priority* asymmetry built into the formulation cannot be stated coherently from outside the field whose subordination it proposes. The view remains a substantial contribution to the philosophy of mind in many of its other respects. Its load-bearing derivation claim is what the argument here pressures.

The diagnosis does not depend on Searle's particular positioning within the philosophy of mind. Versions of biological naturalism that handle emergentism differently, that adopt different accounts of downward causation, or that calibrate the ontological commitments of brain-talk differently from Searle face the same structural pressure. The pressure does not turn on the philosophical apparatus a biological naturalist deploys; it turns on the form of the dependence claim. Wherever a biological naturalist asserts that consciousness depends on the brain in a way the brain does not depend on consciousness, the asymmetry built into the formulation cannot be marked from outside the field within which "the brain," "consciousness," and "dependence" are intelligible as the items they are. The chapter's argument here is not against Searle but against the form of the position.

The same procedure can be applied to other proposals — to identity theories, to higher-order theories, to

global workspace theories, to integrated information theory, to any view that articulates an ontological derivation of awareness from some prior X. The procedure is the same in each case: identify the prior, identify the dependence relation, identify the independence asymmetry, ask whether the determinate articulation of these places the claim inside the field whose subordination it proposes. In each case, the answer the chapter contends for is the same. The derivation does not reach over the field. It articulates within it.

To make the portability of the diagnosis concrete, consider one schematic variant. Suppose a view identifies awareness with some physical or functional item I, where the specification of I is whatever the proponent prefers. The view holds that awareness just is I. The diagnosis applies, with the appropriate adjustments. To make the identification do work, the proponent must specify I determinately enough for the identity to be marked rather than merely declared. The specification proceeds within the conceptual resources of whatever discipline frames I — neuroscience, computational theory, dispositionalist semantics, or other. Those resources are intelligibly deployed; their intelligibility is a feature of meaning-bearing activity within the field. The identification, considered as a claim that awareness is exhausted by I, requires that "I" and "awareness" both be specifiable at the level of determinacy the identification needs — and the level of determinacy the identification needs is the level at which the items the identification connects are intelligible as the items they are. That intelligibility is what the chapter has been pointing at as the field. The identification, considered as a claim about awareness, sits within the field whose subordination it proposes. The pressure is the same as for derivation.

The diagnosis is robust because it operates at the level of what any such claim has to do, not at the level of the philosophical apparatus the claim deploys. A view that takes I to be a physical state and a view that takes I to be a functional role differ in their metaphysical commitments; both face the same structural pressure on the identification's load-bearing terms. The pressure is not about the philosophical merits of any specific approach. It is about what an identification claim about awareness — whatever the kind of item the proponent identifies awareness with — has to do to function as the kind of claim it is.

The treatment in this section shows the portability of the Chapter 3 diagnostic — not the chapter's full engagement with content-architectural identifications. Chapter 4 takes up that engagement in primary-text detail, addressing Block's distinction between access-consciousness and phenomenal consciousness and Carruthers's dispositionalist higher-order thought theory. The schematic identification considered here is meant only to demonstrate that the structural pressure of Chapter 3 carries beyond ontological derivation to identification, with the appropriate adjustments. The fuller treatment, with the load-bearing premises and specific argumentative responses each identification theorist deserves, is the business of the next chapter.

3.8 What this chapter does and does not establish

A short closing paragraph, in place of the symmetrical caveat the prior draft had at the end of each section.

The chapter has done one thing. It has argued that no ontological derivation of awareness can be stated without presupposing the field in which the derivation's key terms become intelligible. The chapter has not argued that no reality exists independently of awareness. It has not argued that empirical inquiry into the brain, the body, or the world is invalid. It has not argued that awareness is a substance, a mysterious cosmic stuff, or an item of any metaphysical inventory. It has argued that one specific form of ontological subordination — derivation from a prior X — cannot be coherently articulated. The chapters that follow extend the same form of argument to other forms of subordination: identification,

temporal conditioning, and the multiplication of awareness into countable subjects. Each of those is a separate argument and a separate chapter.

Chapter 4 — Against Identification

4.1 Identification as a subtler mislocation than derivation

Chapter 3 addressed the most direct way of subordinating awareness: derivation. The rival view there claimed that awareness was ontologically derivative from some prior item — a brain process, a computational structure, a biological organisation. The chapter argued that no such derivation could be coherently stated without articulating its terms within the field whose subordination it was meant to demonstrate. The conclusion was narrow and structural. It did not deny the existence of brains, computations, or biological systems. It denied that any of them could serve as a prior item in an ontological derivation of awareness without the derivation's articulation drawing on awareness in the way the field-of-intelligibility argument identified.

The present chapter takes up a different form of subordination, or rather a family of moves that can be made in place of derivation. Where derivation places awareness ontologically below some prior item, the moves examined here place awareness within an inventory of items the inventory itself does not specify. Identification is the central case: the proposal that awareness simply *is* some determinate feature of cognition — a kind of conscious state, an access relation, a content-architectural arrangement, an availability relation between perceptual contents and consumer systems. Identification need not be reductive in the crude sense of denying awareness's reality. The most sophisticated identification accounts in contemporary philosophy of mind affirm the reality of awareness while taking it to consist in some specifiable feature of mental organisation.

This is a subtler form of mislocation than the one Chapter 3 addressed. The derivation argument was structurally cleaner because the rival view had to specify an asymmetric priority relation between two relata, and the rival's load-bearing terms could be examined in their roles as prior or dependent. An identification claim has a different structure. It does not place awareness below something more basic. It places awareness *among* a particular class of cognitive items — phenomenally conscious states, higher-order representational relations, contents made available to particular consumer systems, properties of mental episodes that figure in introspective recognition. The argument's pressure point cannot be the priority asymmetry, because there is no asymmetry in the identification claim. The argument's pressure point must be located differently.

The chapter argues that the structural form of the book's central argument applies to identification claims as well, but operates at a different level than the determination of priority. What is at issue with identification claims is not whether something is prior to something else. It is whether the items the identification proposes are the kinds of items awareness, as the field of intelligibility, can be identified with. The chapter argues that they are not — not because identification is methodologically illegitimate, and not because the items in question are unreal, but because the items in question are determinate features within an inventory whose intelligibility as an inventory presupposes the field the identification

proposes to specify.

Two of the most articulate contemporary engagements with phenomenal consciousness serve as the chapter's targets. They occupy different positions on the spectrum the chapter is concerned with. The first is Ned Block's distinction between access-consciousness and phenomenal consciousness, developed in his 1995 *Behavioral and Brain Sciences* target article. Block does not himself identify phenomenal consciousness with anything in particular; on the contrary, he resists the reduction of phenomenal to access. But his careful sorting of kinds within the inventory of conscious states is itself a position the book's argument bears on, because the sorting operates at the inventory level. The second is Peter Carruthers's dispositionalist higher-order thought theory, developed in his 2000 book *Phenomenal Consciousness: A Naturalistic Theory*. Carruthers performs the constitutive identification Block declines to perform: phenomenal consciousness, on Carruthers's account, is the availability of analog perceptual contents to a higher-order-thought-wielding consumer system, with consumer semantics conferring categorical dual content. The book's pressure point on Carruthers is the form of an identification argument; on Block it is the structurally adjacent point about the level at which the careful inventory-sorting takes place.

The chapter takes Block first. The reason is dialectical. Block's careful distinction-defending posture is closer to the book's own discipline than Carruthers's identificatory move, and the engagement with Block can begin with substantial agreement before locating the structural pressure. With that agreement and pressure in view, the engagement with Carruthers can then locate a different kind of move — one Block himself has not made — and bring the book's argument against identification to bear on it directly. The chapter close brings the two engagements together as engagements with two positions at different points on the same dialectical landscape.

Three preliminary remarks are needed.

First, the chapter is concerned only with identification claims about awareness. It does not engage identification claims about other things — about persons, about social entities, about mathematical objects, about anything for which identification with a determinate item is proposed. The book's central argument has nothing to add to those other identification debates. The argument's specific bearing is on the case in which awareness, understood as the field of intelligibility, is identified with a determinate feature of cognition. The argument concerns the structural mismatch between field and feature, not the legitimacy of identification claims in general.

Second, the chapter is not an argument for first-order representationalism. The dispositionalist higher-order theorist's case against first-order accounts has its own internal cogency, and the book's engagement with Carruthers should not be misread as a defence of Dretske, Tye, or the broader first-order representational family against him. The book's pressure point operates at a level neither first-order nor higher-order theories address. A first-order representationalist would face structurally similar pressure from the book's argument: identifying phenomenal consciousness with first-order representational content places phenomenal consciousness inside the inventory of content the field of intelligibility discloses. The book is not in the higher-order-versus-first-order dispute. It is at a level both sides of that dispute presuppose.

Third, the chapter is not an argument that the cognitive-scientific research programmes opened by either Block's distinction or Carruthers's identification will fail. They may succeed, and the book has no investment in their failing. The book's argument is structural and operates at a level the success of the programmes does not bear on. Block's distinction has shaped three decades of debate in productive ways; Carruthers's account has motivated substantial work in higher-order theory and consumer semantics; neither contribution is at issue here. The chapter contests one specific philosophical move

sometimes attached to each: the move from "this is the structure of conscious states" or "this is what phenomenal consciousness consists in" to a specification of what awareness is, in the sense the book has been developing.

With those preliminaries in view, the engagement with Block can begin.

4.2 Block: a real distinction, still inventory-level

4.2.1 The argument

Ned Block's 1995 target article in *Behavioral and Brain Sciences*, "On a Confusion about a Function of Consciousness," is the canonical statement of a distinction that has structured debates in philosophy of mind and consciousness science for three decades. The article's central thesis is that "consciousness" is what Block calls a *mongrel concept*: the word names different phenomena, conflated in ordinary use and in the philosophical and scientific literature. The conflation is not innocent. It produces fallacious reasoning. Block's central case in point is what he calls the "target reasoning" about blindsight: the inference from "blindsight patients lack consciousness of stimuli in the blind field and cannot use the information to guide action" to "a function of consciousness is to enable information to guide action."

Block argues that this inference equivocates on "consciousness" (Block 1995, p. 227). What is missing in blindsight is both phenomenal consciousness *and* access-consciousness of stimuli in the blind field. An obvious function of access-consciousness — making information available for reasoning, reporting, and rational control of action — is illicitly transferred to phenomenal consciousness. To diagnose the fallacy properly requires distinguishing the kinds of consciousness, and that distinction is the article's central methodological achievement.

The distinction itself is between two kinds of conscious state, or two kinds of feature of conscious states (Block 1995, pp. 230–232).

Phenomenal consciousness, or *P-consciousness*, is experience — what it is like to be in a state. Block acknowledges that he cannot define P-consciousness in any "remotely noncircular way" (Block 1995, p. 230); the definition consists in pointing to it with rough synonyms: experiential properties, the totality of what it is like to be in the state. Block lists examples — visual experiences of colour, auditory experiences of sound, pains, feelings, thoughts that have phenomenal aspect. He treats P-conscious properties as "distinct from any cognitive, intentional, or functional property" (Block 1995, p. 230), and acknowledges this distinctness as the controversial part of his account.

Access-consciousness, or *A-consciousness*, is defined functionally. A state is access-conscious if, in virtue of one's having the state, "a representation of its content is (1) inferentially promiscuous… that is, poised for use as a premise in reasoning, (2) poised for rational control of action, and (3) poised for rational control of speech" (Block 1995, p. 231). A-consciousness is system-relative; what makes a state A-conscious is what a representation of its content can be used for. A-conscious states are "necessarily transitive" — they are about something. The paradigm A-conscious states are propositional attitudes; the paradigm P-conscious states are sensations.

Block lists three differences between the two kinds (Block 1995, pp. 231–232). First, P-conscious content is phenomenal, while A-conscious content is representational. Second, A-consciousness is a functional notion (and so system-relative), while P-consciousness is not. Third, there is a P-conscious *type* or *kind* — the feel of pain is a type — while A-consciousness is state-specific.

Block also argues for the conceptual possibility of dissociation between the two kinds. *A-consciousness without P-consciousness* is illustrated by the full phenomenal zombie (a computational duplicate with-

out phenomenal experience) and, more concretely, by Block's hypothetical "superblindsight" patient — a blindsight subject trained to spontaneously report on contents in the blind field without ever having phenomenal experience of them (Block 1995, pp. 232–233). *P-consciousness without A-consciousness* is illustrated by the case of the pneumatic drill outside the window of which one was unattendingly P-conscious for some time before noticing it, and by Sperling iconic-memory cases where the apparent richness of the visual array outstrips what is reportable (Block 1995, pp. 233–234).

The argument's dialectical context matters. Block develops the distinction in the course of resisting a particular family of identifications: the functionalist tradition that would identify consciousness with access, broadcasting, or some other functional-cognitive feature. Block's argument is that this tradition has been operating with an undifferentiated concept of "consciousness" that conflates the access-functional with the phenomenal — and that the conflation has driven a generation of arguments about the function of consciousness that mistake A-consciousness's evident functional role for a function of P-consciousness. The distinction is Block's instrument for separating what can be functionally explained from what (he argues) functional explanation does not reach.

Block also distinguishes two further kinds of consciousness that are not the chapter's primary concern but should be noted for completeness. *Self-consciousness* is the possession of the concept of self and the ability to use it in thinking about oneself (Block 1995, p. 235). *Monitoring-consciousness* is various forms of inner perception or internal scanning (Block 1995, p. 235). Block rejects the identification of P-consciousness with either, on grounds that ordinary laptop computers can perform internal scanning without being P-conscious, and that dogs and babies can be P-conscious without being self-conscious.

4.2.2 The argument reconstructed

The article's central argument can be reconstructed in numbered form.

1. The word "consciousness" names different phenomena, conflated in ordinary use and in the philosophical and scientific literature.
2. At least two distinct kinds of consciousness can be carefully identified: phenomenal consciousness (P-consciousness) and access-consciousness (A-consciousness).
3. P-conscious content is phenomenal; A-conscious content is representational and access-functional.
4. The two kinds are conceptually distinct, and their dissociation is conceptually possible (the zombie case, the superblindsight case, the pneumatic-drill case, the Sperling case).
5. Therefore the conflation of P-consciousness with A-consciousness is not innocent: the two are not the same thing.
6. The standard reasoning about a function of consciousness from blindsight to a functional role for phenomenal consciousness equivocates on "consciousness" and is fallacious.
7. The correct conclusion is narrower: what blindsight shows about the function of consciousness concerns A-consciousness, not P-consciousness.

The distinction's load-bearing role in Block's argument is to give the article a structural lever against the functionalist tradition: by separating the access-functional from the phenomenal, Block prevents the access-functional from doing work that requires the phenomenal in order to be a function *of* consciousness in the relevant sense.

4.2.3 The load-bearing premise

The load-bearing premise for the present engagement can be stated explicitly.

> Block's A-consciousness / P-consciousness distinction distinguishes kinds or features of conscious mental states: access-consciousness by functional availability for reasoning, action, and report; phenomenal consciousness by experiential or phenomenal character. The distinction is real and important, but it remains a distinction within the inventory of conscious states rather than an account of the field within which any such inventory is intelligible.

The premise is exhibited throughout the article. It is operative when Block lists examples of P-conscious states (perceptual experiences, sensations, feelings, some thoughts) and contrasts them with paradigm A-conscious states (propositional attitudes). It is operative when Block specifies the three differences between the kinds. It is operative when Block argues for the conceptual possibility of dissociation between the kinds, since dissociation presupposes that the kinds are kinds of conscious state that can come apart while both being kinds of conscious state.

The premise is well-defended in Block's article and the book agrees with it as far as it goes. The book's engagement is not that Block has failed to mark a genuine distinction. The engagement is structural and operates at a level Block's article does not directly address: the level of the field within which the kinds Block distinguishes are intelligible as kinds.

4.2.4 The pressure point developed

The first thing to mark is what the pressure point is not. It is not an objection to the A/P distinction. The distinction is real and Block's case for it is convincing. It is not a claim that access and phenomenality are the same thing. The book agrees with Block that they are not. It is not an objection to Block's resistance to the functionalist tradition. The book agrees with Block that consciousness, in the phenomenal sense, is not reducible to access. The pressure point is structurally adjacent to Block's own argument and operates at a different level than the dispute he is engaged in.

The book's argument and Block's argument converge on one point and diverge at another. They converge in resisting the reduction of consciousness to access-functional notions. Block makes this resistance through the A/P distinction: separating phenomenal from access, defending the phenomenal as a distinct kind of consciousness, and treating attempts to collapse the two as fallacious. The book's resistance is differently located. The book argues that what makes either kind of conscious state — A-conscious, P-conscious, or both — intelligible *as* a conscious state is the field within which such states have their standing as states. The field is not a kind of conscious state. It is not a property that conscious states have. It is not what the A/P distinction sorts.

Consider what the A/P distinction sorts. P-conscious states are states with phenomenal character — visual experiences, auditory experiences, pains, sensations, some thoughts. A-conscious states are states with access-functional content — propositional attitudes poised for use in reasoning, reporting, and the control of action. Both kinds of state are determinate items in an inventory of mental states. Each can be described, contrasted with neighbours, modelled in cognitive-architectural terms, investigated empirically. Block's distinction is a sorting of this inventory into two kinds of item, and the distinction's value lies in making the sorting fine-grained enough to prevent the kinds of conflation Block diagnoses.

What the distinction does not do is what the book has been concerned with throughout. The distinction does not address the field within which any state — P-conscious, A-conscious, or both — has whatever standing it has as a state of the kind it is. Whatever a P-conscious state is, it is intelligible as a P-conscious state within the field. Whatever an A-conscious state is, it is intelligible as an A-conscious state within the field. The distinguishing of one from the other, in Block's careful manner, takes place within the

field. The dissociation cases Block constructs — the zombie, the superblindsight patient, the pneumatic drill, the Sperling array — take place within the field. The arguments by which Block defends the distinction take place within the field. The field is not a kind of conscious state, and the distinction does not specify it.

This is structurally important to mark precisely. The book's argument is *not* that there is a third kind of consciousness, alongside P-consciousness and A-consciousness, that Block has missed. There is no third kind. The book is not adding a kind to Block's inventory. The book is observing that Block's inventory — however carefully sorted, however many kinds the sorting recognises — is an inventory of items the field of intelligibility discloses. The field is not what gets sorted. It is what makes the sorting a sorting.

This pressure becomes sharper if one looks at Block's own treatment of how A-consciousness and P-consciousness relate. At the conclusion of his §4 (Block 1995, p. 232), Block makes a point worth quoting carefully: "Although I make a firm distinction between A-consciousness and P-consciousness, I also want to insist that they interact." Block then gives examples. Attention to a peripheral content — the feel of the shirt on your neck — can switch what was in the background to the foreground, changing one's phenomenal state. A figure-ground switch involves both an A-consciousness change (in what is being accessed) and a P-consciousness change (in what is being experienced). The two kinds are conceptually distinct, in Block's view, but empirically often co-instantiated and interactive.

The interaction claim is important because it shows that Block is not crudely separating A and P into sealed compartments. The kinds are distinct but related; they can come apart in principle but often go together in practice. The book's engagement should not flatten this nuance. Whatever the empirical interaction between A and P, however the two kinds co-instantiate and influence each other, the interaction remains an interaction among kinds of conscious state within the inventory the field of intelligibility discloses. The field is not what the interaction operates between; it is what the interaction itself is intelligible within. Even granting Block's most sophisticated treatment of A-P interaction — and the book grants it — the structural observation about the inventory and the field stands.

A second textual hinge is worth marking. In a footnote at Block 1995, p. 232, Block writes: "Those who are uncomfortable about P-consciousness should pay close attention to A-consciousness, because it is a good candidate for a reductionist identification with P-consciousness." The footnote is Block being characteristically generous to interlocutors who might be inclined to identify P with A. Block does not himself endorse such a reduction; he treats it as a possibility worth taking seriously without being committed to it. The book's engagement should note that Block's openness to such reductionist identifications is a place where the book and Block differ. The book does not regard A-consciousness as a candidate for a reductionist identification with P-consciousness, because the book regards no determinate kind of conscious state as a candidate for identification with awareness, in the sense the book has been developing. Block is in a different dispute than the book is in.

The position the book takes on Block can be stated cleanly. Block is right that A-consciousness and P-consciousness should not be conflated. Block is right that the standard reasoning about a function of consciousness from blindsight equivocates between them. Block is right to resist the functionalist tradition's reduction of consciousness to access. The book agrees with all of this. What the book adds is that the kinds Block distinguishes are kinds within the inventory of conscious states, and that the field within which any such kind has whatever standing it has as a kind of conscious state is not what the A/P distinction specifies.

4.2.5 What the critique establishes, and what it does not

The scope of the critique should be stated carefully.

The critique establishes that Block's A/P distinction, valuable as it is for diagnosing conflations in the philosophy and science of consciousness, operates at the level of kinds of conscious state within an inventory. It does not address the field within which such states have their standing as states. The book's argument is structurally adjacent to Block's own resistance to functionalist reduction and locates the structural pressure at a different level than the level at which Block operates.

The critique does not establish that Block's distinction is misconceived. On the contrary, the distinction is well-defended and the book has already drawn on the distinction's dialectical importance in its engagement with other positions, notably the Frankish engagement in Chapter 5. The book has nothing to add to the diagnostic value of the distinction within debates about the function of consciousness.

The critique does not establish that A-consciousness and P-consciousness are the same thing. The book agrees with Block that they are not. The book's argument neither presupposes nor implies any such identification.

The critique does not establish that Block has missed the field-question. He has not addressed it in the 1995 article, but he is not concerned with it; the article's concerns are the conflation diagnosis and the careful defence of a distinction within the inventory. The article does not pretend to address the field, and the book's critique should not pretend that Block intended otherwise.

The critique does not establish that there are no further useful distinctions to be drawn within the inventory of conscious states. Block himself recognises self-consciousness and monitoring-consciousness as further kinds. The book has no investment in the question of how many kinds the inventory should contain or how they should be individuated. The book's argument bears on the relation between the inventory (however richly populated) and the field, not on the internal structure of the inventory.

What the critique establishes is one thing. It is that the resources of careful kind-distinguishing within the inventory of conscious states, however refined, do not specify awareness as the field of intelligibility — not because the distinguishing is wrong, but because the field is not what such distinguishing is structured to specify.

4.3 Carruthers: constitutive HOT identification

4.3.1 The argument

Peter Carruthers's *Phenomenal Consciousness: A Naturalistic Theory*, published by Cambridge University Press in 2000, develops one of the most careful contemporary identifications of phenomenal consciousness with a determinate feature of cognitive architecture. The book defends what Carruthers calls *dispositionalist higher-order thought theory*. The central thesis is stated in the book's Preface (Carruthers 2000, p. xiii):

> Phenomenal consciousness consists in a certain sort of intentional content ('analog', or fine-grained), held in a special-purpose short-term memory store in such a way as to be available to higher-order thoughts about the occurrence and nature of those contents; and that in virtue of such availability (given the truth of some or other form of 'consumer semantics') all of those contents are at the same time higher-order ones, acquiring a dimension of seeming or subjectivity.

This is a constitutive identification of phenomenal consciousness with a specifiable content-architectural arrangement. Phenomenal consciousness, on Carruthers's view, *is* the availability of analog perceptual contents in a special-purpose short-term memory store to a higher-order-thought-wielding consumer system that understands the is-seems distinction. The identification is naturalistic and reductive in a precise sense: phenomenal consciousness is what such availability under consumer semantics *consists in*.

The view is positioned against four families of competing accounts. *Mysterianism* (Nagel, McGinn, Chalmers) treats phenomenal consciousness as irreducibly resistant to naturalistic explanation; Carruthers argues in earlier chapters of his book that mysterian arguments commit fallacies of ambiguity. *First-order representationalism* (Dretske, Tye) takes phenomenal consciousness to be constituted by certain first-order perceptual contents alone; Carruthers argues at length (Chapters 5 and 6 of the book) that first-order accounts cannot mark the distinction between conscious and non-conscious experience. *Higher-order experience theory* or "inner sense" theory (Armstrong, Lycan) takes phenomenal consciousness to be constituted by inner sensings of first-order experiences; Carruthers argues that this is computationally extravagant and evolutionarily implausible. *Actualist higher-order thought theory* (Rosenthal) requires that an actual higher-order thought target a state for that state to be phenomenally conscious; Carruthers argues that this requires an implausible amount of on-line higher-order cognitive activity and that dispositional availability is sufficient.

The dispositional view's distinctive contribution is the move from actual to dispositional availability. The contents of the special-purpose short-term memory store need not be *actually* targeted by higher-order thoughts. They need only be in the right kind of relation to the higher-order-thought-wielding consumer system. This dispositional structure, Carruthers argues, allows the account to explain the richness of phenomenally conscious experience without requiring a rich array of actual higher-order thoughts at every moment of conscious experience (Carruthers 2000, Chapter 8).

The view's most distinctive technical move is the appeal to *consumer semantics* to explain how dispositional availability can confer categorical phenomenal feel. Chapter 9 of the book develops this in detail. The version of consumer semantics Carruthers endorses is a localist *inferential role semantics*: the immediate inferential connections entered into by a representation are partially constitutive of its content (Carruthers 2000, p. 263). The mind-reading consumer system, by virtue of its capacity to deploy higher-order recognitional concepts and its understanding of the is-seems distinction, transforms the content of perceptual states made available to it. Each percept of red is at one and the same time an analog representation of *reda* and a representation of *seems-reda* (Carruthers 2000, p. 264). The dual content is categorical, not merely dispositional; it is what the analog perceptual content *is*, when held in a memory store that makes it available to the relevant consumer system.

Carruthers's account of subjective feel turns on this dual-content claim. Where Block insists that phenomenal consciousness is distinct from any cognitive, intentional, or functional property, Carruthers argues that phenomenal consciousness is constituted by a specific kind of intentional content — the higher-order analog content (*seems-reda*, *experience of reda*) that perceptual states acquire by being available to the mind-reading consumer system. The subjective feel of experience is, on Carruthers's account, the property of having such higher-order analog content.

At the close of Chapter 9, in his discussion of how to unify the treatment of outer perception and bodily sensation, Carruthers offers a canonical characterisation of phenomenal subjectivity (Carruthers 2000, p. 270):

> The relevant kind of subjectivity attaching to all phenomenally conscious states, which constitutes them as phenomenally conscious, is that they should possess properties of a sort

> which can be available for immediate introspective recognition.

This sentence is the book's most exposed identification claim and the textual hinge for the present engagement. Phenomenal subjectivity is identified, in this sentence, with availability for immediate introspective recognition. The chapter's pressure point will turn on what this identification can and cannot do.

4.3.2 The argument reconstructed

The central argument of Carruthers's book can be reconstructed in numbered form.

1. Physicalism and naturalism are default assumptions, defeasible only by very powerful contrary arguments.
2. The mysterian arguments against naturalising phenomenal consciousness commit fallacies and do not constitute such arguments.
3. First-order representationalism fails because it cannot mark the distinction between conscious and non-conscious experience without either drawing the line arbitrarily or positing phenomenally conscious states to which the subject is blind.
4. Therefore some form of higher-order representational theory is the right approach.
5. Higher-order experience ("inner sense") theory is computationally extravagant and evolutionarily implausible.
6. Actualist higher-order thought theory requires too much on-line higher-order activity and is also evolutionarily implausible.
7. Dispositionalist higher-order thought theory, combined with consumer semantics, explains both the conscious–non-conscious distinction and the subjective feel of conscious experience.
8. Dispositional availability to a HOT-wielding consumer system suffices for phenomenal consciousness because consumer semantics — specifically, inferential role semantics — confers categorical dual content on the analog perceptual states made available.
9. Phenomenal subjectivity is identified with availability for immediate introspective recognition; the dual-content claim is what makes the identification do the explanatory work it is required to do.
10. Therefore phenomenal consciousness is reductively explained in naturalistic terms: it is the availability of analog perceptual contents to a HOT-wielding consumer system, with consumer semantics conferring a dimension of seeming or subjectivity.

The premises that do most of the work for the engagement that follows are premise 8 (the consumer-semantic dual-content claim) and premise 9 (the identification of phenomenal subjectivity with availability for immediate introspective recognition). The dispositional structure of the view, the rejection of competing higher-order accounts, and the broader naturalistic framing are not contested by the book's argument. The chapter's engagement turns on what the identification at premises 8 and 9 specifies and does not specify.

4.3.3 The load-bearing premise

The load-bearing premise for the present engagement can be stated explicitly.

> Carruthers's dispositionalist HOT account identifies phenomenal consciousness with a determinate content-architectural feature: analog perceptual contents, held in a special-purpose short-term memory store, made available to a HOT-wielding consumer system under consumer semantics. The book's pressure is not that this architecture fails by

> Carruthers's own standards, but that such an architecture remains an item within the intelligible inventory rather than the field within which any such architecture has its standing.

The premise is exhibited throughout the book. It is stated explicitly in the Preface. It governs the rejection of first-order representationalism (the case for which is that first-order accounts lack the higher-order availability needed to mark the conscious–non-conscious distinction). It governs the rejection of higher-order experience theory (a competing architectural arrangement is found computationally implausible). It governs the rejection of actualist HOT (a stronger version of the same architectural arrangement is found to require too much cognitive activity). It governs the consumer-semantic explanation of phenomenal feel (consumer semantics is what underwrites the categorical-content claim). It governs the closing characterisation at Carruthers 2000, p. 270 (phenomenal subjectivity is identified with availability for immediate introspective recognition).

The premise is reasonable within the research programme Carruthers is engaged in. The point of identifying it as the book's load-bearing premise is not to suggest that Carruthers has chosen the wrong identificatory target. He has not, by the lights of his programme. The point is to mark precisely what the identification does, and what it does not, deliver.

What the identification delivers, if it succeeds, is a content-architectural reconstruction of phenomenal consciousness. The reconstruction explains how analog perceptual contents acquire higher-order significance under consumer semantics; how the conscious–non-conscious distinction can be marked by reference to the special-purpose short-term memory store and the consumer systems it feeds; how the richness of conscious experience can be preserved without an unrealistic burden of actual higher-order thoughts at every moment; how qualia-as-intrinsic-properties can be denied while the temptation to believe in them is explained away. This is a substantial reductive achievement, on Carruthers's terms.

What the identification does not deliver is the field-question the book has been pursuing. The field within which the special-purpose memory store, the analog contents, the consumer systems, the higher-order recognitional concepts, the inferential dispositions that constitute consumer semantics, and the dual content of perceptual states are all intelligible as the items they are is not what Carruthers's content-architectural reconstruction specifies. The field is not a content-architectural feature. It is not what consumer semantics articulates. It is not what is meant by availability-for-immediate-introspective-recognition. It is the standing within which all such things have their standing.

That is the pressure point. The section develops it through Carruthers's textual hinge at p. 270 and his dual-content claim at p. 264.

4.3.4 The pressure point developed

The first thing to mark is what the pressure point is not. It is not an objection to the dispositional structure of Carruthers's account. The shift from actualist to dispositionalist HOT is well-motivated within higher-order theory, and the book has no objection to it. It is not an objection to consumer semantics. The book takes no position on whether inferential role semantics is the correct theory of how representational content is determined, and the engagement does not require contesting it. It is not an objection to the dual-content claim. The book grants, for the sake of argument, that perceptual contents made available to a HOT-wielding consumer system acquire categorical dual content under consumer semantics. The pressure point operates after these are granted.

The pressure point is on the identification claim itself. Phenomenal subjectivity is identified, at Carruthers 2000, p. 270, with availability for immediate introspective recognition. The dual-content claim

at p. 264 is the mechanism by which Carruthers's account discharges this identification: the dual content of perceptual states (*reda + seems-reda*) under consumer semantics is what such availability consists in.

Now consider what the identification specifies. Availability for immediate introspective recognition is a relation between two items within a cognitive architecture: a perceptual state held in a special-purpose memory store, and a higher-order-thought-wielding consumer system that includes a mind-reading faculty with higher-order recognitional concepts. The relation between these two items is what makes the perceptual state phenomenally subjective, on Carruthers's account. The dual content is the categorical feature of the perceptual state's content that emerges from this relation.

Each of these specifications operates within a fabric of cognitive-architectural intelligibility. A *perceptual state*, on Carruthers's account, is a determinate item: it has analog content, is held in a determinate memory store, has specifiable functional connections to other states. The *memory store* is determinate: it has a function (making contents available to consumer systems), a structure (short-term, special-purpose), a place in the cognitive economy of the organism. The *consumer system* is determinate: it has a mind-reading or theory-of-mind faculty, a capacity for higher-order recognitional concepts, an understanding of the is-seems distinction. The *availability relation* is determinate: it is the relation between contents in the memory store and consumer systems that can use those contents. The *dual content* is determinate: it is the content the perceptual state has, under inferential role semantics, by virtue of its inferential dispositions in the context of the available consumer systems.

Each of these is an item or a feature of items in the inventory of cognitive architecture. Each is something Carruthers can describe, characterise, contrast with neighbours, defend by appeal to evolutionary and functional considerations. The whole content-architectural arrangement is the subject of theoretical inquiry, and Carruthers's book is, in important part, a defence of one such arrangement against competitors. The identification at p. 270 turns on this arrangement: phenomenal subjectivity is the availability-relation within it.

The book's argument has been concerned with a different level of question. What makes a cognitive-architectural arrangement intelligible as a cognitive-architectural arrangement? What makes the relation between perceptual state and consumer system intelligible as an availability relation? What makes the dual content intelligible as a content with two aspects rather than as some other kind of representational fact? Each of these is a question about the standing within which the arrangement, the relation, the content, and their characterisations have whatever standing they have. The field is not what is specified by the items the inventory contains. It is what makes the items intelligible as items.

This is structurally important to mark. The book's argument is not that Carruthers has overlooked a hidden ingredient in his cognitive architecture. There is no hidden ingredient. The book is not saying that Carruthers needs to add something to his C-store, his consumer systems, his inferential role semantics, his dual-content claim. The architecture, as Carruthers describes it, is internally complete within its own descriptive ambitions. The book's argument is that what makes the architecture an architecture *of phenomenal consciousness*, rather than of some other content-organisational fact about cognition, is the field within which the architecture has its standing as such. Identification with the architecture does not specify the field. It uses the field.

This pressure can be brought to bear on the consumer-semantic move directly. Consumer semantics, on Carruthers's account, is what underwrites the categorical-content claim. Perceptual contents acquire higher-order significance because of the inferential dispositions that hold between them and the HOT-wielding consumer system. The dispositions are categorical, the contents are categorically dual, and so phenomenal consciousness — identified with the dual content — is categorical. Even if consumer

semantics succeeds in delivering categorical dual content, the question the book's argument presses is: what makes the categorical dual content intelligible *as* phenomenal consciousness rather than as some other categorical content-architectural feature? The dual content, considered as a content-architectural feature, is what the cognitive architecture makes available. The intelligibility of the dual content *as phenomenal* is not what consumer semantics articulates. It is what consumer semantics presupposes for its own descriptions to function as descriptions of representational content.

The same pressure applies to the identification at p. 270. Availability for immediate introspective recognition is a determinate relation within a cognitive architecture. To call this relation phenomenal subjectivity is to make an identificatory claim: the relation just is what phenomenal subjectivity consists in. The book's pressure is not that the relation does not exist, or that the relation does not have the cognitive-architectural properties Carruthers attributes to it. It is that what makes such a relation intelligible *as* phenomenal subjectivity, rather than as a relation that has some other status, is not what the cognitive-architectural specification gives us. Even if availability-for-immediate-introspective-recognition explains what a HOT-capable consumer system can do with perceptual contents, it does not specify the field within which availability, recognition, perceptual content, and subjectivity are intelligible as such.

The position the book takes on Carruthers can be stated cleanly. Carruthers's dispositionalist HOT account, with its consumer-semantic underpinnings and dual-content claim, is a sophisticated and internally rigorous identification of phenomenal consciousness with a specifiable feature of cognitive architecture. The book agrees that the account, by its own lights, does substantial explanatory work. The book agrees that the dispositional structure has advantages over actualist alternatives. The book takes no position on consumer semantics or on the dual-content claim within consumer semantics. What the book adds is that the architecture Carruthers specifies — however internally well-defended — remains an item within the inventory of cognitive architectural arrangements that the field of intelligibility discloses, and that the identification of phenomenal consciousness with such an architecture does not specify the field.

Three brief acknowledgements are needed before this section closes, to discharge the defensive obligations the engagement with Carruthers's wider project requires.

First, on language. In Chapter 10 of his book Carruthers argues that phenomenal consciousness is independent of and prior to natural language. The argument is largely defensive — against his own earlier reflexive thinking theory and against Dennett's higher-order description theory. The book's argument does not contest the independence of phenomenal consciousness from language. The book has not argued, and does not argue, that phenomenal consciousness requires natural language. Whatever the right account of the relation between phenomenal consciousness and language, the book's pressure point on Carruthers's identification claim operates at a level the language question does not bear on.

Second, on the Cartesian theatre. In Chapter 11 Carruthers defends his commitment to a special-purpose memory store with determinate contents against Dennett's multiple-drafts theory and the charge that any account positing such a store reinstates a Cartesian theatre. The book takes no side in this dispute. The book's argument is compatible with various views about the determinacy or fragmentariness of conscious content. The pressure point on Carruthers's identification operates at a level the dispute about determinacy does not bear on; whatever view of the C-store the reader prefers, the identification still places phenomenal consciousness within an inventory whose intelligibility presupposes the field.

Third, on first-order representationalism. Carruthers's dispositionalist HOT view is partly motivated by his case against first-order representationalism in Chapters 5 and 6 of his book. The case has its own

internal cogency, and the book's engagement with Carruthers should not be misread as a disguised defence of Dretske, Tye, or the broader first-order representational family. The book's argument operates at a level neither first-order nor higher-order theories address. A first-order representationalist who identified phenomenal consciousness with first-order representational content would face structurally similar pressure from the book's argument: identification with a determinate content-level feature, whether first-order or higher-order, places phenomenal consciousness inside the inventory the field of intelligibility discloses. The book is not in the dispute between first-order and higher-order theories. It is at a level both sides of that dispute presuppose.

4.3.5 What the critique establishes, and what it does not

The scope of the critique should be stated carefully.

The critique establishes that Carruthers's identification of phenomenal consciousness with availability-for-immediate-introspective-recognition, however well-defended within his cognitive-architectural research programme, identifies phenomenal consciousness with a determinate feature of cognitive architecture. The architecture is an item within the inventory of intelligible items. The field within which the architecture has its standing is not what the identification specifies. The book's argument is structurally adjacent to Carruthers's account and operates at a level Carruthers's apparatus does not address.

The critique does not establish that Carruthers's account is incorrect within its own research programme. Carruthers's case for dispositionalist HOT over actualist HOT, over higher-order experience theory, and over first-order representationalism is internally cogent; the book takes no position on the merits of any of these higher-order disputes. The critique respects the architecture Carruthers builds and presses on what the architecture, taken as an identification of phenomenal consciousness, can and cannot deliver.

The critique does not establish that consumer semantics is wrong as a theory of how representational content is determined. The book takes no position on consumer semantics. The dual-content claim, granted for the sake of argument, is the mechanism by which Carruthers's account discharges its identification of phenomenal subjectivity with availability for immediate introspective recognition. The book grants the mechanism and presses on the identification.

The critique does not establish that there is no phenomenal consciousness in Carruthers's sense. The book is not eliminativist about phenomenal consciousness, however the term is used. The book's argument concerns what makes any of the items Carruthers's account specifies intelligible as the items they are, not whether the items exist.

The critique does not establish that Carruthers has failed to take account of awareness as the book characterises it. He has not addressed the field in the book's sense, but he is not concerned with it; the cognitive-architectural research programme operates at a different level. The critique respects this and adds that the difference of level is itself the relevant philosophical observation.

What the critique establishes is one thing. It is that an identification of phenomenal consciousness with a determinate feature of cognitive architecture, however internally well-supported, does not exhaust the philosophical work that an account of awareness has to do, because the identification operates within a field whose status is left untouched by it. The cognitive-architectural reconstruction is one thing; the field-question is another; the rigor of the first does not bear on the standing of the second.

4.4 What identification accounts can and cannot show

The two engagements of this chapter have addressed two of the most articulate contemporary engagements with phenomenal consciousness. They sit at different points in the philosophy of mind, and the chapter has been concerned to treat them in their own terms before bringing them together.

Block distinguishes access-consciousness from phenomenal consciousness and resists the functionalist tradition's reduction of consciousness to access. The book agrees with the resistance and adds that the distinction, valuable as it is, sorts kinds or features of conscious states within an inventory whose intelligibility as an inventory presupposes the field. Carruthers identifies phenomenal consciousness with a determinate higher-order content architecture — analog perceptual contents made available to a HOT-wielding consumer system, with consumer semantics conferring categorical dual content. The book grants the architecture and adds that the identification, however internally well-defended, operates on an item within the inventory rather than specifying the field within which any such architecture has its standing.

The two positions are not the same. Block resists reduction; Carruthers performs a sophisticated identification of the kind Block, in his 1995 footnote, treated as an open possibility worth taking seriously. The book's engagement with each is calibrated accordingly. Block is engaged as a sympathetic ally on the resistance to functionalist reduction, with the structural pressure located at a level Block does not address. Carruthers is engaged as the more directly pressured target, because his identification places phenomenal consciousness within the architectural inventory the book's argument is structured to mark as left untouched by such placement.

Beneath the differences, the chapter's two engagements converge on a structural observation. Identification accounts of phenomenal consciousness — whether they identify it with one kind of conscious state among others, or with a specific content-architectural feature of cognition, or with some other determinate item in the inventory — operate at a level the book's argument has been concerned to mark as inventory-level. The inventory may be richly populated, finely sorted, carefully defended; its kinds may be conceptually dissociable, empirically interactive, theoretically refined. None of this bears on whether the inventory, taken as a whole or in any of its parts, specifies the field within which the inventory and its items have their standing as items. The field is not what the inventory contains. It is what makes the containing intelligible as containing.

Block and Carruthers show, in different ways, how much philosophical work can be done within the inventory of conscious states. Block distinguishes access-consciousness from phenomenal consciousness; Carruthers identifies phenomenal consciousness with a higher-order content architecture. The present chapter has not denied the value of either project. It has argued that both remain within the inventory whose intelligibility they presuppose. Identification can sort, distinguish, and model structures within conscious life; it cannot specify awareness as the field in which such structures have standing at all.

A short closing word on what the chapter does not claim. It does not claim that Block or Carruthers have made errors by their own lights. Both have produced careful and influential work. The chapter has been concerned to treat each at their strongest, with the book's argument adjacent rather than opposed to their projects. It does not claim that the identification strategies are exhausted by the two engagements here; other contemporary identification accounts exist — global workspace theories, integrated information theory, predictive processing accounts of consciousness — and would face structurally similar observations. It does not claim that the cognitive-scientific work the identification accounts open will fail. They may succeed in their own terms. What the chapter claims is that their success, if achieved, would not bear on the question the book has been pursuing: the question of awareness as the field of

intelligibility within which the items the identification accounts specify have their standing as items at all.

The chapters that follow extend the same structural argument to two further forms of subordination. Chapter 5 considered illusionist and meta-problem reframings of the hard problem. Chapter 6 will take up temporal conditioning. Chapter 7 will take up the multiplication of awareness into countable subjects. The cumulative result of the chapters is not a new metaphysical posit but a clarified order of dependence: the inventory remains available for description and inquiry, while the field within which the inventory is intelligible as an inventory is not what such description and inquiry specify.

Chapter 5 — Illusionism and the Meta-Problem

5.1 Two responses to the hard problem

The hard problem of consciousness, in the form David Chalmers gave it in the 1990s, asks how physical or functional organisation could give rise to phenomenal experience — to the qualitative character of seeing red, of feeling pain, of tasting coffee. The problem is hard, in Chalmers's sense, because even after one has explained the cognitive, behavioural, and discriminatory features that surround such experiences, a residue remains: why there is something it is like to undergo them at all.

Two of the most articulate responses to the hard problem in the present generation operate not by solving it but by reframing it. Keith Frankish's illusionism replaces the hard problem with the *illusion problem*: the problem of explaining why experiences seem to have phenomenal properties when, on Frankish's view, they do not. David Chalmers's meta-problem shifts attention to a second-order question: why we judge, report, and intuit that there is a hard problem in the first place. Each is a response to the same pressure. Each restates part of the pressure as a question contemporary cognitive science is in principle equipped to handle. Each does so with care, with sophistication, and without straw-manning the position it leaves behind.

This chapter argues that neither response is wrong as a research programme. The illusion problem is a real problem; the meta-problem is a real problem. What the chapter argues is that each programme leaves untouched a question its own framing presupposes. Frankish's account of why experiences seem to have phenomenal properties operates within the field in which the seeming is intelligible as a seeming, the introspective representation is intelligible as a representation, and the explanation is intelligible as an explanation. Chalmers's account of why we form the judgements we form about consciousness operates within the field in which judgements, reports, and intuitions are intelligible as the kinds of acts they are. The field in question is what the book has been calling awareness. The chapter does not argue that either philosopher denies the existence of awareness. It argues that neither programme is structured to address what awareness is, in the sense the book has secured.

The two engagements are presented separately because the philosophers are doing genuinely different work, even though they operate at structurally adjacent points. Frankish's argument is metaphysical: phenomenal properties, as standardly understood, are not instantiated. Chalmers's argument is methodological: a theory of the judgements and reports about consciousness can be developed independently of any commitment to the underlying metaphysics. The chapter takes Frankish first, because his programme makes the sharper commitment.

5.2 Frankish on illusionism

5.2.1 The argument

Keith Frankish's case for illusionism is developed in "Illusionism as a Theory of Consciousness," published in 2016 in the *Journal of Consciousness Studies* and reprinted in his 2017 edited volume of the same name. The article is the cleanest contemporary articulation of strong illusionism about phenomenal consciousness. The engagement in this section is with the 2016 article itself; the commentaries in the 2017 volume are not treated here.

Frankish's central thesis is announced on the first page. Theories of consciousness, he notes, typically address the hard problem: they accept that phenomenal consciousness is real and aim to explain how it comes to exist. There is another approach. It holds that phenomenal consciousness is an illusion and aims to explain why it seems to exist (Frankish 2016, pp. 11–12). Frankish calls this approach illusionism and devotes the article to motivating it and defending it against the standard objections.

The article distinguishes three approaches to phenomenal consciousness (Frankish 2016, pp. 12–13). Radical realism takes phenomenal properties to be real and inexplicable in physical terms: dualists, neutral monists, mysterians, and certain panpsychists fall here. Conservative realism takes phenomenal properties to be real and explicable within the resources of contemporary cognitive science: most physicalist theories of consciousness fall here, including the various representational theories. Illusionism takes phenomenal properties to be unreal; what is to be explained is the introspective appearance of them rather than the properties themselves. Frankish distinguishes his strong illusionism from a weak version that accepts the reality of phenomenal properties while denying some of the features standardly attributed to them — ineffability, intrinsicality, privacy, infallible knowability. He argues that weak illusionism is unstable. It must either employ a concept of phenomenality stronger than the merely quasi-phenomenal, or collapse into strong illusionism. He doubts the stronger concept is available; in his view, only strong illusionism and radical realism are coherent options (Frankish 2016, pp. 13–14).

The technical move on which strong illusionism turns is the notion of a *quasi-phenomenal property*. Frankish defines it as "a non-phenomenal, physical property (perhaps a complex, gerrymandered one) that introspection typically misrepresents as phenomenal" (Frankish 2016, p. 14). Quasi-phenomenal redness, for example, is the physical property that typically triggers introspective representations of phenomenal redness. There is nothing qualitative about quasi-phenomenal properties. They present no special explanatory difficulty. Strong illusionists hold that the introspectable properties of experience are merely quasi-phenomenal. That is, the properties exist as physical features of sensory states, but introspection systematically misrepresents them as having a qualitative character they do not in fact possess.

The hard problem, on Frankish's account, is replaced by the *illusion problem*: the problem of explaining how the introspective representations that misrepresent experiences as phenomenally charged are produced, and why they are so vivid, so universal, and so resistant to correction (Frankish 2016, pp. 11, 13, 31–32). Frankish describes the shape of an illusionist research programme. It requires an account of the sensory states involved (probably modality-specific analogue representations of stimulus features such as position in a quality space, location, and intensity), an account of the introspective mechanisms that misrepresent these states, and an account of why the misrepresentation has the qualitative-seeming character it does. He cites Dennett's user-interface analogy, Humphrey's "ipsundrum" proposal, Rey's projection account, and Pereboom's analogy with secondary qualities as examples of the theoretical work that needs doing (Frankish 2016, pp. 15–17).

The motivating case for illusionism, as Frankish develops it in §2 (Frankish 2016, pp. 21–25), turns on

two arguments. The first is structural. Radical realism faces the threat of epiphenomenalism — phenomenal properties, conceived as non-physical, can have no causal effects in a physically closed world. Conservative realism, meanwhile, is unstable. The reductive explanations it produces tend, on closer examination, to be illusionist in disguise. Functional accounts of phenomenality, when they succeed, succeed by explaining why we are disposed to judge that experiences have qualitative dimensions, rather than by exhibiting the qualitative dimensions themselves. The second argument is positive and abductive. Frankish notes that even most realists accept that our claims and beliefs about consciousness can be fully explained without invoking phenomenal properties. Phenomenal zombies, lacking phenomenal properties but functionally identical to us, would make the same claims and form the same beliefs about consciousness that we do. If our claims and beliefs can be explained without invoking phenomenal properties, then the simplest hypothesis is that the properties are not, in fact, there.

In §3 (Frankish 2016, pp. 25–31), Frankish takes up the standard objections. The four addressed are: denying the data, the no-appearance–reality-gap objection, the question of who or what is the audience for the illusion, and the problem of how phenomenality can be represented if there are no phenomenal properties to represent. The treatment of the first develops a distinction that will be central to the engagement below: between *introspective subjectivity* (the introspective awareness of one's experiences generated by representational mechanisms, which illusionists accept) and *intrinsic subjectivity* (a subjective dimension not produced by introspective mechanisms but arising from what we are, which illusionists reject). The treatment of the second develops the reply to Searle and Kripke on the no-gap objection, which will be the main pivot for the engagement.

Frankish's closing claim, in §4 (Frankish 2016, pp. 31–32), is modest. The illusion problem is hard, but not impossibly so. It is the kind of problem cognitive science is in principle equipped to address, once the realist commitment is dropped.

5.2.2 The argument reconstructed

For clarity, the central argument of the article can be reconstructed in numbered form.

1. Phenomenal consciousness, as standardly conceived, is the having of qualitative properties — ineffable, intrinsic, private, immediately apprehended — that determine what it is like to undergo experiences.
2. There are three responses to the question of how phenomenal consciousness exists: radical realism, conservative realism, and illusionism.
3. Radical realism faces the threat of epiphenomenalism and requires theoretical innovation of a kind the rest of science does not require.
4. Conservative realism is unstable: its reductive explanations of phenomenal properties either covertly become illusionist or fail to exhibit the qualitative residue they were meant to explain.
5. Therefore the realist options face severe internal difficulties.
6. Phenomenal zombies, lacking phenomenal properties, would make all the same reports and form all the same beliefs about consciousness that we do. So our reports and beliefs about consciousness do not require phenomenal properties for their explanation.
7. By an abductive inference from anomalousness, if a property resists physical explanation, is detectable only from a single perspective, and our beliefs about it can be fully explained without invoking it, the simplest hypothesis is that the property is illusory.

8. Therefore phenomenal properties are illusory. What is to be explained is not their existence but the introspective appearance of them.
9. The illusion problem replaces the hard problem as the central explanandum of a theory of consciousness.
10. The illusion problem is hard but tractable: cognitive science can in principle address it.

Premises 6 and 7 do most of the argumentative work. Premise 6 is the explanatory-dispensability move: phenomenal properties are not needed to explain our claims about phenomenal consciousness. Premise 7 is the abductive-inference-from-anomalousness move: an unneeded posit that resists physical explanation is more economically treated as illusory. Together they license the move from "phenomenal properties resist physical explanation and our beliefs about them have alternative explanations" to "phenomenal properties are illusory." The engagement below does not contest the abductive form of the argument. It presses on the conception of the target the premises share.

5.2.3 The load-bearing premise

Frankish's argument, on examination, depends throughout on a particular conception of the target it eliminates. The conception can be made explicit.

> Frankish's eliminative argument depends on the target explanandum being phenomenal property realism: the claim that experiences instantiate intrinsic, qualitative, private phenomenal properties of the sort realists take to generate the hard problem.

The conception is operative throughout the article. It is operative when Frankish characterises phenomenal properties as "simple, ineffable, intrinsic, private, and immediately apprehended" (Frankish 2016, p. 12). It is operative when he defines the quasi-phenomenal property as "a non-phenomenal, physical property that introspection typically misrepresents as phenomenal" (Frankish 2016, p. 14) — the misrepresentation is misrepresentation *as having* the intrinsic-qualitative features the realist attributes. It is operative in the structural arguments against radical and conservative realism, which proceed by showing that neither can do justice to the realist conception of phenomenal properties: radical realism cannot make the properties causally potent while preserving their qualitative character; conservative realism cannot make them genuinely qualitative while preserving their physical tractability (Frankish 2016, pp. 21–23). It is operative in the no-gap reply, where Frankish's response to Searle and Kripke is that the introspective representation of a greenish experience can be non-veridical because the representation itself does not have qualitative content (Frankish 2016, pp. 27–28).

The conception is reasonable. It is the conception most realists in fact deploy, and Frankish is not straw-manning when he uses it. The point of identifying it as the article's load-bearing premise is not to suggest that Frankish has missed a more subtle realist position. It is to mark, precisely, what Frankish has and has not done by the end of his argument.

What Frankish has done, if the argument succeeds, is to eliminate phenomenal property realism. Phenomenal properties, conceived as intrinsic-qualitative-private features of experiences, are not instantiated. The introspective seeming of such properties is to be explained by representational mechanisms whose content is determined by non-phenomenal causal-functional factors. The hard problem dissolves into the illusion problem.

What Frankish has not done is something the book has been concerned to mark.

> Even if Frankish successfully eliminates phenomenal properties so construed, it does not

> follow that the field in which experiences, introspective representations, misrepresentations, and seemings are intelligible has been eliminated or explained.

He has not addressed the field within which the introspective representation, the seeming, the misrepresentation, and the explanatory practice itself are intelligible as the kinds of acts and states they are. The illusionist programme operates within that field throughout. The field is not a phenomenal property in the realist sense. It is not the kind of thing Frankish's eliminative argument is structured to address. And it is not addressed.

That is the pressure point. The section develops it through Frankish's own two clearest acknowledgments of what survives the elimination: his treatment of introspective subjectivity in §3.1, and his treatment of the no-gap objection in §3.2.

5.2.4 The pressure points at §§3.1 and 3.2

The textual hinge for the pressure point lies in §3.1 of Frankish's article (Frankish 2016, pp. 25–27), in his response to David Chalmers's acquaintance theory. In that response, Frankish introduces the distinction between introspective subjectivity and intrinsic subjectivity already noted. *Introspective subjectivity* is the introspective awareness of one's experiences generated by representational mechanisms; illusionists accept it. *Intrinsic subjectivity* is a subjective dimension not produced by introspective mechanisms but arising from what we are, considered as physical systems; illusionists reject it as a relic of acquaintance theory's anti-physicalist commitments.

Frankish's position is that introspective subjectivity does all the work that needs doing. There is no need to posit a further dimension of intrinsic subjectivity. The appearance that there is one is to be explained by the representational mechanisms that produce the seeming. Acquaintance theory, by contrast, posits intrinsic subjectivity as a non-representational relation between a subject and a phenomenal property — and Frankish rejects this both for the structural reasons noted earlier and on grounds of theoretical economy. He writes: intrinsic subjectivity, were it real, would be "a shadowy companion of physical systems, and we could imagine any object possessing it, as panpsychists do" (Frankish 2016, p. 27). The argument by which Frankish dismisses intrinsic subjectivity is that it does no explanatory work; a creature's reports of what its experiences are like will be the product of introspective mechanisms and will thus manifest introspective subjectivity only.

The position is internally consistent and, given the conception of phenomenal properties Frankish is operating with, well-motivated. The question is whether the position addresses what the book has been concerned with. It does not.

The book's argument is not that there is a non-representational relation between a subject and phenomenal properties of the kind Chalmers's acquaintance theory posits. The book has not argued for acquaintance theory. The book's argument is structural: the field within which introspective representations occur, within which they are intelligible as representations, within which the question of whether they accurately represent their objects is a question with content, is not itself a property — qualitative, intrinsic, or otherwise. The field is not a candidate for elimination by an argument that targets reified properties of experiences.

Frankish's denial of intrinsic subjectivity, on this reading, is correct as a denial of a particular realist posit and incorrect as a denial of what the book has been calling awareness. The two are different. The realist posit is a phenomenal property of experiences, conceived as intrinsic and immediately known. The book's awareness is the field in which experiences, properties, representations, and representational-content-questions are intelligible as the items they are. The former can be coherently

eliminated; the latter cannot, because the elimination would itself occur within the field whose existence the elimination would purport to deny.

The argument from explanatory dispensability, which is decisive against intrinsic subjectivity construed as a property, does not bear on the field as the book has characterised it. The field is not a "shadowy companion" of physical systems whose addition would or would not make a difference to their explanation. It is not the kind of thing that figures in a list of items whose explanatory work can be assessed. It is the standing in which lists, items, explanations, and assessments of explanatory work are intelligible as what they are.

This pressure becomes sharper in §3.2 (Frankish 2016, pp. 27–29), where Frankish replies to the no-gap objection associated with John Searle and Saul Kripke. The objection runs that for qualitative states, there is no gap between seeming and being. If it seems to me that I am having a greenish experience, then I am having a greenish experience. The reply, in Frankish's hands, is that "seems" equivocates. If "seems to have a greenish experience" means "introspectively represents oneself as having a greenish experience," there is a clean gap. The introspective representation can be non-veridical. Its content is determined by non-phenomenal causal-functional factors, not by any qualitative character of the representation itself.

The reply is forceful within Frankish's conception of phenomenal properties. The objection presupposes that the introspective representation of a greenish experience must itself be greenish — must itself have the qualitative character of the experience it represents. Frankish denies this on functional-representational grounds, and his denial is well-motivated within the realist conception of phenomenal properties he is targeting. The objector who wants to maintain the no-gap claim must hold that the introspective representation has its own qualitative content, and Frankish has independent arguments against that move.

What the reply does not address is the further question. In what does the very intelligibility of the representation as a representation, the seeming as a seeming, the misrepresentation as a misrepresentation, consist?

That question is not the no-gap objection in its Searle–Kripke form. It is not the claim that the seeming and the being collapse for qualitative states. It is the structural observation that the entire vocabulary Frankish deploys in his reply — "introspective representation," "non-veridical," "content determined by causal-functional factors," "the misrepresentation," "the seeming" — operates within a field of intelligibility whose status is left untouched by the elimination of phenomenal properties. The introspective representation is intelligible as a representation. The non-veridicality is intelligible as a candidate descriptive feature of the representation. The causal-functional determination is intelligible as the kind of relation it is supposed to be. None of these intelligibilities is itself a phenomenal property in the realist sense, and none of them is eliminated by Frankish's argument. The field within which they have their standing is what the book has been concerned with throughout.

The position the book takes can be stated cleanly. Frankish is persuasive against reified phenomenal properties — those conceived as intrinsic, qualitative, private features of experiences, in the way realists standardly conceive them. The question the book presses is whether explaining introspective misrepresentation explains the field in which misrepresentation, seeming, and explanation themselves are available. The answer the book has been working toward is that it does not, and the present section has been concerned to show how this bears on Frankish's specific moves at §§3.1 and 3.2 of the 2016 paper.

5.2.5 What the critique establishes, and what it does not

The critique of Frankish here is structural, not metaphysical in the realist sense, and its scope should be stated carefully.

The critique establishes that strong illusionism's eliminative move on phenomenal property realism does not, by itself, deliver the further metaphysical conclusion that nothing relevant to a theory of consciousness survives the elimination. The field within which the elimination is articulated, within which introspective representations are intelligible as representations, within which the seeming is intelligible as a seeming, is not a phenomenal property in the realist sense and is not eliminated by Frankish's argument.

The critique does not establish that Frankish's elimination of phenomenal property realism is incorrect. On the contrary, the position defended in the book is compatible with that elimination. Reified phenomenal properties — conceived as intrinsic, qualitative, private features of experiences — are not the field, and Chapter 4's identification arguments have already declined to identify awareness with any determinate property, phenomenal properties of the realist kind included.

The critique does not establish that the illusion problem is misconceived. On the contrary, the illusion problem is a legitimate cognitive-scientific question. How introspective representations are produced, why they have the content they have, why they are so vivid and so resistant to correction — these are real questions that admit of empirical and theoretical investigation. The section has nothing to add to those investigations.

The critique does not establish that Frankish has missed phenomenal consciousness in the realist sense. Frankish is precisely the philosopher who has argued, at length, against the realist sense. The critique respects the argument and adds: even granting it, a question remains about the field within which the argument occurs.

The critique does not establish that phenomenal properties of the realist kind exist. They are not the field. Whether the realist conception of phenomenal properties has any defender who could mount a successful response to Frankish is a question the book takes no position on.

What the critique establishes is one thing, narrowly. It is that an eliminative argument against reified phenomenal properties does not exhaust the philosophical work a theory of consciousness has to do, because the elimination operates within a field whose status is left untouched by it. The illusion problem is real; the field-question is also real; the two are not the same; and the cognitive-scientific tractability of the first does not bear on the philosophical standing of the second.

5.3 Chalmers on the meta-problem

5.3.1 The argument

David Chalmers's "The Meta-Problem of Consciousness" was published in 2018 in the *Journal of Consciousness Studies*. The paper introduces what Chalmers calls the meta-problem of consciousness and develops a research programme around it. It does not advance a single solution to the meta-problem. Its first half (§§1–2) defines the meta-problem and surveys thirteen candidate explanations of how it might be solved. Its second half (§§3–7) considers what the meta-problem implies for theories of consciousness, presents debunking arguments from the meta-problem to illusionism, considers six broad reactions to the meta-problem, and offers an argument against strong illusionism. The conclusion identifies the meta-problem as a tractable research project for everyone, illusionists and realists alike.

The meta-problem itself is the problem of "explaining why we think that there is a problem of consciousness" (Chalmers 2018, p. 6). To a second approximation it is the problem of explaining our "problem reports" — verbal reports, judgments, and dispositions to make judgments that reflect our sense that consciousness poses a special problem (Chalmers 2018, pp. 7–8). The hard problem of consciousness is the problem of explaining how physical processes give rise to phenomenal experience. The meta-problem is the problem of explaining why we think the hard problem is hard.

The distinction matters because, as Chalmers stresses, the meta-problem is "strictly speaking one of the easy problems of consciousness" (Chalmers 2018, p. 7). It concerns a behavioural fact — that human beings systematically produce problem reports about consciousness — which is in principle addressable by the standard methods of cognitive science. It can be approached through neural and computational mechanism, through psychology and linguistics, through experimental philosophy. The meta-problem opens a large empirical and philosophical research programme that does not require commitment to any particular metaphysics of consciousness.

The reason for this tractability is a constraint Chalmers makes explicit in §1.2 (Chalmers 2018, pp. 16–18). He introduces what he calls the *topic-neutrality* constraint on explanations of problem reports. An explanation of problem reports counts as a solution to the meta-problem only if it is given in terms that do not mention consciousness or its cognates — qualia, awareness, subjectivity. Physical and functional explanations meet this constraint. So do representational, rational, historical, and structural explanations, where these do not invoke consciousness. The constraint allows non-reductionist theorists to participate in the research programme; even those who think consciousness plays a causal role in generating problem reports can give topic-neutral structural explanations of those reports without thereby denying consciousness.

Topic-neutrality is the methodological hinge of the entire paper. It is what makes the meta-problem an easy problem in Chalmers's sense. Without it, an explanation that appealed to consciousness in order to explain reports about consciousness would not constitute progress on the meta-problem; the explanandum would have shifted back to the hard problem. With topic-neutrality in force, the research programme has a definite shape: produce a topic-neutral account of why problem reports are made, why they have the structure they have, why they are widespread, and why they are resistant to revision.

The §2 survey of thirteen candidate explanations (Chalmers 2018, pp. 19–35) operates within the topic-neutrality constraint. Each proposal is a candidate for how problem reports could be produced without appealing to consciousness in the explanation. The proposals include: introspective models that misrepresent simpler cognitive states; phenomenal concepts with special properties; primitive quality attribution; primitive relation attribution; the user illusion (Dennett); historical and cultural explanations. Chalmers identifies the first seven as most promising and his own preferred line as a combination of primitive relation attribution and acquaintance (Chalmers 2018, p. 39). The proposals are not all mutually exclusive; a complete solution may combine several.

The §3 "meta-problem challenge for theories of consciousness" (Chalmers 2018, pp. 35–40) is the most consequential methodological move in the paper. Chalmers argues that any theory of consciousness that proposes a mechanism M as the basis of consciousness must explain how M plays a central role in generating problem reports. He applies the challenge to integrated information theory, global workspace theory, higher-order thought theory, biological theories, quantum theories, and panpsychism. In each case, he asks whether the proposed basis of consciousness can be shown to do explanatory work in generating phenomenal intuitions; in each case, the answer requires substantial further argument. The challenge is itself topic-neutral: it asks how the proposed mechanism, characterised in topic-neutral terms, explains the topic-neutrally characterised reports.

The §4 six-reactions taxonomy (Chalmers 2018, pp. 40–44) maps the possible theoretical positions in light of the meta-problem. Three are non-reductionist: meta-problem nihilism (no topic-neutral solution exists); correlationism (consciousness correlates with meta-problem processes but plays no causal role); realizationism (consciousness realises meta-problem processes). Three are reductionist or illusionist: strong illusionism (phenomenal consciousness does not exist); lower-order weak illusionism (consciousness is the lower-order state that the meta-problem processes attribute primitive properties to); higher-order weak illusionism (consciousness is the higher-order introspective state itself). Chalmers's own view is realizationism plus a primitive-relation-attribution account.

The §5 debunking arguments (Chalmers 2018, pp. 44–49) deserve particular attention because Chalmers's treatment anticipates much of the dialectical terrain the present engagement enters. The debunking argument runs: there is a correct explanation of our beliefs about consciousness that is independent of consciousness; if so, those beliefs are not justified; therefore those beliefs are not justified. Chalmers gives three replies. The most important for the present engagement is the third: a distinction between four senses of "independent." Solving the meta-problem yields *descriptive* independence (the explanation does not mention consciousness) and arguably *modal* independence (the elements of the explanation could obtain without consciousness, on a zombie scenario). It does not yield *causal* independence (consciousness may play a causal role in the explanation's elements) or *constitutive* independence (consciousness may constitute elements of the explanation). Premise 2 of the debunking argument, Chalmers argues, requires causal and constitutive independence to be plausible. The realist can deny both, particularly on realizationism. The debunking argument therefore does not go through.

The §6 discussion (Chalmers 2018, pp. 49–53) returns to the question of what kind of illusionism follows from the meta-problem, and the §7 argument against strong illusionism (Chalmers 2018, pp. 53–56) presents Chalmers's own Moorean argument: people feel pain, strong illusionism denies that they do, therefore strong illusionism is false. The argument is not the centrepiece of the paper, but it makes clear that Chalmers is a realist about phenomenal consciousness. His engagement with illusionism is methodologically friendly — he writes that he has "more sympathy with [illusionism] than with most materialist views" (Chalmers 2018, p. 9) — but it is not endorsement. Chalmers's preferred position is realism plus realizationism.

The closing §8 (Chalmers 2018, pp. 56–57) marks the difficulty of the dialectic. Both illusionism and realism, in light of the meta-problem, face charges of absurdity — illusionism for denying the obvious, realism for treating the obviousness as a fortunate coincidence between the metaphysics and the meta-problem processes. Chalmers takes the realist side but recognises that the realist's position needs further work.

5.3.2 The argument reconstructed

The central argumentative arc of the paper can be reconstructed in numbered form.

1. There is a hard problem of consciousness: why and how do physical processes give rise to conscious experience.
2. There is also a meta-problem: why we make problem reports about consciousness — judgments, reports, and dispositions to make them that reflect our sense that consciousness poses a special problem.
3. The meta-problem is strictly an easy problem, because problem reports are a behavioural fact that can in principle be explained by the standard methods of cognitive science.

4. To make the meta-problem tractable, explanations of problem reports must be given in topic-neutral terms — terms that do not mention consciousness or its cognates.
5. Many candidate explanations meeting the topic-neutrality constraint can be offered, and a combination of several is most promising.
6. Any theory of consciousness must meet the meta-problem challenge: it must explain how its proposed basis for consciousness plays a central role in generating problem reports.
7. The meta-problem can be leveraged into a debunking argument against realism, but the realist can resist the argument by appealing to causal or constitutive dependence of problem reports on consciousness — most attractively, on realizationism.
8. Strong illusionism is required if illusionism is to dissolve the hard problem, but strong illusionism is false because people feel pain.
9. The most promising line for both realists and illusionists is to integrate consciousness with meta-problem processes — realizationism on the realist side, full strong illusionism on the illusionist side.
10. The meta-problem is a tractable research project either way.

Premises 3 and 4 do most of the work for the engagement that follows. Together they constitute the methodological character of the research programme.

5.3.3 The load-bearing premise

The load-bearing premise for the present engagement can be stated explicitly.

> Chalmers's meta-problem is methodologically constituted by the requirement that problem reports and problem intuitions be explained, or be explanable in principle, in topic-neutral terms — terms that do not mention consciousness or its cognates.

The premise is operative throughout the paper. It is announced at §1.2 (Chalmers 2018, pp. 16–18) and never relaxed. It governs the survey of candidate solutions in §2, each of which is assessed as a topic-neutral proposal. It governs the meta-problem challenge in §3, which asks whether each theory's proposed mechanism explains problem reports under a topic-neutral construal. It governs the six-reactions taxonomy in §4, where each position is characterised by the relation between consciousness (mentioned) and meta-problem processes (topic-neutral). It governs the debunking arguments in §5, where the four senses of "independent" are senses of independence between topic-neutral explanation and consciousness.

The premise is not a mere stylistic preference. It is what makes the meta-problem an easy problem in Chalmers's sense — a problem addressable by the mechanisms-and-functions methods that cognitive science has refined. Without topic-neutrality, an explanation of problem reports could appeal to consciousness in order to explain why we make reports about consciousness, and the methodological progress of the research programme would dissolve back into the hard problem.

The premise is reasonable. It is well-motivated within the research programme it serves. The point of identifying it as the article's load-bearing premise is not to suggest that Chalmers should have framed the meta-problem differently. He should not have. The point is to mark, precisely, what the meta-problem's solution can and cannot deliver.

What a solution to the meta-problem would deliver, if it succeeded, is a topic-neutral account of why problem reports are made. The account would explain why human beings produce judgments about consciousness having the structure those judgments have. It would explain why these judgments are widespread, why they are resistant to theoretical revision, why they have the features they have. It would constitute a substantial cognitive-scientific achievement. Chalmers is right that such an achievement is in principle attainable.

What a solution to the meta-problem would not deliver is something the book has been concerned to mark.

> The field is not the kind of explanandum a topic-neutral account of problem reports is designed to specify.

A topic-neutral explanation of problem reports can succeed on Chalmers's own terms while leaving untouched the field within which reports, intuitions, explanations, and topic-neutral descriptions are intelligible as such. That is the pressure point. The section develops it through Chalmers's own framing of what topic-neutrality is and what its limits are.

5.3.4 The pressure point developed

The first thing to mark is what the pressure point is not. It is not an objection to the meta-problem research programme. The programme is legitimate and Chalmers's case for it is convincing. It is not a denial of topic-neutral explanation. Topic-neutral explanation is exactly the kind of explanation cognitive science is structured to produce. It is not an attempt to debunk Chalmers's realism. Chalmers's own treatment of the debunking arguments in §5 already establishes that the meta-problem does not by itself debunk realism. The pressure point is structural and adjacent to Chalmers's own argument.

The book's argument and Chalmers's argument converge on one point and diverge at another. They converge in resisting the inference from a topic-neutral solution of the meta-problem to the elimination of phenomenal consciousness as a real explanandum. Chalmers makes the resistance through realizationism — the view that consciousness realises meta-problem processes and thereby plays a causal role even where the topic-neutral characterisation of those processes does not mention it. The book's resistance is differently located. The book argues that the field within which the topic-neutral characterisation is itself articulated is not what the topic-neutral apparatus is structured to address. The two resistances are compatible. The book has no quarrel with realizationism as a metaphysics of how consciousness relates to meta-problem processes. Its argument is about what the meta-problem solution would leave untouched even were realizationism true.

Consider what a topic-neutral solution would consist in, in its most articulated form. It would describe, in physical, functional, representational, or structural terms, the mechanisms by which problem reports are produced. It would identify the cognitive architecture that gives rise to the dispositions Chalmers calls problem intuitions. It would show how children acquire these dispositions, how they vary across cultures, how they relate to other intuitive judgments, how they connect to introspective self-modelling, how they are produced by the brain. It would, in short, deliver a complete cognitive-scientific account of why we say what we say about consciousness.

Now consider what the account would presuppose for its own intelligibility. The account is itself a body of claims. It picks out items — cognitive mechanisms, introspective self-models, dispositional states — and characterises them. It identifies relations between them. It makes inferences from observations to mechanisms, from mechanisms to predicted reports, from predicted reports to actual reports. It evaluates competing topic-neutral explanations against one another. It distinguishes correct from incorrect

topic-neutral descriptions of the same set of phenomena. Each of these intellectual operations occurs within the field of intelligibility the book has been marking. The account does not arrive at the field by topic-neutral characterisation, because the field is not one of the items that topic-neutral characterisation distinguishes.

This is not a complaint that the account presupposes consciousness in some general sense. Chalmers explicitly addresses this kind of worry. He acknowledges (Chalmers 2018, p. 18) that on some views, all judgment involves consciousness or all meaningful language is grounded in consciousness. His response is that for purposes of the meta-problem, problem intuitions can be reconstrued as "quasi-phenomenal judgments" that do not require consciousness for their formulation, or as mere propensities to make certain noises and inscriptions. The response is well-motivated for its intended purpose. It allows the research programme to proceed without prejudging the metaphysics of consciousness.

What the response does not address is whether the field, characterised as the book characterises it, is one of the items the topic-neutral move treats as bracketable. The book has not argued that consciousness is required for all meaningful language or for all judgment. It has argued that there is a structural distinction between contents and the field within which contents are determinate as contents. Contents include cognitive mechanisms, introspective self-models, problem reports, judgments, intuitions, dispositions to make noises and inscriptions, and the topic-neutral explanatory accounts that operate on these things. The field is what makes their being-as-content possible. The topic-neutral apparatus operates on contents. It does not address the field, not because the field is hidden, but because the field is not one of the items the apparatus is structured to specify.

Chalmers's §3 meta-problem challenge brings the point out particularly clearly. The challenge asks: for any theory of consciousness that proposes mechanism M as the basis of consciousness, how does M play a central role in generating problem reports? The challenge is well-targeted. Each candidate basis for consciousness — integrated information, global broadcast, higher-order thought, biological structure, quantum process, micro-consciousness — must be shown to explain why we make the reports we make. The challenge is methodologically rigorous and the book has no objection to it.

But the challenge is structured for theories that take consciousness to be a content-level mechanism or substrate of some kind. It is structured for theories that locate consciousness in or alongside the inventory of items cognitive and physical description handles. For such theories, the meta-problem challenge is a real constraint. For a position that takes awareness to be the field within which the inventory is intelligible, the challenge does not apply in the same form. The book is not proposing a mechanism that must be shown to play a role in generating problem reports. It is proposing a structural distinction between what generates reports and the field within which the generation is intelligible as generation. The two are not in the same explanatory market.

The §5 debunking discussion repays a closer look. Chalmers's four-sense distinction between descriptive, modal, causal, and constitutive independence is the most careful treatment in the literature of what a meta-problem solution would and would not establish about consciousness. He concludes that solving the meta-problem yields descriptive and modal independence between the topic-neutral explanation and consciousness, but does not yield causal or constitutive independence — and that the realist can therefore resist the debunking move by holding that consciousness plays a causal or constitutive role in the meta-problem processes.

The book agrees with this resistance and adds a structurally adjacent point. Even granting full causal and constitutive integration of consciousness and meta-problem processes — granting Chalmers's preferred realizationism — the field within which the integration is itself articulated as integration is left untouched. The realizationist framing is articulated topic-neutrally; it identifies meta-problem pro-

cesses by their topic-neutral structure and then claims that consciousness realises them. The framing is a theoretical achievement, not a confusion. But the framing operates within the field. The realizationist move protects the realist from the debunking argument at the level of phenomenal beliefs; it does not bear on whether the field within which phenomenal beliefs and meta-problem processes are both intelligible has been addressed.

A short paragraph on vocabulary, before the closing of this section. Chalmers's §4 vocabulary maps in places onto Frankish's. Chalmers calls the topic-neutrally characterised lower-order states "lower-order meta-problem states"; he notes (Chalmers 2018, p. 41) that these are what Frankish calls "quasi-phenomenal states." The two vocabularies pick out, at the relevant level, the same kind of thing: a content-level state that is the substrate for the introspective representation of phenomenal properties. The two engagements of this chapter therefore operate on a shared landscape, even though the two philosophers occupy different positions on it. Frankish is a strong illusionist; Chalmers is a realist who declines illusionism. What unites their projects, for the purposes of the present chapter, is the structural commitment to a research programme that operates on content-level architecture — Frankish on the illusion-producing representational architecture, Chalmers on the topic-neutrally characterisable processes that produce problem reports. Both projects are real research programmes. Neither addresses the field.

The position the book takes on Chalmers can be stated cleanly. Chalmers is right that the meta-problem is an important research project. He is right that it can be addressed in topic-neutral terms without thereby debunking realism about phenomenal consciousness. He is right that realizationism is a promising direction for realists. The book agrees with all of this. What the book adds is that the field within which the meta-problem's research programme is itself articulated is not what the topic-neutral apparatus is structured to address.

5.3.5 What the critique establishes, and what it does not

The scope of the critique should be stated carefully.

The critique establishes that a successful topic-neutral solution to the meta-problem, on Chalmers's own characterisation of what such a solution would consist in, would not bear on the field-question the book has been pursuing. The field is not what the topic-neutral apparatus is structured to specify. Topic-neutrality is methodologically required for the meta-problem to be an easy problem, but it is also the methodological feature that marks the limit of what such a solution can address.

The critique does not establish that Chalmers's meta-problem is misconceived. On the contrary, the meta-problem is well-defined, the research programme is tractable, and the topic-neutrality constraint is well-motivated for the research programme's purposes. Cognitive science can in principle make substantial progress on the meta-problem, and the section has nothing to add to that work.

The critique does not establish that Chalmers is wrong to be a realist. The book agrees with him on this point. The book also has no objection to realizationism as a way of marking the relation between consciousness and meta-problem processes, although the book is not committed to that metaphysical thesis.

The critique does not establish that Chalmers has failed to take account of awareness as the book characterises it. He has not addressed the field in the book's sense, but he is not concerned with it; the meta-problem operates at a different level. The critique respects this and adds that the difference of level is itself the relevant philosophical observation.

The critique does not establish that the field is something hidden or mysterious that the meta-problem cannot reach by oversight. The field is not hidden. It is the standing in which the meta-problem's research project, its topic-neutral apparatus, its candidate solutions, its debunking arguments, its six-reactions taxonomy, and its conclusions are intelligible as the kinds of acts and states they are. It is presupposed by the apparatus, not concealed from it.

What the critique establishes is one thing. It is that the meta-problem's research programme, however successful, leaves the field-question untouched — not because the meta-problem fails on its own terms but because the field is not what topic-neutral apparatus is structured to address. The meta-problem is a real problem; the field-question is a different real problem; the topic-neutral tractability of the first does not bear on the philosophical standing of the second.

5.4 What the chapter establishes

The two engagements of this chapter have addressed the most articulate contemporary reframings of the hard problem. Frankish's illusionism reframes it as the illusion problem — the problem of explaining why experiences seem to have phenomenal properties when, on Frankish's view, they do not. Chalmers's meta-problem reframes it as the problem of explaining why we make problem reports about consciousness. The two projects are distinct in their metaphysical commitments. Frankish denies phenomenal consciousness; Chalmers affirms it. But they share a methodological structure: each takes an aspect of the hard problem and recasts it as a content-level research question that cognitive science is in principle equipped to handle.

The chapter has been concerned with this shared methodological structure. The engagement with Frankish argued that strong illusionism's elimination of phenomenal property realism is not contested by the book's argument, but that the further metaphysical conclusion — that nothing relevant survives the elimination — relies on a frame the book has independent reason to question. The engagement with Chalmers argued that the meta-problem's topic-neutral research programme is well-motivated and tractable on its own terms, but that its successful solution would not address the field within which the topic-neutral apparatus is itself articulated. The two engagements are not the same. They press on different commitments. They produce different conclusions about what the chapter's targets have and have not accomplished. But they converge on a structural observation.

The observation is this. Each project produces a legitimate cognitive-scientific explanandum. The illusion problem is a real problem. The meta-problem is a real problem. Both can be addressed by methods that have been refined in the cognitive sciences over the past several decades. Neither requires philosophical commitments the book contests. The book's argument is not that the reframings are mistakes. The argument is that the reframings, however successful, do not bear on the field-question.

The field-question, as the book has been formulating it, is the question of awareness understood as the field of intelligibility — the standing within which items, judgments, introspective states, cognitive-scientific theories, topic-neutral characterisations, and philosophical conclusions are intelligible as the kinds of acts, states, and claims they are. The field is not a phenomenal property of mental states. It is not an introspectable feature. It is not a content-level structure that introspection misrepresents or accurately represents. It is not something a topic-neutral apparatus is designed to capture. It is the standing within which all such items and apparatus have whatever standing they have.

When Frankish's research programme succeeds in characterising the introspective mechanisms that generate the illusion of phenomenal properties, what it has characterised is content-level architecture. When Chalmers's research programme succeeds in characterising the meta-problem processes that pro-

duce problem reports, what it has characterised is also content-level architecture. The two architectures are not identical; Chalmers's quasi-phenomenal states and his higher-order introspective states are distinct from Frankish's analogue sensory representations and their misrepresenting introspective mechanisms. But both architectures are inventories of items — items the topic-neutral or representational vocabulary can name, describe, distinguish, compare, model, refine. The field is not one of those items. Neither vocabulary picks it out by mistake or by oversight. The field is not what the vocabularies are structured to pick out.

The chapter's combined result is therefore narrow. It is that the most articulate reframing responses to the hard problem in the contemporary literature do real work at the level they address, and leave a question at a different level untouched. That different-level question is the one the book has been concerned with from the start. The chapter does not press it further here; the previous chapters have already developed the structural distinction the chapter has been using, and the chapters that follow will continue developing what follows from that distinction.

A short final word on what the chapter does not claim. It does not claim that Frankish or Chalmers have made errors by their own lights. Both have produced careful and important work. It does not claim that the reframing strategies are exhausted by the two engagements here; other contemporary reframings exist and would face structurally similar observations. It does not claim that the cognitive-scientific projects opened by the illusion problem and the meta-problem will fail. They may succeed completely on their own terms. What the chapter claims is that their success, if achieved, would not bear on the field-question. The field-question is a different question, and the chapters that follow continue to develop it.

Chapter 6 — Time, Structure, and Content-Experience

6.1 Why time matters

Chapter 3 addressed derivation; Chapters 4 and 5 addressed forms of identification and reframing. The present chapter takes up a different way in which awareness can be subordinated: the way temporal structure is sometimes invoked as a condition of awareness itself.

The move is familiar. It is sometimes argued that whatever else is true of awareness, it must at least be structured temporally — that awareness must take place in time, that conscious experience must unfold through moments of succession, that the unity of any conscious episode must be synthesised across a temporal sequence. The argument is then made that, since awareness must be temporally structured, temporal structure itself conditions awareness; whatever conditions an item also constrains it; therefore awareness is subordinate to whatever conditions temporal structure has. The shape of the move is to find a transcendental condition that applies to awareness, treat that condition as more basic than awareness, and so subordinate awareness to the condition.

The most articulate and historically influential doctrine in this neighbourhood is Kant's. In the Transcendental Aesthetic of the *Critique of Pure Reason*, Kant argues that time is the form of inner sense — the transcendental condition under which the manifold of inner intuition is necessarily ordered. The doctrine has the philosophical weight it does because it is internally rigorous, because it is paired with a careful distinction between appearance and thing in itself, and because its terms have shaped two centuries of debate about the structure of consciousness, mind, and self.

The present chapter engages Kant's doctrine directly. The engagement has the same shape as the chapter's previous opponent engagements: state the doctrine fairly, identify the load-bearing premise, locate the structural pressure point, and state what the critique does and does not establish. The chapter is not anti-Kantian. It grants Kant's account of time as the form of inner sense within its proper jurisdiction. It does not deny temporal experience, the empirical reality of time, the structure of inner sense, the unity of apperception, or any of the further apparatus Kant develops to articulate the transcendental conditions of human experience. What it adds is a structural observation: awareness, as the book has been developing it — the field of intelligibility within which any item is intelligible as an item — is not an inner appearance, not an item in the temporal manifold, and not the self as it appears to itself under inner sense. Kant's account governs the structure of inner appearance. The field is not what such a structure conditions.

The chapter is structured in four further sections. §6.2 reconstructs Kant's doctrine of time. §6.3 takes up the B-Deduction's account of inner sense, apperception, and self-affection. §6.4 develops the pressure point. §6.5 states what the critique does and does not establish.

A brief note before turning to Kant himself. The chapter's engagement is with Kant's text. Contemporary commentators on Kant — Guyer, Strawson, Allais, Longuenesse, Onof, and others — have refined and contested the interpretation of the relevant passages in important ways.[3] The chapter does not enter those interpretive disputes. Its engagement is with the doctrine as Kant articulates it, and where the chapter's reading is contested, the contestation is flagged for source-pending future engagement rather than resolved here.

6.2 Kant's doctrine of time

Kant's account of time is developed in the Transcendental Aesthetic, Section II ("On Time"), §§ 4–8 (A30 / B46 to A49 / B73).[4] The doctrine has two correlative sides: time is *transcendentally ideal*, and time is *empirically real*. The two sides are best taken in turn.

Time, on Kant's account, is not "something that would subsist for itself or attach to things as an objective determination" (A32 / B49). It is not a feature of things in themselves, considered apart from the subjective conditions of our intuition. It is not, on Kant's view, even a relation among things in themselves that we abstract from experience. The argument for this claim runs across §§ 4–6 and culminates in the formulation that gives the doctrine its name. Time is "nothing other than the form of inner sense, i.e., of the intuition of our self and our inner state" (A33 / B49). It is the subjective form under which the manifold of inner intuition is necessarily ordered.

The doctrine of transcendental ideality follows directly. If time is nothing other than the form of inner sense, then "if we abstract from our way of internally intuiting ourselves… and thus take objects as they may be in themselves, then time is nothing" (A34 / B50–B51). Time has no application beyond the subjective conditions of human sensible intuition. Kant marks the point with characteristic firmness: time is "merely a subjective condition of our (human) intuition… and in itself, outside the subject, is nothing" (A35 / B51).

The doctrine is not idealist in any sense that would deny the reality of time within experience. Kant pairs the transcendental ideality of time with its empirical reality: time has "objective validity in regard to all objects that may ever be given to our senses" (A35 / B52). The empirical reality of time secures

[3] The interpretive disputes are extensive and continue to be live. P. F. Strawson, *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason* (London: Methuen, 1966), Part Two Chapter II §§1–2, pp. 39–50, distinguishes an "austere" reading of the Aesthetic — on which space and time are necessary structural features of any conception of experience we can render intelligible — from the "transcendental idealist" reading on which space and time are "in us, prior to experience" and the objects of awareness are appearances of unknowable things in themselves. Strawson defends the austere reading and presses on the wider metaphysical apparatus in Part Three Chapter III §3, "The Complications of Transcendental Idealism," pp. 173–177. Lucy Allais, *Manifest Reality: Kant's Idealism and his Realism* (Oxford: Oxford University Press, 2015), defends a moderate metaphysical interpretation on which Kant's idealism is genuinely metaphysical but not phenomenalist, and on which the Aesthetic argument for transcendental idealism does not proceed via the simple synthetic-a-priori-justification route that Stroud and others have attributed to Kant (see Allais, Chapter 1, pp. 2–36, and Chapter 8, pp. 176–204). Paul Guyer and Allen W. Wood, in their Introduction to the Cambridge edition of the *Critique of Pure Reason* (Cambridge: Cambridge University Press, 1998), note that the meaning of transcendental ideality "remains subject to debate" (p. 8) and that the unity of apperception is "obscure and controversial, and continue[s] to generate lively philosophical discussion even after more than two centuries of interpretation" (pp. 9–10). Further commentary on the figurative synthesis, the dual self, and the apperception-centred reading of the Transcendental Deduction lies outside the scope of the present chapter's footnote apparatus. The chapter's structural argument is calibrated to operate at a level adjacent to these interpretive disputes rather than within them: the field-thesis does not depend on adjudicating between austere and metaphysical readings of transcendental idealism, between phenomenalist and moderate metaphysical construals of appearance, or between conceptualist and non-conceptualist accounts of intuition.

[4] Kant's *Critique of Pure Reason* is cited by standard A/B pagination. Translations are from Immanuel Kant, *Critique of Pure Reason*, trans. and ed. Paul Guyer and Allen W. Wood (Cambridge: Cambridge University Press, 1998), unless otherwise noted.

the objective validity of temporal claims about appearances. Every object that can be given to us in experience necessarily stands in temporal relations; every alteration we observe necessarily occurs in time; every claim about the temporal order of events has its sound objective correctness within the domain of appearance. The two doctrines — transcendental ideality and empirical reality — are not in tension; they articulate a single position from two complementary directions.

Kant also extends the inner-sense doctrine to outer appearances by what is sometimes called the "mediate condition" argument. Although space, not time, is the form of outer sense, all representations of outer objects are themselves determinations of the mind, and as such belong to the inner state, which is itself under the formal condition of inner intuition. So time is "an a priori condition of all appearance in general, and indeed the immediate condition of the inner intuition (of our souls), and thereby also the mediate condition of outer appearances" (A34 / B50–B51). All objects of the senses, inner and outer, necessarily stand in temporal relations.

In § 7 (A36–A38 / B53–B55), Kant addresses an objection that bears on the present engagement. The objection runs:

> Alterations are real (this is proved by the change of our own representations, even if one would deny all outer appearances together with their alterations). Now alterations are possible only in time, therefore time is something real. (A37 / B53)

Kant's reply is precise. Time is "certainly something real, namely the real form of inner intuition. It therefore has subjective reality in regard to inner experience" (A37 / B53). The reality of time is the reality of the form under which inner appearances are necessarily ordered. But this reality is not the absolute reality of something that would obtain independently of the subjective conditions of our intuition. Time is "the way of representing myself as object," not an object in itself. The reality of alterations is preserved within the form of inner intuition; the reality is not the reality the objector wants, but it is the reality the phenomenon supports.

The doctrine's most exposed formulation, for the present engagement, is at A38 / B54–B55: "It is nothing except the form of our inner intuition. If one removes the special condition of our sensibility from it, then the concept of time also disappears, and it does not adhere to the objects themselves, rather merely to the subject that intuits them."

Kant's General Remarks on the Transcendental Aesthetic (§ 8, A39–A49 / B55–B73) extend the doctrine in ways that bear less directly on the present chapter but that mark the framework within which the temporal doctrine operates. Throughout, Kant maintains the distinction between appearance and thing in itself, insists that the transcendental ideality of space and time is paired with their empirical reality, and resists confusing the transcendental level with either dogmatic realism (which would treat time as adhering to things in themselves) or empirical idealism (which would deny the reality of objects given in temporal experience).

The doctrine, on Kant's terms, is well-defined and well-defended within the framework of transcendental idealism. The book's engagement does not contest the doctrine within that framework.

6.3 Inner sense, apperception, and self-affection

The doctrine of time as the form of inner sense is developed further in the B-edition Transcendental Deduction, particularly in §§ 24–25 (B150–B158). The further development matters for the present engagement because it makes explicit what is implicit in the Aesthetic: the structure of self-affection, the

distinction between apperception and inner sense, and the role of the figurative synthesis in determining the temporal manifold.

The starting point is what Kant himself calls a paradox. Inner sense, on the Aesthetic doctrine, presents the self to consciousness only as the self appears to itself. We intuit ourselves only as we are internally affected (B153). This appears paradoxical because it seems to require the self to be related to itself passively — the active subject is affected by something, but what is doing the affecting if not the subject itself? Kant's resolution is to distinguish apperception from inner sense and to introduce the figurative synthesis as the determining action.

The paradox is worth pausing on, because Kant's framing of it as a paradox is doing structural work. The active subject is the subject that does the affecting; the passive subject is the subject that is affected. Standard philosophical accounts of affection separate the two — the affecting object is external, the affected subject is internal. Inner sense, on Kant's account, requires a self-affection in which the affecting and the affected are the same subject under different aspects. If the subject is identical with itself across the affecting and the being-affected, the question becomes: in what way can the subject be active in producing the very inner intuitions it then receives? The paradox is not a confusion in Kant's account; it is a feature he names openly and proceeds to resolve through the distinction the next paragraphs articulate. The chapter's engagement requires that the paradox be marked precisely, because the resolution is what introduces the apparatus the chapter will then grant within its jurisdiction.

Apperception, on Kant's account, is the "I think" — the unity of self-consciousness that accompanies all my representations. It is a *spontaneity*: an active, determining capacity, not a passive receptivity. Apperception applies to the manifold of intuitions in general, prior to all sensible intuition, "under the name of the categories" (B154). It is not itself an intuition; the representation "I think" is "a thinking, not an intuiting" (B157).

The notion of spontaneity is doing specific work here. Kant distinguishes the active contribution of the understanding to cognition from the passive reception of sensible content; spontaneity is the active side. Apperception is spontaneity in the strict sense: the I that thinks, the unity that accompanies all representations, the source of the categories' application. Apperception is not given to itself in intuition; it is what makes intuition's contents available as the contents of a single subject's experience. Kant's claim at B157 that "I think" is "a thinking, not an intuiting" is the load-bearing statement of this asymmetry. The thought of myself as a thinking subject does not give me cognition of what I am; it gives me consciousness only that I am. Cognition of what I am requires intuition, and intuition gives me only the self as it appears to itself.

The asymmetry between apperception and inner sense is therefore not symmetrical at all. Apperception, as spontaneity, is the source of the unity that any inner intuition has as the inner intuition of a single subject. Inner sense, as receptivity, is the form under which the manifold of inner intuition is necessarily ordered. The two faculties are not parallel terms in a coordinated description of self-knowledge; they are differently structured contributions to a complex self-relation in which the active subject affects itself, through the figurative synthesis, in order to be given to itself in inner intuition.

Inner sense, by contrast, is "the mere form of intuition, but without combination of the manifold in it" (B154). It contains the temporal form under which the manifold of inner appearances is necessarily ordered, but it does not by itself contain any determinate intuition. The determinacy comes from somewhere else.

That somewhere else is the *figurative synthesis* — what Kant also calls "the transcendental synthesis of the imagination" (B151–B152). The understanding, as spontaneity, "can determine the inner

sense through the manifold of given representations in accord with the synthetic unity of apperception" (B150). The figurative synthesis is the "effect of the understanding on sensibility and its first application... to objects of the intuition that is possible for us" (B152). It is what determines the temporal manifold into the unified content of an inner intuition.

What the figurative synthesis does, structurally, is bring the manifold of inner intuition under the unity of apperception. Inner sense, considered apart from the figurative synthesis, is what Kant calls "the mere form of intuition, but without combination of the manifold in it" (B154) — the temporal form, but not yet a determinate temporal manifold. The figurative synthesis is what determines: it is the action through which the understanding combines the inner manifold into the kind of structured temporal sequence that has unity, and through which the subject can be given to itself as a single subject across that sequence. The structural shape of the move is that an active faculty (the understanding) affects a receptive faculty (inner sense) by performing the synthesis that produces what the receptive faculty then receives. The architecture is reciprocal: understanding and inner sense are not parallel faculties operating independently; they are jointly engaged in the production of inner intuition, with the figurative synthesis as the structural hinge.

For Kant, this reciprocal architecture is required by the very possibility of cognition. The categories, as products of the understanding, acquire their objective reality only through being applied to sensible intuition. The application is mediated by the figurative synthesis, which is what allows the understanding's active contribution to reach sensibility's receptive contents. Without the figurative synthesis, the understanding's categories would remain merely formal, and inner sense would remain merely an unstructured form of receptivity. Together, through the synthesis, they constitute the temporally determinate manifold that any cognition of an object — including the self as an inner object — requires.

The clearest expression of how this works is Kant's geometric example (B154–B155). We cannot think of a line without drawing it in thought. We cannot represent time without describing it as a straight line in inner intuition, attending to the action of the synthesis through which we successively determine the inner sense. Time, in its determinacy as an ordered manifold, is the product of the understanding's affecting of inner sense through the figurative synthesis.

The consequence Kant draws for self-knowledge is direct (B155–B156):

> How the I that I think is to differ from the I that intuits itself (for I can represent other kinds of intuition as at least possible) and yet be identical with the latter as the same subject... is no more and no less difficult than how I can be an object for myself in general and indeed one of intuition and inner perceptions.

The I that thinks (apperception) and the I that appears to itself (inner intuition under temporal form) are distinguished but identified as the same subject. Kant's § 25 sharpens this. In transcendental apperception, "I am conscious of myself not as I appear to myself, nor as I am in myself, but only *that* I am" (B157). The "I think" gives me consciousness *that I am*, but not cognition of *what* I am. Cognition of myself requires intuition, and intuition gives me only the self as appearance, ordered under the temporal form of inner sense.

The structural shape of Kant's account of self-knowledge is therefore distinctive. Kant does not reduce the self to either of its two aspects; he does not say that the self is *really* the thinking subject and only *apparently* the appearing subject, or the other way around. The two are identified as the same subject and distinguished as objects of different kinds of consideration. The thinking subject is the subject as it figures in the unity of apperception, available only to thinking; the appearing subject is the subject as it figures in inner intuition, available only to intuiting. Neither side of the distinction reduces to the

other, and neither side, on Kant's account, gives us cognition of the self as it is in itself.

This is the structural move that blocks the rationalist temptation to derive cognition of the self from the unity of the "I think." The unity of apperception gives us *that* I am, not *what* I am; and the temporal appearance of the self gives us *what* I appear to myself to be, not what I am in myself. The two together constitute the structurally complete account Kant takes himself to be able to give of the self's accessibility to cognition under the conditions of human sensibility and understanding.

The architecture is intricate. Three of its features matter for the chapter's pressure point.

First, time, in its determinacy as an ordered manifold of inner appearances, is *produced* by the figurative synthesis. Time is not a pre-existing structure that the understanding finds in inner sense and operates on. It is the form under which the manifold is ordered when the understanding affects inner sense.

Second, the self as it appears to itself is not the whole story about the self. There is also the self that thinks — apperception, the "I think," the unity of self-consciousness as spontaneity. The thinking self and the appearing self are identified as the same subject but distinguished as objects of different kinds of consideration (one of thinking, one of intuition).

Third, the doctrine of self-affection is not an oddity at the periphery of Kant's system; it is structural to his account of how cognition is possible at all. The categories acquire their objective reality only through the figurative synthesis, and the figurative synthesis operates by determining inner sense. The doctrine of time as the form of inner sense and the doctrine of the understanding's affecting of inner sense are reciprocally articulated.

Taken together, these three features make the architecture a single integrated structure rather than three independent commitments. Time-as-product, the dual self, and structural self-affection are not three things that Kant could in principle have held in different combinations; they are mutually constitutive elements of his account of how the self is given to itself. The chapter's engagement requires that the architecture be granted as integrated, not piecemeal — because the chapter's pressure point is at a level adjacent to the architecture as a whole, not at any one of its individual features.

This is the architecture the book's engagement must respect.[5] The book grants it within its proper jurisdiction. What it adds — and the next section develops — is that the field of intelligibility within which the architecture is itself intelligible as an architecture is not what the architecture's apparatus reaches.

6.4 The pressure point

The pressure point can be stated cleanly.

> Kant's doctrine of time identifies time as the form of inner sense: the transcendental condition under which inner appearances are ordered. The book grants that doctrine within its

[5] The architecture of the B-Deduction §§ 24–25 — the figurative synthesis, the distinction between apperception and inner sense, the self-affection of sensibility by the understanding, the dual self — is among the most extensively contested in contemporary Kant scholarship. Strawson treats the self-affection apparatus and the noumenal-self model as "shaking itself to pieces" (*Bounds of Sense*, p. 177) and argues that Kant's analytic insight is recoverable without that apparatus. Allais argues, conversely, that the apparatus admits of a moderate metaphysical reading on which the dual self is not phenomenalist (*Manifest Reality*, Chapter 1, pp. 2–36; Chapter 8, pp. 176–204). The chapter's structural argument does not rest on adjudicating between these readings or between further interpretations of the figurative synthesis that lie outside the scope of the present chapter's footnote apparatus. Whichever interpretation of the figurative synthesis and the dual self the reader prefers, the level-difference the chapter is marking remains the same: the field of intelligibility within which any such architecture is articulated as an architecture is not what the architecture's apparatus reaches.

> proper jurisdiction. Its pressure is that awareness-as-field is not an inner appearance, not an item in the temporal manifold, and not the self as it appears to itself under inner sense.

Each clause of the pressure point matters and should be developed in turn.

The field is not an inner appearance. Kant's doctrine of time applies to "appearances" in the precise sense his framework gives the term: what is given through sensibility under the forms of intuition. Inner appearances are the appearings of the self to itself, ordered temporally under the form of inner sense. The field of intelligibility within which any such appearance is intelligible as an appearance is not itself an appearance in this sense. The field is not given through sensibility; it does not have a form of intuition; it does not stand in the relation Kant's framework specifies between inner sense and the manifold inner sense presents.

What it means to say the field is "not itself an appearance" is worth stating carefully, because the claim could be misread as a metaphysical posit about something behind appearances. The claim is structural, not metaphysical. Kant's category of appearance picks out one specific kind of thing: what is given through sensibility under the forms of intuition. To say the field is not an appearance is to say that the field is not in that category — not that it is in some further category Kant's framework opposes to appearance. The field is at a different level of consideration than the level at which Kant's appearance/thing-in-itself distinction operates. Kant's doctrine of time is a doctrine about what falls within the category of appearance; the chapter's claim is that the field is not what such a doctrine reaches, not because the field is hidden, but because the field is not the kind of item that a doctrine about appearances would reach.

The field is not an item in the temporal manifold. The temporal manifold, on Kant's account, is the determinate content of inner intuition: the ordered succession of representations, the unified sequence of inner states, the structured array that the figurative synthesis produces by determining inner sense. The field is not such a content. It does not occupy a position in the temporal sequence. It does not stand in relations of succession, simultaneity, or duration to other items. Granting Kant's account of how the temporal manifold is constituted leaves a further question untouched: what makes the manifold intelligible *as* a manifold, what makes the temporal ordering intelligible *as* an ordering, what makes the inner appearances intelligible *as* appearances. These are not questions about further items in the manifold. They are questions about the standing within which the manifold has its standing.

The structural reason for this is precise. The temporal manifold is what the figurative synthesis produces by determining inner sense. The chapter grants the figurative synthesis its constitutive role; it is what makes the manifold a manifold. But the field is not what the figurative synthesis produces or could produce. The figurative synthesis operates on inner intuition to produce a structured temporal sequence; the field is not a structured temporal sequence and is not a candidate output of the synthesis. The synthesis's products are inner appearances ordered temporally; the field is at the level adjacent to those products, the level at which any such product has its standing as the product it is. Kant's apparatus does not address the field because Kant's apparatus is structured to address the production of inner appearances — and the field is not such a product.

The field is not the self as it appears to itself under inner sense. Kant's distinction between the I that thinks and the I that appears under temporal form is one of the most carefully articulated structural insights in the *Critique*. The self as it appears to itself is not the self in itself; Kant insists on this throughout § 25 and elsewhere. The field, on the book's argument, is not the appearing self. But neither is it the self in itself in Kant's sense. The book does not claim, and should not be read as claiming, that the field is what would remain if one abstracted from the subjective conditions of human sensibility — that would be a metaphysical posit about the thing in itself, and the book is not making such a posit. The field is

something different: it is the standing within which Kant's whole distinction between appearance and thing in itself, between the I that thinks and the I that appears, is itself intelligible as a distinction.

The point bears restating because the temptation to locate the field at one or the other side of Kant's distinction is the chapter's most likely misreading. A reader who follows the chapter's denial that the field is the appearing self may take the chapter to be claiming, by elimination, that the field is the self in itself. That move would commit the book to a metaphysical posit about the thing in itself, which the book does not make. The chapter's denial works at a different level: the field is neither the self as it appears nor the self in itself, because the field is not the kind of item that Kant's appearance/thing-in-itself distinction is structured to sort. The distinction operates on items the framework picks out; the field is at the level adjacent to the whole sorting operation. To say this is not to introduce a third item alongside appearance and thing in itself. It is to mark the standing within which the two-item distinction itself has its standing as a distinction.

This last point is where the structural analogy with Kant's own moves matters, and the analogy must be marked carefully. Kant himself distinguishes the level at which the transcendental architecture operates from the level of empirical objects, and he distinguishes within the transcendental architecture between apperception (spontaneity, the I think) and inner sense (receptivity, the form under which the I appears). The book draws a structurally analogous distinction at a different level. The book distinguishes the level at which the field of intelligibility operates from the level at which Kant's whole transcendental architecture is articulated. The field is not what Kant's distinctions sort. It is what makes Kant's distinctions intelligible as distinctions.

The structural analogy operates only at the level of move-shape. Kant marks a level-difference within his framework — between the empirical level at which we make ordinary judgements about objects and the transcendental level at which we make philosophical judgements about the conditions of those ordinary judgements. The chapter marks a further level-difference, at a level adjacent to Kant's whole transcendental level. The shape of the chapter's move resembles the shape of Kant's: in both cases, a distinction is drawn between an operative level and a level at which the operative level's standing is in question. But the resemblance is at the level of shape, not at the level of content. Kant's transcendental level is articulated within a framework of transcendental conditions and synthetic a priori cognitions; the chapter's adjacent level is articulated as the standing within which any such framework is itself intelligible. The shape of the move is similar; what the move produces is different.

The level-difference the chapter marks is not a level Kant denies, nor a level Kant affirms. It is a level Kant does not address, because Kant's questions — about the conditions of cognition, about the form of sensibility, about the synthesis of the temporal manifold — are at the level the chapter then asks the standing-question about. Kant's apparatus operates within its own questions; the chapter's question is one level adjacent. This is the structural relation the chapter has been marking throughout, and it is the structural relation the analogy with Kant's own grant-and-limit move helps to display.

The analogy with Kant should not be pressed into agreement. The book is not extending Kant. Kant did not draw the distinction the book draws; the book is not claiming that Kant secretly held such a distinction or that the field is what Kant should have been pointing to under the name of apperception. The "I think" of transcendental apperception is, for Kant, a specific philosophical posit: a spontaneity, a unity of self-consciousness, the source of the categories' application to the manifold of intuition. The field, on the book's argument, is not a spontaneity in this Kantian sense. It is not a unity that is contributed to experience by the active understanding. It is not a faculty among other faculties. It is the standing within which the very distinction between active and passive, between contributed and given, between transcendental and empirical, is intelligible as a distinction. The analogy is structural; it uses Kant's

distinctions to model the level-difference the book is pointing at without identifying the field with any of the items Kant's apparatus contains.

A short remark on what the pressure point is not. It is not the claim that Kant overlooked something. Kant was not concerned with the field as the book characterises it, and his transcendental architecture is not deficient by his own lights for failing to specify it. The point is structurally adjacent and operates at a level Kant's apparatus does not address. The concession is real; the limit is structural.

One additional Kantian feature should be marked. Kant's reply to the objection from the reality of alteration (A37 / B53) is structurally instructive. Kant grants the empirical reality of alteration and time and pairs the granting with the doctrine of transcendental ideality. The book makes a structurally analogous move at a different level. The book grants Kant's account of time as the form of inner sense — the analog of Kant's granting empirical reality — and pairs the granting with a structural observation about what Kant's doctrine reaches and does not reach. Both moves take the form of careful concession within a wider structural argument; both insist on the reality of what the granting concedes while marking the limit of what the granting establishes.

The shape of this analogy is worth marking precisely. Kant's two-step move at A37 / B53 — granting the empirical reality of alteration and time, marking its transcendental ideality — is a grant-within-jurisdiction paired with a limit-of-jurisdiction. The grant is what makes the empirical claim defensible; the limit is what prevents the empirical claim from being extended into a metaphysical claim about things in themselves. The chapter's two-step move has the same structural shape at a different level. The grant is to Kant's whole transcendental doctrine, taken within its proper jurisdiction; the limit is the recognition that the jurisdiction is the jurisdiction of an account of organised inner appearance, not of an account of the field within which any such organisation is itself intelligible.

What the analogy does not do is import Kant's transcendental idealism into the chapter's argument. The chapter is not claiming that the field is to Kant's transcendental architecture what Kant's things in themselves are to empirical objects. The thing-in-itself is a Kantian posit within Kant's framework; the field is not a posit at all. The analogy is at the level of how a grant-and-limit move is structured, not at the level of what the grant or the limit involves doctrinally. The chapter borrows the move-shape; it does not borrow the metaphysics.

This is the engagement's central move. The book takes Kant seriously and grants what Kant's doctrine establishes within its proper jurisdiction. The book draws a structurally analogous distinction at an adjacent level, without claiming that Kant himself drew it.

6.5 The limit of the Kantian engagement

The Kantian engagement has a narrow scope.

The critique establishes that Kant's account of time as the form of inner sense, however well-defended within the framework of transcendental idealism, identifies time with the transcendental condition under which the manifold of inner appearance is necessarily ordered. The field of intelligibility within which Kant's whole transcendental architecture is articulated is not what such a condition reaches. The book's argument operates at a level adjacent to but distinct from the level at which Kant's doctrine operates.

The critique does not establish that Kant is wrong about time. Within its proper jurisdiction — the conditions under which the manifold of inner intuition is necessarily ordered — the doctrine is internally rigorous and the book takes no position against it. The empirical reality of time, the transcendental ide-

ality of time, the doctrine that time is the form of inner sense, the doctrine of the figurative synthesis, and the doctrine that inner sense is affected by the understanding all remain available within their own framework.

The critique does not deny temporal experience. We do experience time. Conscious episodes do unfold through succession. Memory, anticipation, perception of motion, recognition of change — all of these are real features of our cognitive lives. The chapter has nothing to add or subtract from any of them. The structural observation about awareness-as-field operates at a level that does not bear on whether our experience has temporal structure; it concerns whether the field within which any such structure is intelligible as a structure is itself what such a structure conditions.

The critique does not deny inner sense. The faculty Kant calls inner sense, the form under which we represent ourselves to ourselves, the manner in which we appear to ourselves under temporal ordering — none of this is contested. The book grants Kant's account of how the self appears to itself under inner sense and does not propose a competing account.

The critique does not claim that awareness is a timeless substance, an eternal object, or any other metaphysically inflated item. The field, as the book has been developing it, is not a substance, not a kind of stuff, not an item that would obtain in the absence of time. The field is the standing within which items, substances, durations, and the relations among them are intelligible as the items they are. To deny that the field is conditioned by temporal form is not to claim that the field is timeless in any positive metaphysical sense. It is to mark a structural distinction between what is temporally ordered and the standing within which temporal ordering is intelligible as ordering.

The critique does not identify the field with transcendental apperception or with the "I think." Kant's apperception is a specific philosophical posit: a spontaneity, a unity, the source of the categories' application. The field is not this. The book does not propose to identify its field with anything in Kant's architecture, and any reading that takes the engagement as a thinly veiled appeal to apperception misreads the book.

What the critique establishes is one thing. It is that temporal conditioning of content-experience, as Kant articulates it, does not specify awareness as the field of intelligibility — not because Kant's account is wrong, but because the field is at a level the account does not reach. Time conditions inner appearance; the field is not an inner appearance; granting Kant's account leaves the standing of the field at the level it occupies, neither contested nor specified by what the account establishes.

The chapters that follow extend the same form of argument to one further structural move. Chapter 7 takes up the multiplication of awareness into countable subjects — the move from the plurality of perspectives to the plurality of awarenesses. Chapter 8 consolidates the chapter's positive ontological commitment and engages contemporary panpsychism and cosmopsychism on the question of how widely awareness can be distributed within an inventory of items. The cumulative result, again, is not a new metaphysical posit but a clarified order of dependence: the inventory remains available for description and inquiry; the field within which the inventory has whatever standing it has is not what such description and inquiry specify.

Chapter 7 — Perspectival Plurality

7.1 The problem of plurality

Other people are real. They have bodies, points of view, memories, expressions, lives. Their perspectives are not modifications of mine; they are theirs. When my interlocutor tells me what they saw from where they were standing, what they remember, what they care about, I am not learning more about my own perspective. I am learning what is true from theirs. Any account of awareness that obscures this fact loses the world.

The argument of this book has been a series of structural denials. Awareness is not ontologically derivative from anything determinate enough to play the role the derivation would assign it. Awareness is not identical to any content, content-relation, function, representation, or access architecture. Awareness is not conditioned by temporal structure in the way conditioning relations between content-items operate. Each of these arguments has been narrow and procedural. None of them has been a metaphysical thesis about the unity of awareness or the singularity of the world.

But a reader who has followed the arguments to this point may worry that the cumulative effect is precisely such a metaphysical thesis. If awareness is not derivative, not identical to anything determinate, and not temporally conditioned, what is one supposed to say about the manifest plurality of perspectives? Are the many perspectives so many awarenesses? Or are they appearances within one awareness? Or is the question itself confused?

The chapter argues that the question is confused, but not in a way that denies what the question is responding to. The plurality of perspectives is real. The plurality of bodies, of sensory access points, of memories, of expressive lives, of standpoints from which the world is encountered — all of this is real and structurally describable. The chapter takes for granted everything the phenomenological tradition has carefully articulated about how perspective-loci are individuated, how embodied subjects encounter one another, how the experiences of others are given to me with their own characteristic limits. None of this is denied.

What the chapter resists is a particular inference: from the plurality of perspective-loci within the field of intelligibility to the plurality of fields themselves. Perspective-loci pluralise; the field does not. Or, more carefully: there is no constructible principle by which the field within which any perspective is intelligible as a perspective could itself be counted. Any attempt to specify such a principle would require articulating a difference between fields, and the articulation of any such difference would already be intelligible content within the field within which the difference is articulated. The question "how many fields?" cannot be made determinate enough for its terms to do the work the question presupposes.

This is not a doctrine of cosmic unity. It is not the claim that all minds are one, that all conscious beings share a single inner life, that the apparent multiplicity of subjects is illusion. The chapter has

no metaphysical thesis about how many subjects there are. The chapter has a procedural thesis about what the question can mean. The procedural thesis is compatible with everything that ordinary speech, empirical psychology, and phenomenological description tell us about the plurality of perspectives.

The chapter draws principally on the phenomenological tradition. Three figures stand out for the chapter's purposes. Husserl articulates the structural features of pure consciousness with the greatest architectural elaboration, particularly through his analysis of intentionality and his account of monadological intersubjectivity. Sartre develops a position structurally close to the chapter's own — that consciousness at its primary level does not contain an I, that the Ego is a transcendent object of reflection rather than an inhabitant of consciousness, and that solipsism is dissolved not by populating a domain of egos but by removing the ego from consciousness in the first place. Zahavi provides the most articulate contemporary synthesis, drawing on Husserl, Sartre, Merleau-Ponty, and others to develop an account of pre-reflective self-awareness, minimal selfhood, and multidimensional intersubjectivity that operates at a level adjacent to the chapter's own argument.

These three are not opponents in the sense of the chapters that preceded this one. The chapter does not press against Husserl, Sartre, or Zahavi the way it pressed against Carruthers's higher-order theory or Frankish's illusionism. Each of the three is largely an ally on the matters this chapter must address. The dialectical work is therefore more delicate. The chapter must grant what these three give, mark the level at which their analyses operate, and indicate precisely what the chapter's structural claim adds to — or rather, sits adjacent to — the work they have already done.

The structure of the chapter follows. Section 7.2 says what a perspective-locus is, and why the structures that individuate perspective-loci are genuinely pluralising structures the chapter does not contest. Section 7.3 draws on Husserl, Sartre, and Zahavi for the phenomenological articulation of consciousness within a perspective-locus, marking what each grants the chapter and where each operates within the field. Section 7.4 takes up the intersubjectivity question directly: how the chapter preserves the reality of multiple perspectives, multiple bodies, multiple memories, multiple expressive lives, without conceding that any of these pluralities yields a plurality of awareness-fields. Section 7.5 states what the chapter as a whole has established and what it has not.

7.2 Perspective-loci

The chapter takes "perspective-locus" as its working term for the structures that individuate one point of view from another. A perspective-locus is not a metaphysical posit. It is a name for what is already plain in ordinary description: a body in a particular place, with particular sensory access, particular memories, particular dispositions, particular expressive resources, located within particular social and linguistic relations. The world has many such loci. Each is a real centre of orientation. The plurality of them is part of the basic furniture of any honest account of experience.

A perspective-locus is individuated by several distinguishable structures. None of these structures, taken individually or in combination, is what the chapter will resist. They are the unproblematic content of perspectival plurality.

Embodiment. Each perspective-locus is given through a particular body, located somewhere, in some posture, with some sensory configuration. The body is not incidental to perspective; it is partly what a perspective is a perspective from. Embodiment is irreducibly plural: there are many bodies, and any account of perspective must reckon with that fact.

Sensory access. What is presented depends on where the body is and how it is configured. My field of vision excludes what is behind me; the deaf person's auditory field is different from mine; the infant's

perceptual world is differently organised from the adult's. Sensory access is differential and locally specific. It pluralises perspective-loci in a fine-grained way.

Memory and biographical orientation. Each perspective-locus has its own remembered past, its own anticipations, its own characteristic ways of bringing prior experience to bear on current encounters. The novelist and the surgeon and the child each remember differently, attend differently, anticipate differently. Memory and biographical orientation are individuating without being arbitrary.

Expressive resources. Each perspective-locus has a characteristic expressive vocabulary — linguistic, gestural, affective, professional. The expressive resources of one perspective do not exhaust those of another. The translator and the dancer and the surgeon's hands have different expressive idioms.

Intersubjective correction. No perspective-locus is closed in on itself. Each is open to correction by other perspectives, to learning what others see and hear and remember, to being told one is mistaken, to discovering what one had missed. The structures of correction — through testimony, through joint perception, through expert testimony, through shared deliberation — are themselves intersubjectively constituted. They pluralise perspective-loci by giving each a particular openness to others.

These five structures, taken together, do most of the work in individuating perspective-loci. They are familiar from ordinary speech, from empirical psychology, from the phenomenological tradition this chapter will engage. The chapter does not contest any of them. The body is many; sensory access is differential; memory is biographically thick; expression is locally idiomatic; intersubjective correction is constitutive of how each perspective is open to the others. All of this is granted.

The chapter's question begins where these structures end. Granted that perspective-loci are richly pluralised by embodiment, sensory access, memory, expressive resources, and intersubjective correction — does this plurality entail or even suggest that there are correspondingly many fields of intelligibility? Are there as many fields as there are bodies? As many fields as there are sensory configurations? As many fields as there are biographical histories?

The question is not rhetorical, and it has a tempting answer in the affirmative. One might think: surely each perspective-locus must have its own field within which the world is given to it, and surely my field and yours are distinct because we are distinct subjects. The temptation is to read the plurality of perspective-loci as already a plurality of fields, with the further question of how the many fields are related — whether they all give onto the same world, whether they communicate, whether they are merely parallel or genuinely intersect — being the substantive question.

The chapter argues that this reading mislocates the field. The structures that individuate perspective-loci — body, sensory access, memory, expression, intersubjective correction — are structures *within* the field of intelligibility. They are intelligible to us as the structures they are because they are articulable within the field of intelligibility. To take the plurality of perspective-loci as already a plurality of fields is to mistake what the structures individuate. They individuate perspective-loci. They do not individuate the field within which the perspective-locus is intelligible as a perspective-locus.

This is not a deep claim. It is a procedural observation. To say that two perspective-loci are distinct is to say something determinate about how they differ — that the first body is here and the second is there, that the first remembers what the second forgets, that the first speaks in a register the second does not share. These determinate differences are stateable, comparable, correctable. They are stateable, comparable, and correctable because they sit within the field within which the comparing, stating, and correcting is intelligible as such.

The chapter's argument in the sections that follow is that the most articulate philosophical accounts of

perspective and self-awareness — Husserl on intentionality, Sartre on the pre-reflective cogito, Zahavi on the minimal self — themselves articulate perspectival plurality at this within-field level. None of them, however careful, articulates a principle by which the field itself could be pluralised. That observation is not a criticism of any of these accounts. They are not in the business of pluralising the field. The chapter's positive claim, drawn from this observation, is that the work they do is the work that needs doing, and that nothing further — no principle by which the field itself is counted — is available, because no such principle is constructible.

7.3 Phenomenological clarifications: Husserl, Sartre, Zahavi

The phenomenological tradition is where the most careful philosophical articulation of perspective and self-awareness lives. Each of the three figures the chapter draws on contributes a distinguishable layer. Husserl articulates the structural features of pure consciousness — the noetic-noematic correlation, the field of mental processes disclosed by the phenomenological reduction, the principle that consciousness is intrinsically of something. Sartre articulates the pre-reflective level at which consciousness is non-positionally self-aware without containing an I or Ego. Zahavi articulates the minimal first-personal givenness — the "mineness" — that pervades any conscious experience and that does not require higher-order representation to obtain.

The three are presented in this order because the dialectical distance from the chapter's argument decreases in this order. Husserl is the furthest, because his apparatus is the most architecturally elaborate and the chapter must mark the level-difference most carefully. Sartre is closer, because the pre-reflective cogito is a structural precedent for the kind of move the chapter makes elsewhere. Zahavi is closest, because his minimal-self thesis resists the same family of identifications the chapter resisted in Chapter 4 and converges on the chapter's own resistance to higher-order accounts.

In each case, the chapter grants what the phenomenologist achieves within the proper jurisdiction of the analysis, and marks the level at which the analysis operates. None of the three is presented as an opponent. None is presented as having held the chapter's view secretly. None is presented as having failed within their own programme. The work is to acknowledge what each has done and to indicate where the chapter's structural claim sits adjacent to that work.

7.3.1 Husserl on intentionality

Husserl's *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy*, Book I (1913) is the canonical articulation of his mature account of intentionality. The book proceeds by stages. The natural attitude is described and then methodologically bracketed by the phenomenological reduction or *epoché* (§§31–32). What remains after the bracketing — the phenomenological residuum — is pure consciousness, investigated in its own essential structures (§33 and following). The principal structural feature of pure consciousness, Husserl argues, is intentionality: the directedness of every mental process upon something.

The thesis is announced at §84 in characteristically firm form: "Under intentionality we understand the own peculiarity of mental processes 'to be consciousness of something' " (Husserl, *Ideas I* §84, Kersten p. 200). The directedness is not a feature added to mental processes from outside; it is constitutive of what they are. A perceiving is a perceiving of something, a judging is a judging of something, a valuing is a valuing of something, a remembering is a remembering of something. The variation across regions — perceptual, judgmental, emotional, volitional — is variation in the mode of directedness; intentionality itself runs through them all.

The structural articulation of this directedness is the noetic-noematic correlation. Every intentional mental process has a noetic moment (the act-aspect: the perceiving, the judging) and a noematic moment (the content-aspect: the perceived as perceived, the judged as judged). The two are correlative; they are abstractable only in analysis, not in fact. The same object can be intended under different noematic senses, and a single noematic sense can be the content of different noetic acts. The noetic-noematic structure is the principal theme of Husserl's analysis through the second half of *Ideas I* (§§87–95).

A further structural feature is the duality of hyle and morphe (§85, Kersten pp. 203–207). Within pure consciousness, Husserl distinguishes the non-intentional sensuous stuffs — the colour-data, touch-data, tone-data — from the intentional moments that animate them into "consciousness of something." Sensation, considered in itself, is the hyletic stratum; the intentional act forms the hyletic stuff into a noematic content. The duality marks a structural distinction within consciousness between what is not yet intentional and what is intentional.

The book grants the entirety of this analysis within its proper jurisdiction. Intentionality is a real and pervasive structural feature of conscious life. The noetic-noematic correlation is a careful articulation of how mental processes relate to their intended contents. The phenomenological reduction is a genuine methodological achievement; the field of mental processes it discloses is a genuine field of investigation. The book does not contest any of this.

The load-bearing premise of the chapter's engagement with Husserl can be stated directly. Husserl's *Ideas I* identifies the principal structural feature of pure consciousness as intentionality — the noetic-noematic correlation, exhibited across all mental processes. The book grants the centrality of this structure within the field of intentional life. Its pressure is that intentionality, however pervasive and structurally fundamental within conscious life, is a content-structure within the field rather than the standing within which any such structure has its standing.

What does this mean concretely? Consider Husserl's own example of the perceiving of a tree. The perceiving is an act of consciousness directed upon the tree; the tree-as-perceived is the noematic content; the variation across profiles, perspectives, and modes of giving is the variation in noematic determinations; the synthetic identity of the tree across these variations is constituted intentionally. Each component of this analysis is articulated within the field of mental processes Husserl's phenomenological reduction has disclosed. The whole architecture — *epoché*, residuum, hyle, morphe, noesis, noema, the analysis of intentionality across regions — is laid out for inspection within a field of intelligibility in which such laying out is itself intelligible.

The book's structural argument does not propose a further item beyond this field. It does not posit a noetic principle prior to the noetic-noematic correlation. It does not propose a deeper stratum below hyle. It does not introduce a further level of consciousness behind the pre-reflective. The book observes, rather, that the field of intelligibility within which Husserl's whole apparatus is articulable as an apparatus is not itself one of the structures the apparatus articulates. To say this is not to add a further item. It is to mark a difference in level between articulated structures of conscious life and the standing within which any articulation has its standing.

The analogy with Chapter 6's engagement with Kant is structurally exact. Kant articulated time as the form of inner intuition and showed that anything appearing in inner sense is temporally structured. Husserl articulates intentionality as the principal structural feature of pure consciousness and shows that any mental process is intentionally structured. In both cases, the book grants the analysis within its jurisdiction and presses on the move from the structural feature to a condition governing the field. In both cases, the pressure is not against the philosopher's analysis of the structural feature; the pressure

is against an additional move from the feature to a standing it cannot bear.

A Husserlian response to this is anticipated and worth stating. The phenomenological residuum is not a substrate behind the field of mental processes; it is the field of mental processes considered as the field of phenomenological investigation. If the book is gesturing toward something behind or beyond the residuum, the gesture is unphenomenological. Husserl himself would resist any such move; the residuum is what phenomenology investigates, and claims about something further are not phenomenologically grounded.

The reply: the book is not gesturing toward anything behind or beyond the residuum. The field of intelligibility is not a further item or domain. It is not a substrate. It is not a deeper structure. To say that the field of intelligibility within which Husserl's apparatus is articulable is not itself one of the structures the apparatus articulates is not to say that something else is. It is to say that the level at which the structures are articulated is not the level at which the standing of any articulation is in question. The two levels do not stand in a part-whole relation, nor in a substance-attribute relation, nor in a foundation-superstructure relation. They stand in the relation of an articulation to the standing within which that articulation is intelligible as one.

Husserl's own apparatus, as it happens, contains distinctions that can be used analogously to mark the level the chapter is concerned with. The distinction between hyle and morphe distinguishes what is not yet intentional from the intentional moments that animate it. The distinction between actional and non-actional cogitations (§35) distinguishes mental processes engaged actively by the pure Ego from those that subsist in the background or fringe. The distinction between immanent and transcendent objects (§38, §42) marks what is genuinely contained within a mental process from what is intended as transcendent to it. Each of these distinctions operates within the field of mental processes. The book uses none of them to identify its field. The field is not the hyletic stratum (hyle is, in Husserl's terms, sensuous stuff that intentional acts form into noematic content). The field is not the non-actional fringe (the fringe is part of the field of mental processes). The field is not the immanent side of the immanent/transcendent distinction (immanence, for Husserl, is a feature of how mental processes contain or fail to contain their objects, and is itself within the residuum).

The book's field is something the analogy with these Husserlian distinctions can model without supplying. It is the standing within which Husserl's whole apparatus — the *epoché*, the residuum, the noetic-noematic correlation, the duality of hyle and morphe, the actional/non-actional distinction, the analysis of intentionality across all regions — is intelligible as the apparatus it is. The book does not extend Husserl. The book does not claim Husserl secretly held the book's view. The book observes that the work Husserl has done is the work that needs doing at the level it operates, and that the question the book is pursuing is a question one level adjacent to that work.

7.3.2 Sartre on the pre-reflective cogito

Sartre's *The Transcendence of the Ego* (1936/37) opens with a thesis directly contrary to Husserl's later position in *Ideas I*: there is no I inhabiting pre-reflective consciousness. The Ego is not formally or materially in consciousness; "it is outside, in the world; it is a being in the world, like the Ego of another" (Sartre, *Transcendence of the Ego*, Brown translation p. 1). The argument of Part I, which is the part the chapter draws on principally, has three connected lines.

The first line is that the unity of consciousness does not require an I as a unifying principle. The unity of consciousness, on Sartre's analysis, is achieved through (a) the transcendent objects toward which consciousness is directed, and (b) the synthetic activity of consciousness in time. The "I think" can

accompany any of my representations, in Kant's formulation, but it does not follow that the I unifies the representations from a position above them. The unity is already in place at the level of consciousness's intentional directedness toward its objects. The I, if it accompanies the representations, does so because the unity is already there; it is not the source of the unity (Brown pp. 3–4).

The second line is that the individuation of consciousnesses does not require an I either. Consciousness, on Sartre's view, is individuated by its own nature. "Consciousness (like Spinoza's substance) cannot be limited except by itself" (Brown p. 4). Two consciousnesses are different because they are different consciousnesses; what makes the first a different consciousness from the second is not that the first has its own I and the second has another. The individuation is intrinsic to each consciousness as a synthetic, individual totality. The Ego, if anything, is an expression of this individuality rather than its condition.

The third line is the most direct. To posit an I within pre-reflective consciousness would, on Sartre's view, "violently separate consciousness from itself" (Brown p. 4). Consciousness is "lightness" and "translucency"; an I would introduce "a centre of opacity" into it. Consciousness's translucency is essential to its being conscious, and any positing of an I within consciousness compromises that translucency. The conclusion: pre-reflective consciousness is impersonal. The transcendental field, before the Ego is constituted, is "without an I" (Brown p. 3).

The famous illustrating passage is at Brown p. 7: "While I was reading, there was a consciousness of the book, of the heroes of the book, but the I did not inhabit this consciousness, it was merely consciousness of the object and non-positional consciousness of itself." The pre-reflective consciousness is consciousness of its object together with a non-positional, non-thetic consciousness of itself. The I does not figure. The I appears only at the reflective level, as a transcendent object given through the reflective act.

The same thesis is developed in different vocabulary in the Introduction to *Being and Nothingness*. Sartre articulates there the structure of pre-reflective consciousness as non-positional self-awareness: "every positional consciousness of an object is at the same time a non-positional consciousness of itself" (Sartre, *Being and Nothingness*, Introduction §III, Barnes translation). The pre-reflective cogito is "the condition of the Cartesian cogito." Consciousness is conscious of itself not by representing itself but by being itself in a particular way — translucently, without the duality of subject and object that reflection would introduce.

The load-bearing premise of the chapter's engagement with Sartre is procedural rather than ontological. Sartre helps dissolve the egological form of the solipsism worry by removing the I from consciousness and treating the Ego as a transcendent object. The book grants the procedural lesson: the field is not owned in the way the private-theatre picture supposes. The book uses Sartre to block ego-ownership, not to identify awareness with Sartrean consciousness.

The book's engagement with Sartre is differently positioned from the engagement with Husserl. With Husserl, the work was to mark a level-difference between the structures of intentional life and the field within which they are articulated. With Sartre, the work is partly to draw on Sartre's own structural distinction — between pre-reflective consciousness and reflection, with the I located only at the reflective level — as a precedent for the kind of distinction the book has been drawing throughout. Sartre is the closest phenomenological precedent for the procedural claim that a tempting picture of consciousness — the picture according to which consciousness is owned by an Ego that inhabits it — is a picture of something other than what consciousness most basically is.

This precedential use of Sartre is, however, not an identification. Sartre's pre-reflective consciousness is a specific phenomenological-ontological posit. It is the for-itself (*pour-soi*), the consciousness that

is what it is not and is not what it is. It has positive structural features: translucency, non-positional self-awareness, intentional directedness, the categorial contrast with being-in-itself. Sartre's analysis articulates these features in detail across the long sweep of *Being and Nothingness*. The book's field is not the for-itself. The field has none of the positive structural features Sartre attributes to pre-reflective consciousness. The field is not translucent in Sartre's sense, not non-positionally self-aware in Sartre's sense, not intentionally directed in Sartre's sense, not the categorial contrast with anything en-soi.

What the book takes from Sartre is procedural rather than ontological. The procedural lesson is that the egological picture of consciousness — the picture according to which consciousness is the private theatre of an I — is a picture of a transcendent object (the Ego) that reflection constitutes, not a picture of consciousness at its pre-reflective level. To remove the I from pre-reflective consciousness is not to deny that there are subjects; it is to relocate the I. The relocation has consequences: solipsism, as a problem, depends on the egological picture; once the picture is dislodged, the problem reformulates. The chapter draws on this reformulation in §7.4 below, where Sartre's procedural lesson is brought to bear on the intersubjectivity question directly.

A worry from within the Sartrean tradition is that the book underestimates what Sartre's pre-reflective cogito already achieves. If pre-reflective consciousness is non-positionally self-aware and translucent, then perhaps it is already what the book has been pointing toward under the name of "field." Why the further distinction? Why not identify the field with the pre-reflective consciousness Sartre has already articulated?

The reply: Sartre's pre-reflective consciousness, however carefully analysed, is a positive structural feature of conscious life. It is the for-itself, with its non-positional self-awareness and its translucency. These features are themselves articulated within a field of intelligibility in which such articulation makes sense — a field in which the contrast between positional and non-positional consciousness is stateable, in which the translucency of consciousness is intelligible as a structural feature, in which the categorial contrast between for-itself and in-itself is operative. The book's field is not the pre-reflective consciousness Sartre articulates. The book's field is the standing within which Sartre's distinction between pre-reflective and reflective consciousness has its standing as a distinction. The two levels are again adjacent rather than overlapping.

7.3.3 Zahavi on the minimal self

Dan Zahavi's *Subjectivity and Selfhood: Investigating the First-Person Perspective* (2005) is the most articulate contemporary synthesis the chapter engages. Drawing on Husserl, Sartre, Heidegger, Merleau-Ponty, and Henry, and in dialogue with contemporary analytic philosophy of mind (Shoemaker, Castañeda, Perry, Nagel, Kriegel, and others), Zahavi develops the thesis that pre-reflective self-awareness is constitutive of phenomenal consciousness — and that this self-awareness involves a *minimal self*, an experiential first-person perspective that is not the narrative self, the social self, or the personal self of identity over time.

The thesis is announced crisply at Zahavi 2005, p. 15: "Phenomenal consciousness must be interpreted precisely as entailing a minimal or thin form of self-awareness. On this account, any experience that lacks self-awareness is nonconscious." The "mineness" of experience — the quality by which an experience is given as the experience of a subject, as an experience that subject is undergoing — is not something added to consciousness by reflection or by higher-order representation. It is built into the structure of any conscious experience at the pre-reflective level.

The thesis is positioned against two main contemporary alternatives. The first is the family of higher-

order theories of consciousness — the Armstrong–Rosenthal–Lycan–Carruthers tradition — that takes the conscious status of a mental state to depend on the presence of a relevant higher-order representation. Zahavi argues that this inverts the proper order of explanation: self-awareness is not what makes a state conscious by a further representational act, but is constitutive of any conscious experience at the pre-reflective level. The second alternative is the family of accounts that links self-awareness to language, to mirror recognition, to theory of mind, or to other thick cognitive capacities. Zahavi argues that these capacities are features of more developed forms of self-awareness, not constitutive of the minimal first-personal givenness that pervades phenomenal consciousness.

The minimal self is the self at the level of any conscious experience whatever. It is not the protagonist of an autobiography. It is not the bearer of a social identity. It is not the substantial soul of dualist tradition. It is the *for-whom* of first-personal givenness — the dative of experience, the experiential subject that any conscious episode involves by its very nature as conscious.

The book's engagement with Zahavi has the most convergence of the three. The Chapter 4 argument against higher-order theories shares Zahavi's diagnosis: representations of representations do not deliver first-personal givenness, and consciousness is not made conscious by being represented. Zahavi's resistance to higher-order accounts is structurally analogous to the chapter's resistance to identification claims that locate awareness within content-architectural relations. The Zahavi engagement is, in this respect, an engagement with a contemporary phenomenologist whose work is the closest analogue to the chapter's own concerns within contemporary philosophy of mind.

The load-bearing premise of the chapter's engagement with Zahavi can be stated in two parts. Zahavi's pre-reflective self-awareness thesis holds that any conscious experience is characterised by a minimal, tacit, first-personal givenness — a "mineness" — that is constitutive of phenomenal consciousness rather than added to it by higher-order representation. The book grants this account of mineness as a structural feature within the field of conscious experience. The chapter's further claim is not a rival phenomenology of mineness but a level distinction: mineness, like the other structural features Zahavi articulates, sits within the field; it does not supply a principle for individuating fields.

The convergence on the diagnosis of higher-order theories is substantial. Zahavi's argument is that the higher-order theorist either runs into a regress (each higher-order state must itself be conscious, generating an endless ladder) or leaves the first-personal givenness of the higher-order state itself unexplained. The chapter argued in Chapter 4 that constitutive identification claims — Carruthers's particularly — locate awareness within a representational relation that the relation itself cannot supply. The two arguments share a structural form: the explanatory move that locates consciousness within a determinate content-relation runs out before it reaches consciousness. Where Zahavi's argument concludes that pre-reflective self-awareness is constitutive of consciousness rather than constituted by it, the chapter's argument concludes that awareness is not the kind of thing such a content-relation can locate.

The divergence is at the level at which these conclusions operate. Zahavi's pre-reflective self-awareness is a positive structural feature of conscious experience. It is the mineness of experience, the first-personal givenness, the for-whom that experience implicates. These are features within consciousness — careful, well-articulated features, accurately diagnosed against the higher-order theorist's misconstruals. The minimal self is the subject of these features. The chapter's argument does not propose a further structural feature of consciousness beyond the minimal self. The chapter's argument is that the level at which Zahavi articulates pre-reflective self-awareness is the level of structural features within the field of conscious experience, and that the field of intelligibility within which any such articulation is intelligible is not one of the features the articulation specifies.

This is not a charge of inadequacy against Zahavi's account. Zahavi is not in the business of articulating

the standing within which his analysis has its standing as analysis. The minimal-self thesis is offered as an account of how phenomenal consciousness is structured; the chapter takes the account at face value and grants what it claims. The chapter's observation is simply that the work the account does is within-field work, and that the question the chapter has been pursuing — about the standing of any such within-field work — is a different question.

A natural worry from within Zahavi's framework is that the chapter's distinction between within-field structures and the standing within which they have their standing reproduces, at a different level, the kind of layering Zahavi has independently shown to be unnecessary in the higher-order case. If the higher-order theorist is wrong to posit a further representational level above pre-reflective self-awareness, perhaps the chapter is wrong to posit a further level beyond it.

The reply is direct. The chapter is not positing a further representational level beyond pre-reflective self-awareness. The standing of intelligibility is not a higher-order anything. It is not a further structural feature of consciousness. It is not a meta-self, a meta-awareness, a transcendental ego in a new costume. The point is the opposite of a level-positing: it is that no further item, no further structure, no further level is being proposed. The chapter is observing that the analyses Zahavi (and the phenomenological tradition more broadly) carries out are within-field analyses, and that the question of the standing within which the analyses are intelligible as analyses is a question their analyses do not address — and need not address, because it is not the question they were designed to answer. The chapter's argument is structurally adjacent to Zahavi's, not at a higher representational level than it.

The three engagements — Husserl, Sartre, Zahavi — together articulate what the phenomenological tradition has done with the structural features of perspective and self-awareness within the field. Husserl gives the most architecturally elaborate analysis of intentionality. Sartre gives the strongest dislodgement of the egological picture by relocating the I from consciousness to reflection. Zahavi gives the most articulate contemporary synthesis, integrating the phenomenological tradition with contemporary analytic philosophy of mind and resisting the higher-order theorist's misconstruals. Each is granted. Each is engaged at the level at which the analysis operates. None of the three articulates a principle by which the field within which the analysis has its standing could itself be pluralised. The next section takes this observation to the intersubjectivity question directly.

7.4 Intersubjectivity without field-plurality

The chapter's most delicate task arrives here. A reader sympathetic to the arguments of the previous section, who has granted that perspective-loci are individuated by body, sensory access, memory, expression, and intersubjective correction, and who has followed the engagements with Husserl, Sartre, and Zahavi, may still ask: but what about intersubjectivity? What about the manifest fact that I encounter other subjects, that they encounter me, that we share a world, that we correct one another, that the second-person address is not a derivation from my first-person experience? Surely intersubjectivity is real, and surely it requires a real plurality of subjects, not merely a plurality of perspective-loci within some single field?

The chapter's answer is yes to the reality of intersubjectivity, no to the inference from real intersubjectivity to a real plurality of awareness-fields. Intersubjectivity is real and structurally rich. Its reality does not yield the conclusion that the field itself is plural. The argument of this section has three parts: a reconstruction of how the phenomenological tradition handles intersubjectivity (drawing on Husserl, Sartre, and Zahavi at this further dialectical level); a statement of why each of these treatments operates within the field rather than pluralising it; and a final paragraph that hands to the chapter's close.

Husserl's *Cartesian Meditations* (1929/1931) contains the most sustained phenomenological treatment of intersubjectivity in the canonical literature. The Fifth Meditation, "Uncovering of the Sphere of Transcendental Being as Monadological Intersubjectivity" (§§42–62), is designed to answer the charge that phenomenological reduction collapses into solipsism. Husserl's argument proceeds in stages. A "peculiar epoché within the universal transcendental sphere" (§44) brackets every sense-component that derives from the otherness of others, leaving the primordial sphere of my own monadic ego — my own animate body, my own stream of subjective process, the primordial Nature given through my pure sensuousness. Within this primordial sphere, the body of another is given as a physical body paired with my own through a passive synthesis Husserl calls *pairing* (§51). The pairing is the structural ground for what Husserl calls *appresentation* (§50): the other's stream of subjective process is appresented along with the perceptual presentation of the other's body, never directly given but co-intended in a way that has its own characteristic style of verification (§52, Cairns p. 113).

The cumulative result of these analyses is what Husserl calls *monadological intersubjectivity*. The other is given to me appresentatively as a co-constituting monad whose own primordial sphere co-constitutes the same objective world. Through the community of monads, the objective Nature is intersubjectively constituted as one — common to all monads, common to all perspectives. At the limit, Husserl states the result strongly: "Actually, therefore, there can exist only a single community of monads, the community of all co-existing monads. Hence there can exist only one Objective world, only one Objective time, only one Objective space, only one Objective Nature" (Husserl, *Cartesian Meditations* §60, Cairns p. 140). Husserl rejects as "pure absurdity" the supposition of two non-communicating intersubjectivities each constituting a separate world.

The Fifth Meditation's survey at §62 frames the achievement explicitly. The illusion of solipsism is dissolved not by leaving the transcendental attitude but by carrying it through. Within my experiencing intentionality, the alter ego is demonstrated as appresentatively mirrored, given as constitutively transcendent — given precisely as someone who is not me, but given to me within the analyses the phenomenological reduction makes available.

The chapter grants this whole architecture. Husserl's monadology is the most architecturally elaborate phenomenology of perspectival plurality on offer. The pluralisation of monads is real, the appresentation structure is carefully articulated, the constitution of the intersubjective world is sustained through detail, and the result — that there can be only one objective world common to the community of all co-existing monads — is the strongest available defence of the unity of world across the plurality of perspectives.

What the chapter observes is that the entire monadological apparatus is articulated within the transcendental attitude. Monads, the sphere of ownness, pairing, appresentation, the constitution of intersubjective Nature, the community of all co-existing monads — these are all structures laid out for inspection within Husserl's phenomenological field. Each monad is individuated through its primordial sphere of ownness; the community of monads constitutes the single objective world from that plurality of primordial spheres; the analysis presents the whole as a single articulation. Husserl's pluralisation is a pluralisation of constituting subjects within the field of constitution; it is not, and cannot be, a pluralisation of the field within which constitution-talk is itself intelligible as such talk.

Sartre's treatment of solipsism in *The Transcendence of the Ego* operates differently. Where Husserl answers the solipsism objection through monadological intersubjectivity, Sartre dislodges the picture that generates the solipsism problem in the first place. The argument is given in the Conclusion of *The Transcendence of the Ego* (Brown pp. 28–29). Sartre's claim is that the egological picture is what generates the solipsism problem in its standard form. So long as the I is a structural feature of consciousness

with a privileged ontological position, the threat that other consciousnesses are inaccessible cannot be dislodged from that privilege. But once the I is relocated — taken out of consciousness and treated as a transcendent object that, like any other transcendent object, falls under the *epoché* — the problem reformulates. "Solipsism becomes unthinkable as soon as the I no longer has any privileged position" (Brown p. 28). The other I is no less certain for consciousness than my own. Mine, Sartre concludes, "is simply more intimate" (Brown p. 29).

The Sartrean strategy is procedural. The plurality of I's is not refuted; the privileged position of any one I in particular is dislodged. The field, which was supposed to be owned by my I (with the solipsism problem arising from the impossibility of accessing other fields owned by other I's), is no longer owned by any I.

The chapter takes Sartre's procedural lesson and uses it narrowly. The lesson is that the egological form of the solipsism problem — the form according to which my private theatre is sealed off from any other private theatre — is generated by a picture of consciousness as Ego-owned. Once that picture is dislodged, the problem of how to reach other Egos from within my own Ego loses its grip. The chapter's structural argument, which has resisted identifying awareness with any determinate item Carruthers, Block, Frankish, or Chalmers proposed, can be read at this point as carrying the dislodgement one step further. If awareness is not Ego-owned, and if the field of intelligibility is not the private theatre of any subject, then the problem of how to count fields by counting Egos was a problem about counting something that was not what the field is.

This is procedural use, not ontological adoption. The chapter does not commit to the for-itself / in-itself apparatus. The chapter does not adopt Sartre's characterisation of consciousness as translucency, as nothingness, as the negative spontaneity that distinguishes itself from being. These are positive characterisations of consciousness within Sartre's phenomenological ontology; they are content-claims of one specific kind. The chapter takes only Sartre's procedural observation: that the solipsism problem assumes the egological picture, and that dislodging the picture dissolves the problem.

Zahavi's *Subjectivity and Selfhood* Chapter 6, "Self and Other," is the most articulate contemporary phenomenological account of intersubjectivity. The chapter reconstructs four complementary approaches found across the phenomenological tradition: (i) empathy as a specific intentional act distinct from perception, imagination, and recollection; (ii) embodied subjectivity with internal alterity, in which the body's touching-touched reversibility prepares the structure of intersubjectivity; (iii) world-mediated co-subjectivity, in which the perceptual object's horizon of profiles given for other possible subjects pre-figures intersubjectivity in the structure of perception itself; (iv) the irreducible transcendence of the other, in which the other is given precisely as the limit of my givenness, with the inaccessibility of the other's first-personal life being constitutive of the otherness of the other rather than a defect in my access to it. Zahavi's argument is that no single approach is adequate, and that a satisfactory account combines all four.

Several Zahavian claims bear directly on the chapter. The first is that the other's first-personal givenness is constitutively inaccessible to me: "Had I had the same access to the consciousness of the other as I have to my own, the other would cease being an other and instead become a part of myself" (cited from Husserl by Zahavi, *Subjectivity and Selfhood* p. 154). The asymmetry between first-person and second-person access is constitutional, not an imperfection. The second is that intersubjectivity has a place in the very structure of perceptual intentionality. On Zahavi's reconstruction, Husserl held that "every appearance that I have is from the very beginning a part of an open endless, but not explicitly realized totality of possible appearances of the same, and the subjectivity belonging to this appearance is open intersubjectivity" (cited from Husserl by Zahavi, *Subjectivity and Selfhood* p. 167). Perception is

intersubjective in its horizon-structure. The third is that mineness — the first-personal givenness constitutive of phenomenal consciousness — is compatible with intersubjectivity rather than opposed to it: mineness and otherness are mutually constitutive structural features of consciousness. Zahavi cites Husserl on the further claim, drawn from his unpublished manuscripts, that "the transcendental subject is only what it is within intersubjectivity" (cited from Husserl by Zahavi, *Subjectivity and Selfhood* p. 173).

The chapter grants Zahavi's reconstruction as the most articulate contemporary treatment of perspectival plurality within the field. The four-approach account is a richer integration than any of its components taken alone. The four approaches handle the manifest reality of intersubjectivity at four distinguishable levels — at the level of the face-to-face encounter, at the level of bodily interlacement, at the level of world-mediated co-subjectivity, at the level of the irreducible alterity of the other — without collapsing into either solipsism (the other is constitutively given, with constitutive inaccessibility) or anonymous monism (the difference between self and other is not derivative of a more fundamental anonymity).

What the chapter observes, again, is that the entire Zahavian articulation operates within the field. Mineness, embodied interlacement, world-mediated co-subjectivity, the transcendence of the other — these are structural features within the multidimensional structure of consciousness Zahavi articulates. They handle the question of how perspective-loci are individuated and how they are open to one another. They do not handle, and were not designed to handle, the question of how the field within which any of these structures is intelligible has its standing as field.

The level-difference, stated with the care this section requires, is this. Each of the three accounts — Husserl's monadology, Sartre's procedural dislodgement of the egological picture, Zahavi's multidimensional account of intersubjectivity — articulates the plurality of perspective-loci within a field. The pluralisation is real and is handled at the level the accounts operate. Husserl's monads are individuated through ownness; Sartre's consciousnesses are individuated through their objects; Zahavi's perspective-loci are individuated through mineness, embodiment, and the constitutive transcendence of the other. None of these articulations, however careful, supplies a constructible principle by which the field within which the articulation has its standing could itself be counted. To articulate such a principle would require specifying a difference between fields, and any specification of a difference would already be intelligible content within the field within which the specification is made. The pluralisation of fields would be carried out within a field. That is not pluralisation of the field.

This is not the assertion of monism. The chapter does not claim that there is exactly one field, or that all minds are one, or that the apparent plurality of subjects is illusion. The chapter makes no count. The chapter observes that the question "how many fields?" cannot be made determinate enough to receive an answer. The structures by which perspective-loci are individuated — body, sensory access, memory, expression, intersubjective correction — do their work at the level of perspective-loci. They do not transfer to the level of fields, because there is no specification of a difference between fields that would not already be a within-field articulation.

The result is not a reduction of others to my perspective, nor a doctrine of one cosmic subject. It is the narrower claim that the structures by which perspectives are individuated — body, access, memory, expression, and intersubjective correction — individuate perspective-loci within the field. They do not provide a principle by which the field itself could be counted.

7.5 The narrow result of perspectival plurality

The chapter's claim is narrow. Perspective-loci are real and richly individuated; the inference from their plurality to the plurality of fields is not constructible. That is the whole of what has been established.

The chapter does not establish a doctrine of one cosmic mind. No claim has been made that all conscious beings share a single inner life, or that there is some single subjectivity within which the apparent multiplicity of subjects is contained. The chapter has made no count of subjects, has posited no metaphysical unity beyond perspective-loci, and has not proposed any positive thesis about how many anything-relevant there is. What the chapter has done is observe that the question "how many fields?" cannot be made determinate enough for an answer to be the answer it would have to be.

Nothing in the chapter denies the manifest reality of multiple perspectives. The arguments of the previous sections have, on the contrary, taken as their starting point that perspective-loci are real and richly individuated. Embodiment, sensory access, memory, expressive idiom, and intersubjective correction are real pluralising structures. The chapter takes them at face value and does no work to reduce any of them. To grant the plurality of perspective-loci is not yet to grant the plurality of fields; the chapter resists only the inference from the first to the second.

The chapter has not held that Husserl, Sartre, or Zahavi failed within their own programmes. Husserl's monadology is the most articulate phenomenology of intersubjectivity in the canonical literature; Sartre's procedural dislodgement of the egological picture is one of the most influential moves in twentieth-century French philosophy; Zahavi's contemporary synthesis is the closest analytical-phenomenological counterpart to the chapter's own concerns. Each is granted within the scope it operates. The chapter has used each as an interlocutor whose work clarifies what perspectival plurality looks like at the within-field level, and has marked where the chapter's structural question sits adjacent to that work.

The chapter has not proposed a higher-order meta-self, a transcendental ego in a new costume, or any further item beyond the structures Husserl, Sartre, and Zahavi articulate. The level-difference the chapter has marked is not a difference between two items, one above the other. It is the difference between an articulation of structural features within the field and the standing within which any such articulation is intelligible as articulation. To mark this difference is not to posit a further item; it is to observe what the items already posited do and do not specify.

Nothing in the chapter has displaced the phenomenological analysis of empathy, embodied interlacement, world-mediated co-subjectivity, or the irreducible alterity of the other. These analyses remain the substantive content of any account of perspectival plurality. The chapter has drawn on them, granted them, and indicated where its own structural observation sits in relation to them. The four-approach account that emerges from Zahavi's reconstruction of the phenomenological tradition is the most articulate available treatment of how multiple perspectives are real and how they relate; the chapter has not contested any part of it.

What the chapter has established is narrow and procedural. There is no constructible principle by which the field within which perspectival plurality is intelligible could itself be counted. Any attempted individuation of fields would have to specify a difference, and that specification would already be intelligible content within the field. The plurality of perspective-loci, embodiment, sensory access, memory, expression, and intersubjective correction is real and is the substantive content of perspectival plurality. The further inference — from that plurality to a plurality of fields — runs into a question whose terms cannot be made determinate enough to bear the answer the inference would require.

The next chapter takes up the question of ontological jurisdiction. If the field within which inventories of items, structures, and relations are articulated is not one more item to be inventoried, then the relation of awareness to ontology needs to be stated. Chapter 8 addresses this directly, and engages with contemporary panpsychism and cosmopsychism on the specific question of how widely consciousness can be distributed across an inventory of items.

Chapter 8 — Ontological Jurisdiction

8.1 Why ontology returns here

The book has been mostly negative. Chapter 3 argued that no ontological derivation of awareness can be stated without presupposing the field within which the derivation's terms become intelligible. Chapters 4 and 5 argued that the principal contemporary identifications and reframings of consciousness — higher-order representation, global workspace, access functions, illusionism, the meta-problem — each locate awareness within a determinate content-architectural feature of conscious life and so relocate the field they purport to explain. Chapter 6 argued that Kant's account of time as the form of inner sense, however successful within its proper jurisdiction, does not extend to a condition governing the field. Chapter 7 argued that the structures by which perspectives are individuated — body, sensory access, memory, expression, intersubjective correction — individuate perspective-loci within the field rather than supplying any principle by which the field itself could be counted.

A reader who has followed this sequence may by now wonder what positive ontological commitment, if any, the book has been holding back. A book that denies that awareness is ontologically derivative, identifiable, temporally conditioned, or pluralisable might appear to be sliding toward a positive thesis about awareness as a metaphysically privileged item — the one item that survives all the eliminative pressure. Such a thesis would have to be stated as an ontological claim of some kind: that awareness is a fundamental property, a fundamental individual, an irreducible substance, a cosmic mind, an absolute. The chapter's task is to state what the eliminative arc actually licenses and what it does not. The short answer is that the arc licenses a *jurisdictional* limit on ontology rather than a positive ontological posit within it. The longer answer is the chapter's work.

Ontology, as the discipline that catalogues what exists and how the items in the catalogue depend on one another, is not abandoned. The chapter is not an attack on ontology as such. Ontological inquiry into items, properties, structures, dependence relations, grounding relations, and fundamental individuals proceeds on its own merits, and the book has no quarrel with such inquiry within its scope. The book's argument is about the *jurisdiction* of ontology — about what it can and cannot include within its catalogue. The standing within which any catalogue is intelligible as a catalogue is not, the chapter argues, one of the items the catalogue can list.

The chapter uses Philip Goff's *Consciousness and Fundamental Reality* (2017) as the central worked case. The choice is not arbitrary. Goff is the contemporary philosopher whose work converges most closely with the negative arguments of Chapters 3–5. He shares the book's resistance to deriving consciousness from non-conscious matter; he shares the book's resistance to identifying consciousness with functional or representational relations within content-experience; he shares the book's resistance to illusionism. He has, moreover, articulated as carefully as anyone the constructive ontological alternative the book must refuse. Constitutive cosmopsychism — the combination of priority monism with panpsychism

— is offered as the residual position after a chain of careful internal filters, and Goff develops it in full architectural detail. The engagement with Goff therefore allows the chapter's positive thesis to be tested against the most demanding contemporary form of the position the book diverges from.

The chapter is structured as follows. §8.2 says what the book does not claim, ruling out the ontological positions readers might suspect the eliminative arc commits the book to. §8.3 reconstructs Goff's constructive cosmopsychism, fairly and in his own terms. §8.4 marks the point of convergence between Goff's view and the arguments of Chapters 3–5. §8.5 marks the point of divergence and engages the most likely objection — that the book's distinction between the field and the ontological inventory is either a further posit or an evasion. §8.6 states the positive thesis of ontological jurisdiction. §8.7 states what the chapter establishes and what it does not.

8.2 What the book does not claim

The negations have to come first. Several positions a reader might reasonably suspect the book of holding are not the book's position. Each can be stated and dispatched briefly.

Not idealism. Idealism, in any of its standard forms, holds that the world is constituted by mind — that the items, structures, and dependencies the world contains are in some sense constructions of, or contents within, a mind or community of minds. The book makes no such claim. It does not deny the observer-independent existence of stars, mountains, biological organisms, or the events in which they participate. The book's argument has been about what *follows from* the standing of awareness within which any ontological discourse is intelligible, not about whether the items the ontological discourse catalogues exist independently of any mind. Idealism is one specific metaphysical thesis within ontology; the book's argument is jurisdictional and does not entail it.

Not dualism. The book does not claim that mind and matter are two equally fundamental substances. It does not claim that awareness is a non-physical thing standing alongside the physical, with the two requiring some mechanism of interaction. Dualism is, again, a specific metaphysical position within ontology — a position about the count and kind of fundamental substances. The book makes no such count and no such kind-attribution. The structural argument of the preceding chapters has refused, throughout, to identify awareness with any positive ontological item that could enter a count.

Not monism in the positive sense. The book does not claim that all reality is one substance, that all distinctions are merely apparent, that the manifest plurality of items and persons is an illusion overlying a substantial unity. This is the most pressing potential confusion, because the arguments against pluralising awareness in Chapter 7 may sound like an assertion of unity at the level of awareness. They are not. The chapter explicitly noted (§7.5) that the chapter makes no count. A refusal to pluralise is not an assertion of one; both pluralisation and counting-to-one operate within the inventory the book's structural argument is about, and the book takes neither side within the inventory.

Not panpsychism. The book does not claim that consciousness is ubiquitous in matter, that fundamental physical entities have phenomenal properties, or that the deep nature of the physical is experiential. Panpsychism is a specific positive thesis about the distribution of conscious properties across an ontology of fundamental items. The book takes no position on that distribution. Whether or not it should turn out that electrons have some primitive form of conscious experience is, on the book's view, a within-ontology question on which the book's structural argument is silent.

Not cosmopsychism. The book does not claim that the cosmos is a conscious subject, that the universe has a unified mental life of which our individual minds are aspects, that there is one cosmic awareness within which all individual awarenesses are subsumed. This negation is particularly important given

the engagement with Goff that follows. The book's argument is not a cosmopsychist argument with the cosmic subject given a different name. The structural argument resists the positing of any item — including a cosmic subject — as the locus of awareness, because awareness, on the book's argument, is not the kind of thing that has a locus within an ontological inventory.

Not mysticism. The book does not claim that awareness is ineffable, that the field is beyond reasoned discourse, that what the book has been pointing toward is something the methods of analytic philosophy cannot reach. The argument has been throughout analytic, structural, and inferential. The book's restraint about positive ontological characterisation of awareness is not a mystical reticence; it is a structural observation that ontological characterisation is not the level at which the question of awareness as field is to be answered.

Not realism as a positive metaphysical doctrine. The book is compatible with realism about ordinary objects, scientific entities, and much else. But it does not itself advance realism as its positive thesis. Realism is a position within ontology about the status of items in the inventory. The book's claim concerns the standing within which any such inventory is intelligible.

Not anti-realism. The book does not deny that there is an external world, that ordinary perceptual claims about the world can be true, that empirical science delivers genuine knowledge about reality, that the items the sciences catalogue exist. The book's argument operates at a level adjacent to all such claims. Empirical science succeeds within its proper jurisdiction; ontological inquiry succeeds within its proper jurisdiction; the book's argument concerns a structural feature of the field within which all such inquiry is intelligible as inquiry.

These eight negations clear the ground. None of the standard ontological positions a reader might reach for is the book's position. What the book is doing, the chapter argues, is articulating a jurisdictional limit on ontology itself — a limit that is not contained in any of the positions just ruled out, because those positions are themselves operative within ontology rather than about ontology's scope.

8.3 Goff's constructive alternative

Philip Goff's *Consciousness and Fundamental Reality* is the most articulate contemporary defence of an ontology that takes consciousness as fundamental. The book begins from what Goff calls *the Consciousness Constraint*: "Any adequate theory of reality must entail that at least some phenomenal concepts are satisfied" (Goff 2017, §1.1.1, p. 3). The constraint is offered as the Archimedean axiom of Goff's metaphysical method. Any ontology, on Goff's view, must accommodate the unrevised reality of phenomenal consciousness as a datum on par with the data of third-person observation and experiment.

Goff's diagnosis of the contemporary situation runs through what he calls *the Galileo argument* (Goff 2017, §1.1.3, pp. 11–14). The success of mathematical physics over the past five centuries, on Goff's reading, is the success of a deliberately restricted inquiry. Galileo, in the *Assayer*, located the sensory qualities — colour, taste, smell, sound — outside the domain of physics, placing them in the soul of the perceiver. What remained — extension, motion, shape, number — could be mathematised. The physics that grew out of this restriction has been extraordinarily successful, but its success is the success of a domain that has already excluded the sensory qualities from its remit. Those qualities still exist, Goff argues, and must be re-incorporated into the ontological picture. Locating them in conscious experience is one move; integrating conscious experience back into the ontology of nature is the further task.

Goff's positive proposal is *Russellian monism*. The view has two components (Goff 2017, §1.2.2, pp. 17–18). The negative component: physical science describes only the relational, dispositional, extrinsic, or

structural properties of matter. There is a deeper categorical, intrinsic, or non-structural nature underlying that structure — the "deep nature" of matter — about which physics is silent. The positive component: the deep nature of matter transparently grounds consciousness. Within the broader Russellian framework Goff is a *panpsychist* (the deep nature of matter is itself experiential, not merely suited to bringing about experience), a *constitutivist* (ordinary consciousness is grounded in, not merely caused by, the deep nature of matter), and a *priority monist* (the cosmos is the one fundamental individual, with everything else grounded in it).

The combination yields *constitutive cosmopsychism*. The chapter takes care to mark that Goff did not reach this position by direct preference. He was driven to it by an internal filter the book should reconstruct fairly. The standard form of constitutive Russellian monism is smallist micropsychism: basic micro-level entities have some primitive form of conscious experience, and human and animal consciousness is grounded in the experiences of the micro-subjects composing the brain. Goff's *Subject Irreducibility* argument (Goff 2017, §8.3, pp. 209–218) rules this out.

Subject Irreducibility is the thesis that "what it is for there to be a conscious subject S cannot be analyzed into facts not involving S" (Goff 2017, §8.3, p. 209). The intuition is articulated by contrast with parties: a party admits of a deflationary analysis in terms of people revelling, where the people are not the party. Subjects do not, on Goff's view, admit of any analogous deflationary analysis. Goff defends Subject Irreducibility through a process of elimination, considering the candidate deflationary analyses available to a constitutive Russellian monist and showing that each either fails outright or collapses into a non-deflationary form. Analytic functionalism is ruled out by the arguments against type-A physicalism (Goff 2017, Ch. 3). A posteriori physicalism is ruled out by Phenomenal Transparency (Goff 2017, Ch. 5). The micropsychist analysis in terms of micro-subjects standing in a relation R generates an infinite regress and lacks a priori availability for R. The co-consciousness analysis, in Dainton's form, either presupposes subjects (incoherent) or reduces to "experienced by a single subject" (not deflationary).

The argument's force is not litigated here. The chapter grants Goff's Subject Irreducibility as the load-bearing conclusion he draws and traces its consequences. Given Subject Irreducibility, the standard mechanism of *grounding by analysis* fails for o-subjects (Goff 2017, §8.3, pp. 216–218). Grounding by analysis requires that the grounded entity's essential conditions be specifiable in the grounding terms. If subjecthood admits of no deflationary analysis, the micro-level facts cannot, by themselves, logically entail the existence of any particular o-subject. Constitutive micropsychism with grounding by analysis is therefore unavailable to the constitutive Russellian monist.

The move to cosmopsychism is, on Goff's reading, the residual option. Emergentism (consciousness causally arises from but is not grounded in the deep nature of matter) is resisted on empirical grounds (Goff 2017, §9.4). Within the constitutive family, what is needed is a form of grounding that does not require the analysability of the grounded entity. *Grounding by subsumption* is the alternative (Goff 2017, §9.1, pp. 220–227): "Entity X grounds by subsumption entity Y iff (i) X grounds Y, and (ii) X is a unity of which Y is an aspect" (Goff 2017, §9.1, p. 221). Subsumption-grounding does not require Y to be analysable in terms not involving Y. The grounded entity is an *aspect* of a unity, and an aspect can be irreducible while still being grounded in the unity of which it is an aspect.

The cosmopsychist view is the application of grounding-by-subsumption to the question of o-subjects. The cosmos is one fundamental unified individual; its deep nature is *consciousness+*, a property with experiential and non-experiential aspects (the consciousness+ apparatus is introduced in §9.2 to block a conceivability objection that pure cosmic consciousness would be conceivable in the absence of any o-subject's consciousness). Ordinary subjects are *vertical aspects* of the cosmos: proper parts that are aspects of the whole. Ordinary conscious states are *horizontal aspects* of vertical aspects: aspects that

abstract away from the richness of the vertical aspect (Goff 2017, §9.3.2, pp. 238–241). The architecture is buttressed by an empirical argument running through Power Realism and Nomic Generality (Goff 2017, §9.4, pp. 243–253) for the conclusion that, given materialism, the only plausible candidates for fundamental entities whose causal powers ground the general laws are micro-level entities or the cosmos. Subject Irreducibility eliminates the micro-level option for the constitutive Russellian monist. The cosmos is what remains.

The shape of Goff's view, fairly stated, is this. The cosmos is the one fundamental individual. Its intrinsic nature is consciousness+. Every ordinary subject is a vertical aspect of the cosmos; every ordinary conscious state is a horizontal aspect of a vertical aspect. The grounding throughout is grounding by subsumption rather than grounding by analysis. The whole architecture is a residual ontological commitment driven by the failure of smallism, the resistance to emergentism, and the demand for some form of constitutive grounding that survives Subject Irreducibility.

8.4 The point of convergence

What the book takes from Goff is substantial. The convergence runs across three of the prior chapters.

With Chapter 3, Goff agrees on the central point: consciousness cannot be derived from non-conscious matter. The arguments differ in detail and vocabulary — Goff's argument against physicalism turns on Phenomenal Transparency and the inadequacy of any a priori or a posteriori functional analysis of conscious states; the book's argument turns on the structural pressure on the load-bearing terms in any ontological derivation claim. But the two arrive at the same place. Physicalism, on both arguments, fails to give consciousness the role it must have; the routes differ. Goff's route runs through Phenomenal Transparency and the failure of physicalist analyses, while the book's route runs through the structural pressure on ontological derivation claims. Goff is, in this respect, the closest contemporary ally on the negative phase of the book's argument.

With Chapter 4, Goff agrees on the inadequacy of standard physicalist identifications of consciousness with content-architectural relations. Higher-order representation, functional role, access architecture, global workspace — each of these, on Goff's argument, fails to deliver Phenomenal Transparency and therefore fails to identify consciousness with anything that could play the role consciousness plays in our cognitive lives. The book's Chapter 4 argument runs along structurally similar lines: identifications collapse the field into one of its contents, because whatever is determinate enough to play the identifying role is already content within the field rather than the field itself. Goff and the book diverge on what to make of this convergence — Goff takes it as motivation for Russellian monism and the deep-nature thesis; the book takes it as a structural observation about the limits of any identification claim — but they converge on the negative point.

With Chapter 5, Goff agrees on the inadequacy of illusionism. The Consciousness Constraint, taken as Goff's axiom, rules out any account that revises or eliminates phenomenal consciousness. Goff is explicit (Goff 2017, §1.1.2, pp. 9–10) that illusionism is the position the Consciousness Constraint is designed to exclude. The book's Chapter 5 argument operates at a structurally adjacent level — it grants illusionism's careful analysis of the meta-problem and presses on what such an account leaves untouched — but the result is convergent: illusionism does not deliver what an account of consciousness needs to deliver.

The cumulative convergence is significant. Goff is not a target of the negative phase of the book's argument. Where the book denies, Goff also denies, and his arguments are often parallel in structure to the book's. The chapter takes this seriously. Goff is engaged here as the most articulate contemporary

expression of the negative phase the book also affirms — and as the philosopher who has gone furthest in trying to specify what positive ontology the negative phase supports.

The divergence is at exactly that step. Goff moves from the negative phase to a positive ontological architecture in which consciousness is the deep nature of the one fundamental individual. The book makes no such move. The divergence is not a disagreement about the negative phase; it is a disagreement about what the negative phase makes available constructively.

8.5 The point of divergence

The divergence has two parts. The first is the substantive structural claim: Goff's constructive ontology places consciousness as an *item* within an ontological inventory, even after Goff's own filters have done their work. The second is the response to a likely objection from within Goff's framework: that the book's distinction between the field and the items the inventory catalogues is either a further metaphysical posit or an evasion.

Take the substantive claim first. Goff's constitutive cosmopsychism, after all the within-framework filtering has been done, is an *ontological architecture*. The architecture has fundamental and non-fundamental items (the cosmos is fundamental; ordinary subjects and material objects are not). It has properties of those items (consciousness+ is the intrinsic property of the cosmos; the relational/dispositional properties physics describes are the extrinsic side). It has relations among the items (grounding by subsumption rather than grounding by analysis; the vertical/horizontal aspect distinction; the relations of dependence captured by Power Realism). It has laws (the general laws governing animate and inanimate entities, grounded in the causal powers of fundamental entities). Each of these — items, properties, relations, laws — is a category within the ontological catalogue. The architecture is a thoroughgoing ontology, articulated with great care, and consciousness has been incorporated into the catalogue at the level of the intrinsic nature of the cosmos.

The book's structural argument operates on the entire architecture, not on any one component. The argument is not that Goff has mis-identified consciousness with the wrong item, nor that consciousness should have been placed at a different level of the catalogue, nor that the cosmos should not have been chosen as the fundamental individual. The argument is that the field within which an architecture of this kind is intelligible *as an architecture* is not one of the items the architecture catalogues. The contrast between fundamental and non-fundamental, between intrinsic and relational, between categorical and structural, between item and inventory — every such contrast is intelligible to us as a contrast within a field of intelligibility. The field is what makes each side of each contrast available as a side. It is not itself either side, and it is not a further item alongside the contrasted items.

Goff has not failed within his programme. The chapter does not claim that Subject Irreducibility is mistaken, that grounding by subsumption is incoherent, that consciousness+ is unmotivated, that the cosmos cannot be a fundamental individual, or that the Power Realism argument is unsound. Each of these is a within-framework question on which the book takes no position. What the chapter claims is that the framework, however internally successful, is an articulation within a field of intelligibility, and that the field-question is a question about the standing of any such articulation as articulation.

This brings the chapter to the most likely objection. From within Goff's framework, the book's structural distinction between the field and the items its ontology catalogues may appear to present a dilemma:

> Either the field is a further fundamental posit, or it is an evasion. If it is a posit, then it belongs in the inventory and must compete with consciousness+ as a candidate for the intrinsic nature of fundamental reality. If it is not a posit — if it is something other than a

> fundamental item with a determinate ontological role — then it seems to do no metaphysical work, and the book has merely refused, throughout, to take a metaphysical position.

The objection is serious. It is the natural reading of the book's restraint about positive ontological characterisation of the field, from within a framework that recognises only one place for any metaphysical claim to land: in the catalogue of items, properties, and relations the ontology recognises.

The chapter's reply is that the dilemma misclassifies the book's claim. The field is not a further fundamental item to be added to the inventory, and the book is not refusing to take a metaphysical position. Both horns of the dilemma assume that any metaphysical claim must be statable as a claim about an item within an ontological catalogue — either a positive claim that an item exists with such-and-such properties, or no claim at all. The book's argument does not fit either horn because the book's claim is not a claim about an item.

The field, as the book has been developing it, is the standing within which the contrast between item and inventory, between fundamental and non-fundamental, between intrinsic and relational, between categorical and structural, between unity and aspect, is intelligible at all. Each of these contrasts is operative within an ontology; each of them gives an ontology its shape. None of them is itself a posit on either side of the contrast it draws. The field is what makes these contrasts available as contrasts, and it is not a further posit alongside them.

The book is therefore not adding an item to Goff's ontology. It is not competing with consciousness+ for the role of the intrinsic nature of the cosmos. It is not offering a rival fundamental individual. It is denying that the field-question is the kind of question that is answered by adding, relocating, or deepening any item within ontology. The denial is itself a substantive claim — a structural observation about the limits of ontological articulation — but it is not a claim about a further item.

Is this not, the Goffian might press, simply a refusal to play the metaphysical game? Is the book not saying, in effect, that the field-question is not amenable to ontological treatment, and is that not a kind of evasion? The reply is that the field-question *is* amenable to disciplined philosophical treatment — the book's argument throughout has been such treatment — but that the discipline required is not the discipline of placing awareness within an ontological inventory. The treatment is structural and jurisdictional. It traces what ontology can and cannot include within its catalogue. To say that awareness as field is not an item ontology catalogues is not to refuse to say what awareness is; it is to say something specific about the relation between awareness and the form of articulation ontology offers.

What kind of philosophical work the chapter is doing in saying this is worth marking, because the work has a definite shape that distinguishes it from refusal. Throughout the book's eliminative arc, the book has produced structural diagnoses: the determinate items each rival's claim deploys are themselves articulated within the field whose subordination the claim asserts; the priority asymmetry, the identificatory placement, the conditioning relation cannot be marked without already drawing on what the claim proposes to subordinate. The diagnoses are not refusals to engage. They are sustained engagements with the rival's structure, carried through to the point at which the structure no longer holds. Chapter 8's positive thesis is the cumulative shape of that engagement: ontology can do what it does, but what it does has a definite scope, and the scope reaches its edge at the field-question. Tracing that scope is not abstention from ontology; it is the philosophical work the eliminative arc has been doing all along.

The work is jurisdictional in a precise sense. It does not stand outside ontology's disputes and decline to participate; it traces where those disputes can and cannot bear on a specific question. That tracing is a positive contribution to how the question of awareness is to be philosophically engaged. The con-

tribution is not a new item in any catalogue, not a new property of any item, not a new grounding relation among items; it is a structural observation about what catalogues, properties, and grounding relations can and cannot be made to settle. The observation is substantive because it changes what within-ontology results can be taken to have established about awareness. The chapter is not standing outside the disputes of constitutive cosmopsychists, smallist micropsychists, dualists, and physicalists; it is marking what those disputes can and cannot decide. To say that the field is not what any of them is disputing over is to make a definite philosophical claim about the relation of the disputes to the question, and a definite claim is not an evasion.

A related Goffian worry remains, and it is the subtler one. Even granting that the chapter is doing substantive philosophical work, the worry runs: any sufficiently definite characterisation of the field is going to *be* a positive ontology of the field, whatever the chapter calls it. To say what the field is — even to say that it is the standing within which cataloguing is intelligible as cataloguing — is to attribute a determinate character to a determinate item, and a determinate character attributed to a determinate item is just what an ontology of that item consists in. The chapter, on this reading, has been doing positive ontology in a vocabulary that disguises it.

The reply is that the chapter's substantive claims about the field lack the structural marks a positive ontology has. A positive ontological claim places an item within a catalogue: it names the item, ranks it among fundamental and non-fundamental items, attributes properties to it, specifies the grounding or constitution relations in which it stands to other items, and lets it figure as a subject of further ontological dependence claims. Goff's consciousness+ has each of these marks: it is named; it is placed at the level of fundamental reality; it has the property of being the intrinsic nature of the cosmos; it stands in subsumption-grounding relations to ordinary subjects; it features in within-ontology disputes with rival fundamental-nature candidates. The chapter's claims about the field have none of these marks. The field is not placed within a catalogue of fundamental and non-fundamental items. It is not attributed properties. It does not enter into grounding relations with anything else. It does not figure as a subject of ontological dependence claims about which alternative positions could be taken. A claim about the field can be substantive — can say something definite, contestable, recognisable as a structural observation — without having the form an ontological claim has. The chapter's positive thesis is in the first form; it is not in the second.

The divergence with Goff is therefore not at the level of who has the better ontology. The divergence is about whether the field-question is to be answered by ontology in the first place. Goff thinks it is, and he has done careful constructive work to give the best ontological answer he can. The book thinks the question is jurisdictional, and the careful constructive work it requires is structural rather than ontological.

8.6 Ontological jurisdiction

Ontology is the discipline that catalogues what exists. The catalogue can contain items, properties, structures, dependence relations, grounding relations, laws, and fundamental individuals. The relations among the categories can be many: items can instantiate properties; properties can ground other properties; fundamental individuals can have non-fundamental items as aspects; laws can govern the behaviour of items; structures can be realised by items at various levels of dependence. The internal richness of ontology as a discipline is considerable.

The discipline operates throughout under conditions that make its operations intelligible. To catalogue items, one must be able to tell what an item is; to track properties, one must be able to tell what a property is; to specify a grounding relation, one must be able to tell what grounding is. None of these conditions

is itself one of the items, properties, or relations the ontology catalogues. They are the conditions of intelligibility under which any such cataloguing has a sense as cataloguing.

The point can be put more sharply. To know what an item is, in the sense the discipline needs in order to catalogue items, is to be able to pick the item out, distinguish it from other items, describe its features, and place it in the relations it stands in to other items the discipline tracks. That kind of knowing has a sense within the cataloguing practice; it is what the practice operates on. The same goes for properties, for structures, for grounding relations. Each is operated on by the discipline in a way that requires its being available, at the level of cataloguing's own articulations, as the kind of thing the discipline can pick out. That availability is what the discipline presupposes for its cataloguing to be cataloguing, and it is what the book has been calling the field. The field is not a further item the discipline could go on to pick out; were the discipline to attempt to pick it out as an item, the picking-out would already presuppose it.

The chapter's positive thesis is that *ontological jurisdiction* is the discipline's scope within these conditions, and that awareness as field is not within ontology's jurisdiction as an item — not as something further or higher, but as the standing within which the discipline operates.

The thesis has three components.

First, *ontology catalogues items*. The catalogue's reach is wide. It can include physical particles, biological organisms, abstract universals, mathematical structures, propositions, possible worlds, events, processes, properties, relations, dispositions, fundamental individuals, and the dependence and grounding relations among them. The book's argument takes no position on which items the catalogue should include or how the inventory should be arranged. These are first-order ontological questions on which philosophers of ontology have developed careful and substantial views, and the book's argument operates at a level adjacent to those views.

Second, *ontology cannot catalogue the field as one of its items*. The standing within which an ontological catalogue is intelligible as a catalogue is not, the book argues, the kind of thing the catalogue can include. To attempt to include it would be to treat the conditions of intelligibility for the catalogue as one of the things the catalogue lists. Such an attempt produces a category confusion. The conditions under which any item is intelligible as an item are not themselves further items. The conditions under which any property is intelligible as a property are not themselves further properties. The conditions under which any grounding relation is intelligible as grounding are not themselves further grounding relations. The conditions are jurisdictional rather than catalogue-internal.

Third, *awareness as field is the standing within which ontological cataloguing is intelligible as cataloguing*. This is the positive thesis the eliminative arc of the prior chapters has been working toward. Each prior chapter argued that a candidate ontological location of awareness fails — that awareness cannot be derived from items the catalogue contains, identified with relations among them, conditioned by structural features within them, or pluralised in the way the catalogue counts items. The cumulative result is that awareness is not one of the items the catalogue can list. The positive thesis adds that awareness is, in fact, what makes any such cataloguing intelligible.

This is the strongest positive thing the book says about awareness, and it is worth marking what kind of positive characterisation the formulation is. "Standing within which" picks out awareness in its structural relation to cataloguing — not as something the cataloguing operates on, not as something it produces, not as something alongside the things it distinguishes from one another, but as the standing in which cataloguing is the activity it is. The characterisation is positive because it states something definite about awareness: the chapter does not say only that awareness is not in the inventory; it states

awareness's positive role in jurisdictional terms — the standing in which the inventory has its sense as an inventory, within which the activities that produce and revise the inventory are the activities they are.

The characterisation is not ontological because it does not place awareness as an item among other items in any catalogue. It does not say that awareness is the most fundamental thing, or the deepest layer, or the ground from which other things follow. These are first-order ontological claims with specific commitments — about priority, about the grounding relations in which items stand, about what an inventory of fundamental and derivative reality includes at its base. The chapter's claim is not any of these. To say that awareness is the standing within which cataloguing is intelligible is to say something about the relation between awareness and the cataloguing discipline; it is not to enter awareness as a candidate in any of the within-ontology disputes about what is to be catalogued.

The positive thesis is therefore not a residue left over after the eliminative arc has cleared the candidates away. It is the structural content of what the arc has established. Each prior chapter has shown that a candidate location of awareness fails because the candidate has to be articulated within a field whose subordination the candidate's articulation presupposes. The cumulative result is not merely that awareness is not in those locations; it is that awareness has been operative throughout the articulation of each candidate, in the structurally specific role of being what makes the articulation an articulation of anything at all. To call this role the standing-within-which is to name what each chapter's diagnosis has been pointing at: the standing within which each candidate's terms are intelligible as the terms they are.

The thesis is not a claim that ontology is illegitimate or that ontological inquiry should be abandoned. It is the opposite. Ontology can do its work — and has done extraordinary work — within its jurisdiction. The first-order disputes that ontologists have about whether properties are universals or tropes, whether grounding is primitive or analysable, whether priority monism is more parsimonious than smallism, whether the cosmos is a fundamental individual, whether composition is universal or restricted, whether materialism can accommodate consciousness — all of these continue to be questions the discipline takes up and addresses. The book's argument does not adjudicate them. It marks the limit at which ontology, as the cataloguing discipline, reaches its jurisdictional edge.

To see the limit clearly, take an example from Goff's view. Goff places consciousness+ as the intrinsic nature of the cosmos. This is an ontological claim with substantial implications: it tells us that the cosmos is one fundamental individual, that it has consciousness-involving properties, that ordinary subjects are aspects of it, that the grounding of ordinary subjects is by subsumption rather than analysis. Within ontology, the claim is articulate and contestable. It can be evaluated against rival ontological proposals (smallist micropsychism, panprotopsychism, emergentism, dualism, eliminative materialism). The evaluation is internal to ontology.

The book's claim is at a different level. It is that the standing within which Goff's claim is intelligible as an ontological claim — within which the cosmos can be identified as a candidate fundamental individual, consciousness+ as a candidate intrinsic property, subsumption as a candidate grounding relation, and so on — is not itself an ontological claim. The book is not in the same discipline as Goff at the level of the book's central thesis. The book's thesis is about the discipline's scope.

This is what is meant by "ontological jurisdiction." The discipline has a jurisdiction: a domain within which its methods, distinctions, and claims operate. Within the jurisdiction, ontologists can disagree productively, and the discipline can make real progress. There is also a limit to what the discipline can take as its subject. What awareness as field is — the standing within which the discipline's own operations are intelligible — is not subject to ontology's jurisdiction in the way items, properties, and

grounding relations are. The book's argument is the careful articulation of where that limit falls.

The articulation is jurisdictional rather than positively ontological because the jurisdictional thesis does not name a further item to be added. There is no further fundamental individual the book is naming; no further intrinsic property the book is identifying; no further grounding relation the book is positing. The articulation is a structural observation about where the cataloguing discipline finds the limit of its catalogue.

The book is therefore in a particular kind of relation to Goff's view. The book grants Goff's negative phase. The book does not enter the within-ontology disputes Goff engages in (with the smallist panpsychists, the emergentists, the physicalists, the dualists). The book marks the level-difference between Goff's constructive ontological proposal and the structural question the book has been pursuing. Goff has done his work within the discipline. The book has done its work at the discipline's jurisdictional edge.

What this finally amounts to as a philosophical result is worth saying directly. Marking the discipline's jurisdictional edge is not the same as standing outside the discipline and declining to enter. It is a substantive observation about what the discipline's results can and cannot be taken to settle. Within-ontology disputes — between Goff and the smallist micropsychist, between either and the emergentist, between any of them and the constitutive cosmopsychist's rivals — produce real results within their domain. What they cannot produce, the chapter's argument has shown, is a settlement of the field-question. That this is so is not internal to any of those disputes. It is a structural observation about what kind of question the field-question is, available only from the level at which the chapter has been arguing.

The book's contribution is therefore not the addition of one more position to the debates the book leaves untouched. It is the establishment of a result those debates cannot themselves establish: that the field-question is a question about the standing of any such debate as a debate about awareness in the relevant sense, and that no further item, property, or grounding relation will settle it. The book has not, by establishing this, ranked itself among the metaphysical positions ontology canvasses. It has marked where ontology's canvassing reaches its edge with respect to awareness as field — and the marking is the chapter's positive jurisdictional result, not the absence of one.

8.7 The jurisdictional result

The chapter's positive thesis is jurisdictional, not ontological. Ontology has its scope; the field-question is at the jurisdictional edge of that scope. The result is correspondingly narrow.

The chapter has not established that ontology is illegitimate, that ontological inquiry should be discontinued, or that the disputes among contemporary ontologists are confused. Ontology is one of philosophy's most articulate disciplines, and its first-order disputes have produced real progress and real understanding. The chapter has marked the limit of ontology's jurisdiction with respect to awareness as field, not with respect to ontology generally. Within its jurisdiction, ontology continues to do what it does.

Nothing in the chapter denies that consciousness, considered as a feature of experiential lives, can be the subject of substantive empirical and metaphysical inquiry. The neurosciences, cognitive psychology, phenomenology, philosophy of mind, and metaphysics each take consciousness as a subject and make claims about it that can be evaluated within their respective methods. The chapter's argument has been about the standing within which any such inquiry is intelligible as inquiry into consciousness, not about the inquiries themselves.

The chapter has not held that Goff has failed within his programme. *Consciousness and Fundamental Reality* is one of the most articulate contemporary works in the metaphysics of mind, and its arguments — the Consciousness Constraint, the Galileo argument, the case against physicalism via Phenomenal Transparency, the Subject Irreducibility argument, the move from grounding-by-analysis to grounding-by-subsumption, the constitutive cosmopsychist construction, the Power Realism argument — each repay careful study within their own terms. The chapter has not litigated any of them. Goff has done his work; the chapter's structural observation operates at a level adjacent to that work.

The chapter has not proposed a competing ontology to constitutive cosmopsychism. The field is not a further fundamental individual, not a further intrinsic property, not a further grounding relation, and not a further item within any catalogue. The chapter has resisted the temptation to offer a positive ontological characterisation of awareness throughout, because the central thesis is that awareness as field is not the kind of thing that has a positive ontological characterisation. The resistance is itself substantive — it is the chapter's main contribution — but it is not a competing posit.

Nothing in the chapter implies that awareness is unknowable, ineffable, or beyond reasoned discourse. The book has been treating awareness as a subject of reasoned philosophical inquiry throughout. What the chapter has marked is that the inquiry is not first-order ontological inquiry. The discipline required to articulate the relation of awareness to ontology is structural and jurisdictional. The book's arguments throughout have been examples of that discipline at work.

What the chapter has established is one thing, narrowly. Ontology, considered as the discipline that catalogues what exists and how the items in the catalogue depend on one another, has a jurisdiction within which its methods operate. Awareness as field is not subject to that jurisdiction in the way items, properties, and grounding relations are — not as something further or higher, but as the standing within which the discipline's cataloguing is itself intelligible as cataloguing. The chapter's central claim is that no constructive ontological proposal, however careful and however well-motivated by internal filters, answers the field-question; the field-question is not a question within ontology to which a new item, a new property, or a new grounding relation is the answer.

The next chapter turns to philosophical antecedents, with Advaita Vedānta as the primary source-grounded engagement. Other comparisons — Yogācāra, Husserl, Sartre, apophatic theology — will be made only where source consultation warrants them. The aim will not be to claim that the book recovers or supersedes these traditions, but to state the relation modestly: the thesis has deep precedents; the present book offers a disciplined eliminative argument for it.

Chapter 9 — Philosophical Antecedents

9.1 Why antecedents matter

The book's thesis is not new. Versions or relatives of the claim that awareness is the standing within which any ontological discourse is intelligible — not an item within the inventory of what appears, not a determinate property of any item, not a structure derivable from items more basic — have been articulated for over a millennium in the classical Indian and Buddhist traditions, and have echoes in the apophatic strands of European theological reflection, in the phenomenological tradition, and in parts of contemporary philosophy of mind. The contribution of the present book has been throughout methodological rather than first-doctrinal: a disciplined eliminative argument, run under a local referential constraint, against the principal ways awareness is mislocated within content-architectural features of conscious life.

A reader who has followed the eliminative arc may reasonably ask what bearing the classical articulations have on the book's argument, and what bearing the book's argument has on them. Both questions deserve careful answers.

The first answer is that the classical articulations supply a substantial structural precedent for the position the book defends. They show that the position has been thought through with great care and over long periods; that the analytical resources of the tradition are not exhausted by the contemporary philosophy of mind which has formed the bulk of the book's contemporary engagement; and that the structural shape of the field-thesis is older and more articulate than its contemporary versions sometimes acknowledge. The chapter takes this seriously. The book is not, on the structural matters, in a position to claim originality, and acknowledging the precedents is part of intellectual honesty.

The second answer requires more care. The book does not claim that any classical articulation contains the book's argument in disguise; does not claim that the book's argument recovers what the classical articulations were really aiming at; does not claim that the book supersedes, completes, purifies, secularises, or updates any classical position; and does not claim that the structural kinship licenses the book to take on the classical articulation's wider commitments. The classical articulations are theirs: their soteriologies, scriptural authorities, metaphysical commitments, and historical contexts are part of what they are. The book's argument shares a structural shape with what some classical traditions have articulated, and that shape is what the chapter engages, with the rest left as the tradition's own.

The chapter concentrates on the closest of these classical precedents: Śaṅkara's Advaita Vedānta. Of the available antecedents, Śaṅkara's articulation is the most directly comparable in structural form to the position the book defends. The witness-doctrine, the apparatus of *adhyāsa* (superimposition) as the diagnostic of mislocation, the method of *neti neti* (not this, not so), the careful distinction between the Self and the inner organ — each of these has a clear structural parallel in the book's prior chapters. The chapter does not claim that Śaṅkara's Advaita is *the* tradition antecedent to the book, nor that any other

tradition can be dismissed; it claims that Advaita is the precedent the chapter has the source-grounded basis to engage, and that this engagement is sufficient to demonstrate the structural kinship the chapter aims to mark.

Other comparison points deserve acknowledgement without substantive engagement here. The phenomenological materials considered in Chapter 7 offer nearer modern analogues, especially in Sartre's displacement of the ego from pre-reflective consciousness and Zahavi's account of pre-reflective self-awareness as constitutive of phenomenal consciousness. Advaita differs from these by embedding the witness-structure within a scriptural, soteriological, and metaphysical account of Ātman and Brahman. Yogācāra Buddhism and apophatic theology may offer further comparison points, but they are not engaged here because the source work has not been done.

The chapter is structured in five further sections. §9.2 reconstructs Śaṅkara's central position in his own terms, focusing on the *Upadeśasāhasrī* and the *Adhyāsabhāṣya* with Mayeda's and Gambhīrānanda's translations as guides. §9.3 marks the structural kinship between Śaṅkara's witness-doctrine and the book's field-thesis along four specific lines. §9.4 marks the doctrinal distance — the substantive commitments Advaita has that the book does not take on — along six specific lines. §9.5 says what the comparison licenses, and §9.6 closes by stating what the chapter establishes and what it does not.

9.2 Śaṅkara's Advaita: the witness and non-duality

Advaita Vedānta is not a single timeless block of doctrine. The tradition is over a millennium old and has been productive throughout; it has contained internal disputes between Śaṅkara (eighth century), Maṇḍana, Sureśvara, Padmapāda, Vācaspati Miśra, Sarvajñātman, and many others; it has been reformulated in contemporary scholarship by Ram-Prasad, Timalsina, Fort, and others. The chapter engages Śaṅkara's articulation specifically, with Mayeda's translation of the *Upadeśasāhasrī* and Gambhīrānanda's translation of the *Brahmasūtrabhāṣya* as the primary sources. The references throughout to "Advaita" should be read as "Śaṅkara's Advaita as preserved in these works."

Śaṅkara's central thesis is the non-difference of Ātman and Brahman: the innermost Self of every conscious being is identical with the metaphysical absolute. Mayeda calls this "not only the starting point of [Śaṅkara's] philosophy but also its goal" (Mayeda 1992, Introduction §II, p. 11). The thesis is concentrated in the great Upaniṣadic sentences (*mahāvākyas*) that Śaṅkara reads as expressing the identity: *tat tvam asi* ("Thou art That," Chāndogya Up. VI,8–16) and *ahaṃ brahmāsmi* ("I am Brahman," Bṛhadāraṇyaka Up. I,4,10). The whole of the *Upadeśasāhasrī* is dedicated to expounding this identity and to showing that knowledge of it is the means to final release (*mokṣa*) from transmigratory existence.

The thesis has three approaches in Śaṅkara, which Mayeda systematises as the theological/cosmological, the psychological/epistemological, and the exegetical (Mayeda Introduction §III). The chapter concentrates on the psychological/epistemological approach, since this is where the structural kinship with the book's argument is most direct, and on the diagnostic apparatus of the *Adhyāsabhāṣya*, since this contains Śaṅkara's most direct statement of the diagnostic of mislocation.

The Self as Pure Consciousness. The opening verse of the *Upadeśasāhasrī*'s metrical first chapter (*Caitanyaprakaraṇa*, "Pure Consciousness") sets the keynote: "Salutation to the all-knowing Pure Consciousness which pervades all, is all, abides in the hearts of all beings, and is beyond all objects [of knowledge]" (Upad I,1,1, Mayeda translation). The Self (*ātman*) is *caitanya*, Pure Consciousness, an unqualified subject that pervades experience without being any of its determinate contents. It is "all-knowing" not in the sense of containing infinite contents but in the sense of being the consciousness within which any content is known.

The witness (sākṣin). The Self stands in a particular structural relation to its objects: it sees them; they do not see it. Śaṅkara articulates this in the metrical chapters on the seer (*draṣṭṛ*) and the witness. Verses from chapter 13 of the metrical part state the position with characteristic firmness: "The modification of the *manas*, which is caused by the visual sense and is depicted by form-colour [of its object], is certainly always seen by the constant Seeing of Atman" (Upad I,13,6); "The Seeing of the Seer is, therefore, constant, pure, infinite and alone" (Upad I,13,8). Mental events — perception, memory, the modifications of the inner organ in dreaming — are objects of the witness. The witness itself is not an object of further perception, but is what perceives.

The Self is not the inner organ. The careful distinction between the Self and the *antaḥkaraṇa* (the inner organ, comprising the *buddhi* or intellect, the *manas* or mind, the *ahaṃkāra* or "I"-making, and the *citta*) is one of Śaṅkara's most rigorously developed doctrines. The "I"-notion (*ahaṃpratyaya*) — the everyday thought "I do this, I feel that, I think the other" — is on Śaṅkara's analysis a feature of the *buddhi* acting as the bearer of *ahaṃkāra*. It is not a feature of the Self. The Self is the *object* of the "I"-notion, in the sense that the notion is directed toward the Self, but the Self is not constituted by the notion (Upad II,2,52–53; *Brahmasūtrabhāṣya* II,3,38, Gambhīrānanda p. 545). The discriminating reflective thought "I am the knower, not the object of knowledge, pure, always free" (Upad I,12,14) is itself a feature of the *buddhi* — a state of the inner organ that has the Self as its object. The Self witnesses this discriminating thought, as it witnesses other states of the inner organ, without being any of them.

The negation method. The most familiar of Śaṅkara's methodological tools is the *neti neti* method, derived from Bṛhadāraṇyaka Up. II,3,6: "Not thus! Not so!" The Self is not the body, not the senses, not the inner organ, not the experiencer, not the agent, not any of the items it can be mistakenly identified with. Negation proceeds until what remains is what cannot be negated. The metrical Chapter 2 of the *Upadeśasāhasrī* (*Pratiṣedhaprakaraṇa*, "Negation") opens with this principle: "As [Atman] cannot be negated, [It] is left unnegated [by the Sruti, 'Not thus! Not so!']" (Upad I,2,1). "The Seeing (= Atman) is one alone, self-established" (Upad I,2,3). What cannot be negated, on Śaṅkara's account, is the witness — the Self.

The diagnostic of superimposition. Śaṅkara's most direct statement of the bondage-condition is the *Adhyāsabhāṣya*, his prefatory essay to the *Brahmasūtrabhāṣya* (Gambhīrānanda pp. 1–28). The essay opens with the observation that the Self and the not-Self are absolutely opposed in nature — the Self being subject, the not-Self being object — yet "the mutual superimposition of the Self and the not-Self, known as nescience, is the presupposition on which proceed all the empirical dealings involving the means of knowledge, objects of knowledge, and the results thereof" (Gambhīrānanda p. 5). The paradigm case of superimposition (*adhyāsa*) is the rope mistaken for a snake: in the half-light, a coiled rope is perceived as a snake; the snake is "superimposed" on the rope. Śaṅkara extends this everyday phenomenon to the metaphysical level. The not-Self (the body, senses, intellect, agency, experience) is superimposed on the Self; reciprocally, characteristics of the Self (consciousness, intentionality, agency) are superimposed on the not-Self (the inner organ illuminated by consciousness comes to appear as the bearer of consciousness itself). The result is the everyday self-conception in which one identifies oneself with the body, the inner organ, the agent of action.

The unreality of bondage and the path to release. Bondage, on Śaṅkara's account, is the operation of avidyā (nescience) producing the appearance of separation between Ātman and Brahman. The path to release is knowledge (*vidyā*) — specifically, knowledge of the Ātman–Brahman identity. When this knowledge is firmly grasped, the avidyā that has been operating beginninglessly is dissolved, and what remains is the recognition of what has always been the case: the Self is Brahman. Mokṣa, the final release from transmigratory existence, is not the acquisition of a new state but the recognition of the state that was always already present. The dialogue of the prose Chapter 2 of the *Upadeśasāhasrī* (*Av-*

agatiprakaraṇa, "Awareness") works through this position via the teacher–student exchange: the student asks how release from transmigration is possible; the teacher explains that the cause is avidyā and the means is knowledge; the student presses on what avidyā is and how it could affect the Self if the Self is by nature pure; the teacher works through the apparatus of superimposition and the resolution it offers. The dialogue is the philosophical centrepiece of the work.

This is the central shape of Śaṅkara's Advaita. The next two sections mark the structural kinship and the doctrinal distance between this position and the book's argument.

9.3 Structural kinship: negation, witness, non-identification, self-luminosity

The structural kinship between Śaṅkara's Advaita and the book's central thesis runs along four lines. The chapter notes each in turn, with care to mark that the kinship is structural rather than doctrinal — that the book's argument shares a shape with Śaṅkara's, not a content.

First, the method of negation. The book's prior chapters have proceeded by negation. Chapter 3 negated the possibility of ontologically deriving awareness from any item determinate enough to play the deriving role. Chapter 4 negated the identification of awareness with content-architectural relations (higher-order representation, functional access, global workspace, attention). Chapter 5 negated the illusionist and meta-problem reframings that would dissolve or relocate the problem of awareness rather than answer it. Chapter 6 negated the inference from time as the form of inner appearance to time as a condition governing awareness. Chapter 7 negated the inference from perspectival plurality to the plurality of awareness-fields. Each chapter has proceeded by ruling out a candidate location for awareness, without proposing a positive ontological location in its place.

Śaṅkara's *neti neti* method has the same procedural shape. The Self is reached, on his account, by negating what it is not — by negating the body, the senses, the inner organ, the agent, the experiencer, and the candidates that ordinary self-understanding offers — until what remains is what cannot be negated. The shared procedure is striking: in both cases, the method is eliminative rather than constructive; in both cases, what is being arrived at is not a further posit but the standing within which the posits are negated.

The kinship is structural. It does not entail that Śaṅkara's apparatus is the book's apparatus. Śaṅkara has Sruti as his warrant for the negation method; the book has structural argument. Śaṅkara's negation reaches *brahman* as the unsublatable; the book's negation reaches the field as the standing within which the negation proceeds. The shared shape — eliminative procedure converging on an unnegatable standing — is what the kinship marks. The differences in what the unnegatable standing is taken to be, and in how it is warranted, are real and considerable.

Second, the witness/standing language. The book's prior chapters have spoken of the *field within which* items, properties, structures, dependence relations, grounding relations, and fundamental individuals are intelligible. Chapter 8 made the jurisdictional thesis explicit: ontology has a jurisdiction within which its cataloguing methods operate, and awareness as field is not within ontology's jurisdiction as an item; awareness is the standing within which the cataloguing is itself intelligible as cataloguing. The book's term has been "field"; the structural function of the term has been the standing-within-which.

Śaṅkara's *sākṣin* — the witness — does substantially the same structural work in his apparatus. The witness sees but is not seen; is constant where its objects are variable; is pure where they are mixed; is alone where they are plural; is self-luminous where they are illuminated. The witness is the standing within which any seeing occurs. The structural correspondence to the field-thesis is direct.

Again, the kinship is structural. Śaṅkara's *sākṣin* is loaded with specific metaphysical content (it is the Self; the Self is Brahman; Brahman is the metaphysical absolute) that the book does not take on. The book's field is not the *sākṣin*; the book's field is not the Self; the book's field is not the absolute. But the structural function of the *sākṣin* in Śaṅkara's account — the standing-within-which that is itself not one of the things-within — resembles the structural function of "field" in the book's argument. The resemblance is real and worth marking.

Third, the resistance to identification with the inner organ. The book's Chapter 4 argued that the principal contemporary identifications of awareness with content-architectural relations — Carruthers's higher-order theory, Block's distinctions on access and phenomenal consciousness, the various functional and representational identifications — fail because they locate awareness within a determinate content-architectural feature, and whatever is determinate enough to play the identifying role is itself content within the field rather than the field within which content has its standing.

Śaṅkara's careful distinction between the *ātman* and the *antaḥkaraṇa* is structurally the most direct classical analogue to the book's Chapter 4 argument. The *antaḥkaraṇa* is the inner organ — the buddhi, manas, ahaṃkāra, and citta as developed in classical Indian psychology. It is the seat of cognitive operations: perception, judgement, memory, intention, agency, the "I"-notion. Śaṅkara is rigorous about denying that any of these operations is the Self. The Self is what *sees* the cognitive operations; the Self is what makes them perceivable as cognitive operations; but the Self is not one of them. The verses of the metrical *Upadeśasāhasrī* are clear on this: "Sankara maintains that pain is an object (*viṣaya*) of *buddhis*, which are in turn objects of the inner Ātman" (Upad I,18,201, paraphrased Mayeda Introduction §III.B.4, p. 41). Pain is in the *buddhi*; the *buddhi* is in turn an object for the witness; the witness is not in the *buddhi*.

The structural correspondence to the book's Chapter 4 is close. Both arguments resist the identification of awareness with cognitive content-architectural features; both maintain that the cognitive features are within the field of awareness rather than constitutive of it; both rest on a structural rather than empirical observation about the kind of thing awareness is. The classical four-part inner-organ taxonomy is the closest analogue in classical Indian thought to contemporary content-architectural theory, and Śaṅkara's resistance to identifying the Self with any of the four parts is the closest classical analogue to the book's Chapter 4 resistance to contemporary content-architectural identifications.

Fourth, the self-luminosity of consciousness. The Advaita doctrine of *svayaṃprakāśa* (self-luminosity) is the claim that consciousness is not made manifest by another consciousness, but is manifest in itself. Other items in the world are illuminated by consciousness; consciousness is not illuminated by anything further. Mayeda summarises Śaṅkara's position: "Atman-Brahman is self-evident (*svapramāṇaka*) and self-established (*svataḥsiddha*)… Therefore, Atman-Brahman is by nature independent of the means of knowledge" (Mayeda Introduction §III.C.1, p. 47, citing Upad I,18,203 and Upad II,2,93).

The book's argument at Chapter 4 against higher-order theories shares this structural shape. Representations of representations cannot, on the book's argument, deliver first-personal givenness because first-personal givenness is not delivered by representation in the first place — it is what makes the representations be the representations they are, available as the kind of thing they are to a subject. Zahavi's contemporary phenomenological articulation of this point in the *Subjectivity and Selfhood* engagement (Chapter 7 §7.3.3) traces the same structural claim through Husserl, Sartre, and Henry in the European phenomenological tradition. The Advaita articulation is older and developed in a different idiom, but the structural shape — consciousness as self-luminous rather than illuminated-by-another — is the same.

The four points of structural kinship are substantial. The book's argument has, at the structural level, a

classical articulation of the position it defends. The chapter takes this seriously. The next section marks the doctrinal distance with equal seriousness.

9.4 Doctrinal distance: Brahman, Śruti, mokṣa, māyā, saṃsāra

The structural kinship between Śaṅkara's Advaita and the book's argument is real. The book and Advaita part company at six specific points of doctrine, each of which carries substantial content that the book does not take on. The chapter marks each in turn.

First, Brahman as the metaphysical absolute. Śaṅkara's identity-thesis places Brahman as the ultimate metaphysical reality — the absolute, the substance and cause of the universe, the unconditioned source from which the manifest world proceeds. The *Brahmasūtra* opens with the famous definition: Brahman is "that from which the origination, subsistence, and dissolution of this universe proceed" (BS I,1,2, cited in Mayeda Introduction §III.A.1, p. 18). Brahman is both the efficient and the material cause of the universe (BS I,4,23–27). The creation is "nothing but self-creation" (*ātmakṛti*, BS I,4,26). Brahman is the metaphysical absolute in the strongest sense the tradition has used the term.

The book makes no such claim. The book's argument in Chapter 8 was that ontology has a jurisdiction within which its cataloguing methods operate, and that the field is not within that jurisdiction as an item to be catalogued. To place the field as the metaphysical absolute would be to place it within ontology as the ultimate item — to take precisely the move the book's argument resists. Brahman, on Śaṅkara's account, is the absolute within an ontology of fundamental and non-fundamental reality; the book's field is not such an absolute, because the field is not within the kind of ontology in which "absolute" picks out a candidate posit.

The chapter must not import the ontology vocabulary of Chapter 8 into Śaṅkara, nor import Śaṅkara's metaphysical absolute into the book's field-thesis. The book's argument does not, and could not, deliver Brahman as a conclusion; Brahman is Śaṅkara's conclusion, reached through the apparatus of scriptural authority and metaphysical reasoning that is his, not the book's.

Second, sat-cit-ānanda as positive characterisation. The traditional positive characterisation of Brahman — sat-cit-ānanda, Being-Consciousness-Bliss — is part of Advaita's substantive metaphysics. The book does not characterise awareness in any of these ways. The book has not attributed *being* to the field in the sense in which Advaita attributes being to Brahman; the field is not a fundamental individual, not a substance, not a metaphysical item to which being-predicates apply. The book has not characterised awareness as bliss (the ānanda component); the soteriological framing within which the bliss-claim has its force is not the book's framing, and the chapter must avoid sliding into language that would suggest awareness is positively characterised by anything like ānanda. The *cit* (consciousness) component is the closest to the book's concern, but Advaita's *cit* is the consciousness of Brahman as ultimate reality, not the field as the book has been developing it; the difference matters.

Third, Sruti as the warrant. Advaita's epistemology is grounded in the authority of the Upaniṣads. Mayeda is explicit: "Atman-Brahman, the knowledge of which is the means to final release (mokṣa), is self-evident and self-established. Therefore, Atman-Brahman is by nature independent of the means of knowledge… Then how is the knowledge of Atman-Brahman obtained? It is attained only through the Sruti" (Mayeda Introduction §III.C.1, p. 47). Śaṅkara holds that the Upaniṣads are infallible; that no argument or justification is necessary for their validity; that the other means of knowledge (sense-perception, inference) are valid only within the sphere of nescience and before the attainment of Ātman-knowledge.

The book makes no such appeal. The argument throughout has been structural and analytical; the

warrant has been inferential rather than scriptural. The book does not depend on the Upaniṣads being infallible, on any specific scriptural tradition being authoritative, or on any extra-philosophical source of knowledge. The chapter must not appear to take on Sruti as the warrant for the field-thesis, even implicitly. The kinship with Śaṅkara is structural; the methodological difference about warrant is a substantial divergence, and it is part of what separates the book's argument from the tradition's articulation.

Fourth, mokṣa as the orienting purpose. Advaita is a path. The doctrines of avidyā, adhyāsa, sākṣin, and Brahman–Ātman identity are framed throughout as therapeutic: they aim at the dissolution of suffering and the recovery of the Self's original perfection. *Mokṣa* — release from saṃsāra (transmigratory existence) — is the orienting purpose of Śaṅkara's whole project. The *Upadeśasāhasrī* opens with the question of release; it closes with instructions for the practice (*parisaṃkhyāna* meditation) that prepares the seeker for the recognition of the Ātman–Brahman identity.

The book has no soteriological project. The chapter must not appear to take on the framing within which Advaita's doctrines have their orienting purpose. The structural arguments of the prior chapters did not aim at release; they aimed at structural clarification. The book's argument does not deliver a path, does not name a goal beyond philosophical clarity, does not endorse the soteriological framing within which Advaita's doctrines have their force. The chapter must mark this difference plainly. The book and Advaita share a structural commitment to the witness-doctrine; they do not share a project.

Fifth, māyā as the ontological status of the empirical world. Advaita's mature doctrine treats the empirical world as appearance — real in *vyavahāra* (the empirical standpoint), unreal in *paramārtha* (the standpoint of highest truth). The doctrines of *māyā* and *vivarta* (illusion-transformation, the Advaita interpretation of the world's relation to Brahman) make this substantive ontological claim about the manifest world.

The book takes no such position. Empirical reality is empirical reality, on the book's view. The chapters that preceded this have not argued, anywhere, that the world of items is illusory; the eliminative arguments have been about the location of awareness, not about the reality of the world. The book's argument in Chapter 8 was explicit that ontology continues to do its work within its jurisdiction; the question of which items the catalogue should include and how the inventory should be arranged is a first-order ontological question on which the book takes no position. The chapter must not slide toward Advaita's *māyā* doctrine. The book's argument does not deliver, and the chapter must not be read as endorsing, the unreality of the manifest world.

Sixth, saṃsāra and karman as framing context. Advaita assumes the framework of saṃsāra (transmigratory existence), karman (the law of action and consequence across lives), and rebirth. The *Upadeśasāhasrī* opens with the student's question: "How can I be released from transmigratory existence?" The structural apparatus of the work — pure consciousness as the Self, avidyā as the cause of bondage, knowledge as the means to release — is articulated within the broader cosmological framing of repeated lives in which the consequences of action propagate and from which the Self seeks release.

The book takes no position on this framework. The chapter must engage Advaita without taking on the cosmological assumptions that frame Śaṅkara's project. The structural kinship at the witness-doctrine level is real; the framing context within which Śaṅkara articulates the witness-doctrine is not the book's framing. The chapter must engage Śaṅkara on the witness-doctrine without implicitly endorsing the saṃsāra cosmology.

The six points of distance are substantial. The book's argument does not take over Advaita's metaphys-

ical absolute, its positive characterisation of consciousness as sat-cit-ānanda, its scriptural warrant, its soteriological project, its doctrine of māyā, or its saṃsāra cosmology. The kinship is structural; the doctrinal framework that surrounds it in Śaṅkara remains Śaṅkara's.

9.5 What the comparison licenses

What does the structural kinship license? Three things, the chapter argues, and no more.

First, the comparison licenses *acknowledgement of precedent*. The book's central thesis is not novel in spirit. The structural shape of the position the book defends — awareness as the standing within which any ontological discourse is intelligible, not an item within the inventory, not a determinate property of any item, not a structure derivable from items more basic — has been articulated for over a millennium in Advaita Vedānta. The acknowledgement matters. It locates the book within a long tradition of careful thought on the structural shape of awareness, and it removes any pretence that the book is doing something unprecedented at the level of structural commitments. The contribution of the book has been methodological — a disciplined eliminative argument under a local referential constraint — rather than first-doctrinal. The acknowledgement of precedent is part of that methodological honesty.

Second, the comparison licenses *structural confidence*. A position the book reaches by inferential argument has been reached by classical tradition through different paths — scriptural exegesis, meditation, sustained dialogue across generations of careful thinkers — and articulated with comparable structural specificity. The convergence is not, by itself, a vindication of the book's argument; the classical articulation and the book's inferential argument could both fail, despite the structural convergence. But the convergence is some evidence that the structural shape is not an artefact of the contemporary intellectual context in which the book has been written. The shape has been arrived at independently. The book takes this evidence with appropriate weight: as a defeasible reason to think that the structural claim has stable shape across very different intellectual frameworks.

Third, the comparison licenses *modest dialogue*. The book and Advaita can be in conversation at the structural level without either party needing to take on the other's framework. The book's structural argument can engage Advaita's witness-doctrine as the most articulate classical statement of the position the book defends, while leaving Advaita's wider doctrinal commitments intact. Advaita does not need to be modernised, secularised, purified of its scriptural or soteriological commitments to be in conversation with the book. The book's argument does not need to take on Advaita's commitments to share the structural kinship. The dialogue can be modest on both sides.

What the comparison does not license is more important to state clearly.

It does not license the book to claim that the book has recovered what Advaita was really aiming at. The chapter's engagement is structural; the claim that the book has uncovered Advaita's structural argument while bracketing the doctrinal apparatus would amount to the move the brief has explicitly ruled out. Advaita's articulation has its scriptural warrant, its soteriological orientation, its metaphysical commitments; these are not incidental decoration around a structural thesis the book has now identified. They are part of what Advaita is. The book's argument and Advaita's tradition share a structural shape, and that is what the engagement marks. It is not a claim about what Advaita meant.

It does not license the book to take on Advaita's authority as a warrant for the field-thesis. The book's argument has been structural throughout; its warrants have been inferential. The chapter's appeal to Advaita is an acknowledgement of structural precedent, not an appeal to scriptural or traditional authority. If a reader is not convinced by the book's structural arguments, the existence of an Advaita articulation of the same structural position is not, by itself, a reason for the reader to accept the position.

Conversely, if a reader accepts the book's structural arguments, the structural kinship with Advaita supplies historical depth and contemporary resonance, but it does not strengthen the arguments themselves. The arguments stand or fall on their own structural ground.

It does not license the book to claim that Advaita is a "version" of the book's argument. Advaita is its own tradition with its own commitments and its own internal disputes. The book is not a contemporary form of Advaita. The book is, at the structural level, in agreement with a structural commitment of Advaita, while having its own argumentative structure, its own contemporary engagements, and its own methodological character.

The phenomenological materials considered in Chapter 7 offer nearer modern analogues, especially in Sartre's displacement of the ego from pre-reflective consciousness and Zahavi's account of pre-reflective self-awareness as constitutive of phenomenal consciousness. Advaita differs from these by embedding the witness-structure within a scriptural, soteriological, and metaphysical account of Ātman and Brahman. The chapter does not re-engage the phenomenological material here, having addressed it in Chapter 7; the cross-reference is sufficient to mark that the modern analogues exist and have been engaged in the book's earlier chapters. Yogācāra Buddhism and apophatic theology may offer further comparison points, but they are not engaged here because the source work has not been done.

9.6 The kinship and its limit

The kinship with Śaṅkara's Advaita is real, but limited. The chapter's result is structural, not doctrinal.

The chapter has not established that Advaita is correct. Advaita is a metaphysical, soteriological, and exegetical tradition with substantive commitments — Brahman as the ultimate metaphysical absolute, the authority of the Upaniṣads, the path to mokṣa from saṃsāra, the doctrine of māyā about the empirical world. The chapter has not argued for any of these. The chapter has noted that the witness-structure within Advaita is the closest classical articulation of the structural position the book defends; this is a structural observation, not an endorsement of Advaita's doctrinal framework. A reader who is sceptical of the doctrinal framework can recognise the structural kinship without taking on the doctrines.

Nothing in the chapter holds that the book is a contemporary version of Advaita. Advaita is its own tradition, articulated over a millennium across many thinkers and many disputes; the book is a contemporary argument with its own contemporary engagements, its own argumentative structure, and its own methodological character. The structural kinship at the witness-doctrine level does not make the book a contemporary Advaita, and the chapter has avoided the language that would suggest such a claim.

The chapter has not claimed to recover, complete, secularise, purify, update, or supersede Advaita. The traditional doctrines remain the tradition's doctrines: Ātman–Brahman identity, sat-cit-ānanda, Sruti as warrant, mokṣa as orienting purpose, māyā as the ontological status of the empirical world, saṃsāra and karman as the framing context. The book's structural engagement with the witness-doctrine does not take any of these on; nor does it imply that they need to be left behind for the structural witness-doctrine to be philosophically operative. They are Advaita's commitments; the chapter has marked them as Advaita's.

Nothing in the chapter denies the existence of other possible antecedents or parallels. Yogācāra Buddhism, apophatic theology, Neoplatonism, Mādhyamika, Sufi metaphysics, and Daoist traditions may each warrant careful comparison in other work. This chapter has not undertaken that work. Their absence here is a matter of source-discipline, not a denial of their importance.

The chapter has not entered Advaita's internal disputes. The relation of avidyā to Brahman; the doctrine of *anirvacanīya* (the "indescribable" status of avidyā as neither real nor unreal); the disagreements between Śaṅkara and his immediate disciples Sureśvara and Padmapāda; the later Vivaraṇa and Bhāmatī schools of Advaita interpretation; the modern reformulations by Vivekānanda, Aurobindo, and others; the contemporary scholarly reassessments — none of these has been engaged in the chapter. The chapter has rested on Śaṅkara's primary texts (the *Upadeśasāhasrī* and the *Brahmasūtrabhāṣya*'s *Adhyāsabhāṣya*) and on Mayeda's and Gambhīrānanda's translations as the source-grounded basis of the engagement.

What the chapter has established is one thing, modestly. Śaṅkara's Advaita is the closest classical precedent for the structural position the book defends; the kinship runs along four specific lines (the negation method, the witness/standing language, the resistance to identifying consciousness with the inner organ, the self-luminosity of consciousness); the kinship is structural rather than doctrinal; and the chapter has acknowledged the precedent without taking on Advaita's wider doctrinal framework. The book's argument retains its own structural character, its own contemporary engagements, and its own methodological orientation. Advaita retains its own scriptural, soteriological, and metaphysical commitments. The two are in modest dialogue at the structural level.

The chapter that follows is the Conclusion, where the principal results are restated in modest form, open problems are identified, and the relation of the eliminative arc to further work is briefly indicated.

Conclusion — What Has and Has Not Been Shown

C.1 The thesis restated

The book argues for one thesis.

Awareness is not one item within the inventory of what appears, but the field within which any such inventory becomes intelligible.

The thesis has three connected negative components and one positive observation. Awareness is not ontologically derivative from anything determinate enough to play the role the derivation would assign it. Awareness is not identical to any content, content-relation, function, representation, or access architecture. Awareness is not subject to the kinds of structural conditioning — temporal, perspectival, modal — that operate on items within the field. The positive observation is that awareness, considered as field, is the standing within which any cataloguing of items, properties, structures, and grounding relations is itself intelligible as cataloguing.

The book has not shown that awareness is a substance, an object, a cosmic subject, a metaphysical absolute, or a hidden item somewhere within or behind the items the ordinary inventories of experience and science catalogue. It has shown that the principal attempts to derive, identify, reframe, temporally subordinate, perspectivally multiply, or ontologically place awareness each presuppose the field within which their own terms become intelligible. The thesis is a structural result, not a positive ontological posit. The argument is jurisdictional, not metaphysical in the standard first-order sense.

The contribution of the book is methodological rather than first-doctrinal. The thesis has deep philosophical precedent in Śaṅkara's Advaita Vedānta, in the phenomenological tradition's treatments of intentionality and pre-reflective self-awareness, in parts of contemporary philosophy of mind that resist deriving consciousness from non-conscious matter, and in traditions that may warrant later comparison but have not been engaged here. What the book has offered is a disciplined eliminative argument, run under a local referential constraint, against the principal contemporary ways in which awareness has been mislocated within content-architectural features of conscious life.

C.2 What the eliminative arc has shown

The argument has moved through nine chapters, each engaging a specific candidate way of locating awareness within the inventory the book's structural argument is about. The results of the arc can be gathered in summary form.

Chapters 1 and 2 set the methodological ground. Chapter 1 introduced the basic distinction between awareness and determinate content, and the mislocation diagnosis: the principal contemporary theo-

ries of consciousness, however sophisticated within their domains, take consciousness as an item within an explanatory field whose intelligibility already presupposes awareness. Chapter 2 developed the local referent constraint as a narrowly specified diagnostic tool for testing whether the load-bearing terms in any ontological dependence claim about awareness have enough determinacy to support the claim. The constraint is not a general theory of meaning; it is a local test, deployed only where the book's structural arguments require it.

Chapter 3 — the Bridge against ontological derivation — argued that no derivation of awareness from items more basic in the ontological order can be stated without already presupposing the field of intelligibility within which the items, the dependence relation, and the awareness being derived are themselves intelligible. The chapter engaged the Stroud-style transcendental objection at length and distinguished the book's argument from a standard Kantian move from conditions of knowledge to conditions of being. The book's conclusion is narrower: any proposed derivation of awareness must be made intelligible within the very field it proposes to subordinate.

Chapters 4 and 5 turned the same form of argument on identification and reframing. Chapter 4 engaged Carruthers's higher-order theory and Block's distinction between access and phenomenal consciousness, arguing that identification claims of either kind locate awareness within a determinate content-architectural feature, and that whatever is determinate enough to play the identifying role is itself content within the field rather than the field within which content has its standing. Chapter 5 engaged Frankish's illusionism and Chalmers's meta-problem of consciousness, arguing that explaining our intuitions and reports about consciousness leaves untouched the field within which the intuitions, the reports, and the explanations are themselves intelligible.

Chapter 6 developed the temporal subordination case. Kant's account of time as the form of inner sense was granted within its proper jurisdiction — the form of organised inner appearance — and the argument located the pressure-point at the move from this jurisdiction to a condition governing the field within which any organisation of inner appearance is itself intelligible. The result was structural rather than anti-Kantian: time conditions content-experience; the field is not what such conditioning specifies.

Chapter 7 took up perspectival plurality. Embodiment, sensory access, memory, expressive resources, and intersubjective correction richly individuate perspective-loci within experience. The chapter granted all of this and resisted only the inference from the plurality of perspective-loci to the plurality of awareness-fields. The phenomenological accounts of Husserl, Sartre, and Zahavi were engaged at the level of perspective-loci within a field, with Husserl's monadological intersubjectivity and Zahavi's multidimensional account given particular attention; the chapter argued that no constructible principle individuates fields by counting perspective-loci.

Chapter 8 — Ontological Jurisdiction — consolidated the book's positive thesis about the relation of awareness to ontology. The chapter was structured around the engagement with Goff's constitutive cosmopsychism, the most articulate contemporary defence of an ontology that takes consciousness as fundamental. Goff is the closest contemporary ally on the negative phase of the book's argument: he agrees that consciousness cannot be derived from non-conscious matter, that the standard physicalist identifications fail, and that illusionism does not deliver what an account of consciousness needs to deliver. The book and Goff diverge at the constructive step. Goff places consciousness as the intrinsic nature of the cosmos within an ontological architecture; the book argued that the field-question concerns the standing within which any such architecture is intelligible as architecture, and that the field is not within ontology's jurisdiction as an item to be catalogued. The chapter's central claim is the jurisdictional limit: ontology has its scope; awareness as field is at the jurisdictional edge.

Chapter 9 engaged the closest classical precedent. Śaṅkara's Advaita Vedānta has articulated a structurally adjacent position, in a different idiom and within different doctrinal commitments, for over a millennium. The chapter marked four specific points of structural kinship — the negation method, the witness-doctrine, the resistance to identifying consciousness with the inner organ, and the self-luminosity of consciousness — and six specific points of doctrinal distance: Brahman as metaphysical absolute, sat-cit-ānanda as positive characterisation, Sruti as scriptural warrant, mokṣa as soteriological purpose, māyā as the ontological status of the empirical world, and saṃsāra as framing cosmology. The book takes on none of the six doctrinal commitments. The kinship is structural; the doctrines remain Advaita's.

The cumulative result of the arc is what the book aims for. No individual chapter delivers more than its narrow procedural result; the arc, taken together, marks where ontology reaches its jurisdictional edge with respect to awareness as field.

C.3 What remains open

The book's argument is procedural and structural; it has not resolved, and was not designed to resolve, the substantive questions that lie within the disciplines its argument operates adjacent to. Several open problems deserve explicit acknowledgement.

The relation between the field-thesis and empirical consciousness science. The book has been jurisdictional about ontology, not about the empirical sciences of mind. Neuroscience, cognitive psychology, and the experimental investigation of consciousness continue within their proper domains, and the book's argument does not, and cannot, dictate the results of their investigations. What remains open is the question of how the field-thesis bears on the interpretation of empirical findings: what neural correlates of consciousness are taken to establish, what the integrated-information formalism is taken to model, what the discoveries of cognitive psychology are taken to show about the structure of consciousness as a feature of experiential life. The book's argument is that these findings establish what they establish about content-architectural features of conscious experience, and that what they establish does not extend to the field within which such features are themselves intelligible. The translation of this jurisdictional limit into specific interpretive guidance for the sciences is work the book has not undertaken.

The relation between the field-thesis and first-person phenomenology. The phenomenological tradition has developed careful accounts of the structures of conscious experience — intentionality, temporality, embodiment, intersubjectivity, the pre-reflective givenness of the self. The book has drawn on this tradition in Chapter 7 (Husserl, Sartre, Zahavi) and at points throughout. What remains open is the question of how the field-thesis bears on the future of phenomenological inquiry: whether phenomenology can continue to articulate the structures of conscious experience while taking on board the structural observation that awareness as field is not one of the structures articulated. The provisional answer is that phenomenology can do so, and that the field-thesis is compatible with continued phenomenological work at the structural level the tradition has been operating at. But the detailed working-out of this compatibility is not the book's project.

The status of intersubjective communication after the rejection of field-plurality. Chapter 7 argued that perspective-loci are real and structurally describable, while the inference from perspective-loci to field-plurality is not constructible. What remains open is how to understand the substantive practices of intersubjective communication, mutual correction, testimony, joint perception, and the shared world of human life. The book's position is that these practices are intersubjective at the level of perspective-loci within a field, and that the field is not what they individuate. The detailed articulation of this

position — how the practices work, what conditions enable them, what they presuppose — is left to phenomenology and to the relevant empirical and conceptual investigations.

The relation between structural argument and classical contemplative traditions. Chapter 9 engaged Śaṅkara's Advaita as the closest classical precedent for the structural position the book defends. What remains open is the relation between the book's structural argument and the contemplative practices within which Advaita and related traditions have articulated their structural commitments. The book has been a piece of analytic philosophy throughout; the structural arguments it has developed are inferential rather than experiential. Whether the structural shape the book articulates by inference can be approached by other means, whether contemplative practice illuminates the structural shape in ways argument cannot, and how the two kinds of approach bear on one another — these are questions the book has not addressed and does not undertake to address.

The remaining scholarly engagements. Several engagements named in the brief have been retained as source-pending notes rather than developed into chapter prose: Stroud, Stern, Gomes, and Allais on transcendental argumentation (relevant to Chapter 3); contemporary Kantian commentators including Guyer, Strawson, Allais, Longuenesse, Onof, and Schulting (relevant to Chapter 6). These remain available for footnotes or later revision. The argument as it stands does not depend on them being completed; their inclusion would strengthen the historical and scholarly density of the book's contemporary engagements without altering the structural arc.

Each of these open problems is genuinely open. The book has not solved them and does not claim to. They mark the points at which the book's structural argument hands work back to other modes of inquiry.

C.4 What the argument hands back to inquiry

Neuroscience, cognitive science, phenomenology, and ontology remain valid within their domains. The argument changes what they can be taken to establish about awareness as field.

For neuroscience and cognitive science, the change is interpretive rather than methodological. The empirical investigation of neural correlates, functional architectures, perceptual processes, and conscious states proceeds as it has been proceeding. What changes is the inference from results to conclusions about awareness as field. A finding about the neural correlates of conscious access establishes what it establishes about the correlates of access — the dependencies, the structures, the conditions under which the access-features of experience occur. It does not establish that awareness as field has been located. The empirical findings retain their proper jurisdiction; the inference to conclusions about awareness as field is the move the book's argument has resisted, repeatedly, across the chapters.

For phenomenology, the change is closer to acknowledgement than to interpretive correction. The phenomenological tradition has been articulating the structures of conscious experience with care, and the book's argument has drawn on this articulation. What the book adds is a structural observation about the level at which phenomenological articulation operates: the structures it articulates are within-field structures; the field within which articulation has its standing is not one of the structures articulated. This observation does not displace phenomenological work. It marks a level-difference between the work and the standing of awareness as field.

For ontology, the change is jurisdictional. The first-order disputes of contemporary ontology — about properties as universals or tropes, about grounding as primitive or analysable, about composition as restricted or universal, about priority monism, about the nature of fundamental individuals, about modal commitments, about the materialism question — continue within their proper jurisdiction. The book's

argument does not adjudicate them. What the book argues is that the question of awareness as field is at the jurisdictional edge of these disputes, and that no further item within the ontological inventory will deliver an answer to the field-question. Within the inventory, the ontologists do their work. The field-question is not the kind of question the inventory can answer.

For the philosophical engagement with classical traditions, the change is one of posture. The traditions retain their commitments — their scriptural authorities, their soteriological projects, their metaphysical articulations, their internal disputes. The book's structural argument can be in dialogue with them at the structural level without taking on their commitments and without demanding that they take on the book's. This is the modest dialogue Chapter 9 articulated with respect to Advaita; a similar posture may be appropriate for later source-grounded engagements with Yogācāra, apophatic theology, Neoplatonism, Mādhyamika, Sufi metaphysics, Daoist traditions, or other possible antecedents. The book has not undertaken these further engagements. The posture for doing so is set out in Chapter 9, and the work remains.

The argument also hands work back to first philosophy itself. The structural question the book has been pursuing is not new in philosophy — it is at least as old as the *Adhyāsabhāṣya* and as recent as the contemporary phenomenological and analytic literatures that have approached it from different directions. What the book has offered is one disciplined contemporary articulation, in one available philosophical idiom, with attention to the specific contemporary opponents who have most articulately pressed against it. Other articulations are available. The book has not claimed that its idiom is the only available one or that its engagements exhaust the field. The articulation it has offered is one contribution among many possible contributions to a question that long predates the book and will outlast it.

C.5 Final statement

The book has not added a new item to the inventory. It has tried to show why the inventory was never the right place to look.

The thesis defended is structural, not metaphysical in the first-order sense. The argument has been procedural, not constructive. The conclusion is jurisdictional, not triumphal. The book has been a careful working-through of one structurally simple observation, against the principal contemporary attempts to resist it: awareness is not one item within the inventory of what appears, but the field within which any such inventory becomes intelligible.

What this finally establishes is modest. No constructive ontology, however careful and however well-motivated by internal filters, answers the field-question; the field-question is not a question within ontology to which a new item, a new property, or a new grounding relation is the answer. The book is not a metaphysics of the One, not a doctrine of cosmic mind, not a refutation of physicalism in the standard sense, not a verificationist semantics, not a rival ontology to constitutive cosmopsychism, not a contemporary form of Advaita. It is a structural argument for a jurisdictional limit, with an acknowledgement that the limit has been articulated before, in other idioms, and with the contribution offered as methodological rather than first-doctrinal.

The argument has reached its limit. What follows is the work of others — empirical, phenomenological, ontological, comparative, contemplative — for which the structural result is at most a propaedeutic. The book makes no claim to more than this.

Bibliography

Allais, Lucy. *Manifest Reality: Kant's Idealism and his Realism*. Oxford: Oxford University Press, 2015.

Block, Ned. "On a Confusion about a Function of Consciousness." *Behavioral and Brain Sciences* 18, no. 2 (1995): 227–247.

Carruthers, Peter. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press, 2000.

Chalmers, David J. "The Meta-Problem of Consciousness." *Journal of Consciousness Studies* 25, nos. 9–10 (2018): 6–61.

Frankish, Keith. "Illusionism as a Theory of Consciousness." *Journal of Consciousness Studies* 23, nos. 11–12 (2016): 11–39.

Goff, Philip. *Consciousness and Fundamental Reality*. New York: Oxford University Press, 2017.

Gomes, Anil. "Perception and Reflection." *Philosophical Perspectives* 31, no. 1 (2017): 131–152.

Guyer, Paul, and Allen W. Wood. "Introduction." In Immanuel Kant, *Critique of Pure Reason*, 1–44. Translated and edited by Paul Guyer and Allen W. Wood. Cambridge: Cambridge University Press, 1998.

Husserl, Edmund. *Cartesian Meditations: An Introduction to Phenomenology*. Translated by Dorion Cairns. The Hague: Martinus Nijhoff, 1960.

———. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, First Book: General Introduction to a Pure Phenomenology*. Translated by F. Kersten. The Hague: Martinus Nijhoff, 1982.

Kant, Immanuel. *Critique of Pure Reason*. Translated and edited by Paul Guyer and Allen W. Wood. Cambridge: Cambridge University Press, 1998. Cited by A/B pagination throughout.

Śaṅkara. *Brahma-Sūtra-Bhāṣya*. Translated by Swāmī Gambhīrānanda. Kolkata: Advaita Ashrama, 1965. Tenth impression, 2009.

———. *Upadeśasāhasrī: A Thousand Teachings*. Translated by Sengaku Mayeda. Albany: State University of New York Press, 1992. [Reprint of University of Tokyo Press, 1979.]

Sartre, Jean-Paul. *Being and Nothingness: An Essay on Phenomenological Ontology*. Translated by Hazel E. Barnes. New York: Philosophical Library, 1956. 2011 Open Road Integrated Media edition.

———. *The Transcendence of the Ego: A Sketch for a Phenomenological Description*. Translated by Andrew Brown. London and New York: Routledge, 2004.

Searle, John R. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press, 1992.

Stern, Robert. *Transcendental Arguments and Scepticism: Answering the Question of Justification*. Oxford: Clarendon Press, 2000.

Strawson, P. F. *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason*. London: Methuen, 1966.

Stroud, Barry. "Transcendental Arguments." *The Journal of Philosophy* 65, no. 9 (1968): 241–256.

———. *The Significance of Philosophical Scepticism*. Oxford: Clarendon Press, 1984.

Zahavi, Dan. *Subjectivity and Selfhood: Investigating the First-Person Perspective*. Cambridge, MA: MIT Press, 2005.