

Very much an informal work, continually in progress. Comments, corrections, feedback, and suggestions are more than welcome. Please send to [brian.cantwell.smith@gmail.com](mailto:brian.cantwell.smith@gmail.com)

## A Rough and Ready Guide

*to the tortuous landscape of metaphysics, ontology, and epistemology*<sup>†</sup>

Brian Cantwell Smith<sup>†</sup>  
University of Toronto  
September 29, 2014

“What do you see out the window?”, the master asked the novitiate. “A mountain,” said the novitiate. “Go study,” said the master.

“What do you see out the window?”, asked the master, ten years later. “That which, in virtue of my upbringing, situation, and cultural embedding, I take to be that which I call *mountain*,” said the new monk. “Good,” said the master. “Keep studying.”

“What do you see out the window?”, asked the master when he was very old, many decades later. “A mountain,” said the monk—and smiled.

### Introduction

This document is intended as an informal reference and guide to some of the most common concepts and vocabulary in metaphysics, ontology, and epistemology (MOE). It is not a list of *definitions* (there is nothing to make definitions out of, at this fundamental level). It merely systematizes, in a rough-and-ready way, intuitions that the reader is assumed already to have. Its aims are modest: (i) to organise commonsense into a sufficiently workable system to sup-

### Different world views

A thorny issue faced by any introductory document of this type is to accommodate different metaphysical views on allocation of responsibility for “how things are”—specifically, on whether the nature of the world is determined: (i) by the external world, independent of us; (ii) by us, singly or collectively, independent of the (rest of the) world; or (iii) by us and the world together, in inexorable and inextricable combination. Some people (roughly: naïve realists) believe that the world *is a certain way*, independent of how people or cultures understand it. Others—notably: constructivists, idealists, and other types of *irrealist*—believe that responsibility not only for *how we take the world to be*, but even for *how the world is*, lies in part, or even wholly, in the minds and practices of observers and communities.

It is not the aim of this document to resolve—or even to take a stand on—this spectrum of views. Instead, Part III presents what is intended to serve as relatively neutral language, roughly compatible with, and thus moderately shareable among, a wide variety of them. ...

### Different world views (cont'd)

First, though, setting metaphysical diversity largely aside, Parts I and II lay out the vaguely realist intuitive picture of language, mind, and world that underlies commonsense discourse (a view sometimes called the “natural ontological attitude”).

Thus, as will be perfectly evident to any reader, Parts I and ii are framed in terms of an intuitive subject-object divide. Critics of naïve realism may feel that there is no such divide—and that any deliberations framed in its terms are at best fraught, and at worst fatally flawed. But in my view that reaction is mistaken. First, nothing in Parts I or II assumes that an ontological world independent of subjects or observers. Second, some kind of distinction must be made between what we think (what is in our heads) and what we think about, on pain of brute intelligibility. If one thinks that someone is tall, or says that a painting is red, or notes that it is raining, then whatever is the case about that which is tall, that which is red, or the rain, it cannot in any simple way be the mental idea or image, a pure concept, or anything else wholly cognitive. The question for the irrealist, therefore, is one of accounting for what it is that is being thought (or claimed) tall, red, or wet—or, which may provide room for more radical alternatives, of accounting for what is going on when someone thinks or says such things. Simply to deny the existence of that which is claimed tall, red, etc., without making any other suitable distinctions between a thought and that which it thinks of—i.e., simply to endorse an undifferentiated collapse of mind and world—is a way of ducking, not addressing, the substantive issue.

The most important step in developing MOE expertise, in my view, is to recognize that doing justice to alternative metaphysical world views, including radically non-realist ones, is not to eschew, or to attempt to escape, everyday language. That way lies vapid futility. Rather, the challenge is to understand how ordinary language is used in mundane discourse—what it means, what warrants its proper or worthwhile exercise, and how it should be analysed. It would miss the point, to say nothing of being shallow, to view the debate between realism and antirealism as being about *what (if anything) is real*. A much deeper issue is at stake: *what it is to be real*. This is the question that supports genuine, profound disagreement—and whose study can lead to enlightenment.

port elementary conceptual analysis; and (ii) to mesh sufficiently with the terminology of logic, philosophy, AI, knowledge representation, and other technical disciplines so as to enable the reader to gain some understanding of those discipline’s literatures, proposals, and systems.

A number of challenges face any presentation of this material.

1. **World views:** As described in more detail in the sidebar (previous page and above), it is a struggle to accommodate the inevitable variety of metaphysical views that readers will bring to the subject. Part III of this document addresses such issues explicitly; earlier parts largely set them aside, but in ways that are intended to make room for a wide divergence of underlying views.
2. **Terminology:** Almost all of the terms that are used to talk about ontological, epistemological, and metaphysical issues (including the terms *ontology*, *epistemology*, and *metaphysics* themselves) are used in substantially divergent ways across the intellectual landscape, according to the local customs and predilections of different disciplines, schools, traditions, etc.

The word ‘concept,’ for example, is used non-psychologically in many parts of philosophy, but very much psychologically in other places, including some areas of psychology. One possible interpretation of the situation is to assume that the word has the same “meaning” or “referent” (see Part II), and thus to conclude, when philosopher  $\alpha$  strenuously argues that concepts are not in the head, and psychologist  $\beta$  equally passionately claims that they are, that the two are disagreeing. But it is equally possible that they are talking about different things. More likely yet, the truth of the matter may involve a little of both. And there is nothing unique about the word ‘concept.’ All sorts of MOE terms share a similar fate—of being used in identifiably different but potentially overlapping ways in different disciplines, including *meaning, content, semantics, representation, interpretation, reference, function, term, variable, denotation, correctness,* and myriad others.

These two issues interact in deviously intricate ways, so as to engender all sorts of confusion in a subject matter than in any case would be far from simple to understand. So take the following as *a*—not *the*—story about the MOE landscape. One of the by-now most widely agreed facts about classification and sense-making is that any proposed set of categories in terms of which a subject matter is rendered intelligible will reflect the contingencies, biases, and perspectives of its author. This is certainly true in the present case.

An overarching three-way conceptual division will make this fact bluntly apparent. The document is structured in three parts, according to a relatively straightforward but still not entirely standard reading of the three major terms:

- Part I · **Ontology** — What there is
- Part II · **Epistemology** — How we think and talk about what there is
- Part III · **Metaphysics** — Underlying (or overarching?) world views.

## I • **Ontology** – The Structure of the World

While ‘ontology’ is usually technically characterised as the “philosophical study of the nature of being, existence and/or reality,” in more informal discourse it is used as a label for the various *categories of being*—which is to say, for the categorical structure of the things in the world around us (what is sometimes called the *furniture of the world*). That is, and especially for purposes of representation and classification, ontology is basically (the study of<sup>1</sup>) *what there is*.

The words used to characterise “what there is” are multiple. In fact it is not too much to say that “characterising what there is” is one (if not the) most important functions of language—to describe the world, thereby conveying how it is, or at least how we take it to be (or would like it to be, or are afraid it is, or whatever). This is especially true of simple concrete names, verbs, adjectives, and adverbs. When we call a tree a tree, or ask someone whether they’ve had dinner, or remark on some mind-bollixing prestidigitation, we use the words *tree, dinner, prestidigitation,* and so forth, to, as it is variously said:

---

<sup>1</sup>The ‘ology’ suffix, from the Greek λόγος (‘word’), signifies “the study of.” Thus *musicology* is the study of music; *pharmacology*, the study of the uses and effects of drugs; etc. Strictly speaking, ‘ontology’ should be used for the *study* of what there is, not for what there is simpliciter. Since we lack a simple term for the latter, however, the term ‘ontology’ is often pressed into service.

1. *Name, identify, classify, categorize, sort, or catalogue* an object or phenomenon or event in some way...
3. In the process attributing to said object or phenomenon or event (or claiming that it exhibits or exemplifies or manifests) one or other *property, quality, characteristic, or facet*...
4. In virtue of which it can thereby be assigned to (or be claimed to be a member of) one or other *class or category or group*, and...
5. As a result of the whole process thereby *organize, sort, systematize, catalogue, render intelligible, make sense of*, and so on, the world.

Here, for simplicity, as a way to get started, I will talk primarily about *properties*, without thereby assuming anything special about them—specifically, whether they are essential, intrinsic, well-behaved, socially constructed, vague, political, highly contested, subjective, or anything of the sort. That is: the word ‘property’ will be used as a kind of “ur-concept” —simply to get at the inexorable fact, as old as human history, that we classify or identify things, assign them to named groups, etc. Similarly, I will use the term ‘object’ as neutrally as possible, as effectively a synonym for *thing*<sup>2</sup>—i.e., for anything that can be thought about or referred to or identified, anything that could or might exist. The idea is to develop or construct more complex, nuanced, restricted, and otherwise particular meanings and phenomena on top of this basic notion.

Strong constructivists, irrealists, and some others may object to the apparent subject-object divide in the this picture, which to them may seem pernicious. Three comments. Note, first, as should soon be evident, that no commitment is being made to any of the classical objectivist, absolute, independent-of-human perception characteristics that are normally assumed to be the cause of the demerits of adopting this mundane world view. Second, as mentioned in the sidebar (p. 1), all that is being expressed, in employing this “lightweight” subject/object divide, is what I take to be the natural ontological attitude that underwrites everyday discourse. Third, to make this concrete, imagine saying of something that it is entirely flat, or a thousand miles long. What is being called “ontological” in this section are merely those things that are flat, or a thousand miles long, *whatever they are*. Even if one imagines that the properties “flat” and “a thousand miles long” are mental constructs, it is not the mind, per se, that is being claimed to be flat, or a thousand miles long, on *anyone’s* metaphysics. For present purposes, therefore, it is enough simply to take ontology to be (the study of) the sorts of things that can be flat, long, etc., whatever their metaphysical status.

## I • Basics

- A. **Objects** (equivalently: **entities**) are the basic *things* that populate the world being talked about. Objects are typically assumed to be **countable**, **reidentifiable**, and (at least conceptually) **distinct**.
  1. Physical objects are sometimes assumed to necessarily be spatio-temporally continuous. But in many situations that is an unhelpfully constraining principle. A better (though ultimately circular!) characterization is to describe them as *common causes of multiple effects*, and *common recipients of the effects of multiple causes*.
  2. Objects are paradigmatically thought to be the “locus classicus” of **identity**, in the sense

---

<sup>2</sup>Pace distinctions between the two notions, on Heideggerian, Latourian, and many other lines.

that there is assumed to be a fact of the matter, for any two objects  $\alpha$  and  $\beta$ , as to whether  $\alpha$  and  $\beta$  are the same object or not. Concrete objects, more specifically, are sometimes characterised (e.g., in Strawson's *Individuals*) as “re-identifiable particulars.” Re-identifiable means that it is (necessarily) possible for one and the same object to be identified, from or at a different time and place, and/or from a different perspective, as one and the same (numerically identical) object it was earlier.

When used in this way, and indeed in most logico-philosophical contexts, the term ‘identity’ refers to something like the property, possessed by each and every object, of *being the object that it is*.<sup>3</sup> As such, the notion is different from, and vastly simpler than, the much more complex, psychologically rich notion of identity intended in such sentences as *He doesn't have a robust sense of identity* or *The identity of the aboriginal community has been decimated by the spread of late neo-capitalism*.

3. Some objects are **abstract**—i.e., don't “occur”, don't weigh anything, don't take up any space, and “exist,” if existence can be predicated of them,<sup>4</sup> only in something like Plato's heaven or other equally diaphanous realm. Example of abstract objects would include: *numbers*, the *concept of fairness*, and *types and classes* (as opposed to the items “in” or “of” those types or classes; see §I.D, below). Other objects are **concrete**. But the distinction isn't particularly clear. One can argue whether many familiar things are abstract or concrete—such as *ideas*, *fame*, *détente*, and *arrangements* of blocks (as opposed to the blocks themselves, which are presumably concrete). Though massless, all these things occur, have histories—and will vanish when the universe collapses.

Even if not “cast in stone,” contextually-appropriate versions of the abstract/concrete distinction are often useful in particular situations.

- B. **Properties** are characteristics or facets **exemplified** (equivalently: **held** or **manifested**) by objects. One sometimes also talks of properties being **attributed to** or **predicated of** objects.

1. Classically, some properties were taken to be **intrinsic** to objects, in the sense of being essential to the object's being the object that it is.<sup>5</sup> Others, called **extrinsic**, were viewed as inessential, or contingent. Thus the property of being male, or at least of being human, might be viewed as intrinsic to Ivor, in contrast to having short hair, standing in front of the refrigerator, or living in Madagascar, all of which would be viewed as extrinsic, because they could all change, without his thereby becoming a different person (“he would still be him”). Relational properties were always taken to be extrinsic.

The idea of intrinsic properties commands far less allegiance than it once did, under the influence of various forms of antiessentialism and poststructuralist critique. Nothing in the rest of this document presumes the intelligibility, let alone the validity, of an intrinsic/extrinsic distinction—and indeed it is a not a distinction I myself believe in.

2. Similarly, nothing is assumed, in virtue of an object's being assigned a property, or being **identified** as having a property or of belonging to a category of class (below), that

<sup>3</sup>Cf. Duns Scotus' notion of *haecceity*.

<sup>4</sup>There is a major tradition that argues that existence is not a predicate; but that is for later.

<sup>5</sup>The idea that properties are intrinsic underlies the etymology: classically, a *property* was something understood to be something that was “proper” to an object.

there are any *other* properties that it must necessarily thereby manifest or exhibit.<sup>6</sup>

If an item is identified as a *table*, for example, then on the present account we say that it exhibits the property ‘table’, or is ascribed the property ‘table’. Nothing is presumed about what *warrants* that ascription. In particular no claim is made that there are (one or more) *other properties* that it must have or exhibit, in order to legitimately possess the ‘table’ property, or (equivalently) to be a legitimate member of the class *table*. (Equivalently: nothing is presumed about whether anything at all can be *defined*.)

3. All objects manifest an infinite number of properties (and are thus members of an infinite number of classes). Thus a stopwatch may exemplify the properties of *being a gift*, *weighing 127 grams*, *costing \$119.37*, *not working*, etc. (as well as *weighing less than 128 grams*, *weighing less than 129 grams*, etc. ad infinitum—to say nothing of exemplifying the property *being a stopwatch*).<sup>7</sup>
  4. Properties (like types and classes, below) are abstract and “multiply instantiated,” in the sense that if two different pools of water both exemplify the property of being H<sub>2</sub>O, it is the one and the same (numerically identical) property that they both exemplify. Similarly, if two patches of colour are chartreuse, then there is only one colour that they both are.<sup>8</sup>
- C. **Relations** are like properties, but involve two (or more) objects. Thus the relation *being 7 feet apart* might be exemplified by Pat and Hillary. As well as saying that two (or more) objects *exemplify* a relation, one speaks of them as **standing in** that relation.
1. In general it matters which **role** in the relation each object plays. Thus if the relation *being the father of* holds between Llewellyn and Ewan on one assignment of roles (e.g., with Llewellyn assigned to the ‘father’ role, and Ewan to the role of ‘child’), then it will not hold (claims that it holds will be false) if the assignment to roles is reversed.
  2. A **relational property** is a property made by “filling in” all but one of role in a relation. Thus *living halfway between San Francisco and Palo Alto* is a relational property made out of the 3-role relation *living-halfway-between*(\_,\_,\_), with two of the roles filled in.
- D. **Types** (equivalently: **kinds**; sometimes also called **sortals**) are properties used to sort objects into identifiable groups—which groups are then often called **classes** or **categories**.<sup>9</sup>
1. For example, the type *chair* is basically the *property of being a chair*, used to group together those things that are chairs into the class *chair*, and to separate them from those

---

<sup>6</sup>I.e., no assumption is made that properties, classes, types, categories, etc., are either monothetic or polythetic.

<sup>7</sup>No difference is made here between the property *being a stopwatch* and the property *stopwatch*. Sometimes it is convenient to refer to it one way; sometimes, the other.

<sup>8</sup>To the extent that there is metaphysical sense to talking, in a more fine-grained way, about the object-specific concrete redness of this patch in particular, as opposed to the object-specific concrete redness of that one, it can be approximated in the present picture by distinguishing the state of affairs  $\sigma_1$  of patch<sub>1</sub>’s being chartreuse from the different state of affairs  $\sigma_2$  of patch<sub>2</sub>’s being red. (In philosophy such object-specific “instances” of property exemplification are sometimes called *tropes*, in a wholly different sense term from its ordinary use to refer to a conventional or idiomatic use of a phrase.)

<sup>9</sup>Some people use ‘class’ for an informal grouping, such as ‘chair’ or ‘bread’, ‘category’ or ‘friend’, for a grouping that results from a deliberate and more organized process of establishing categories, such as those that result from biological processes of constructing taxonomies. In line with the overall aim of staying at as simple and fundamental a level as possible, no distinction is made here between these notions.

things that are not chairs—including not only (most) tables, rugs, lamps, milkshakes, etc., but also the colour green, perestroika, and the number seven.

2. In the general case, types are said to have **instances**, and (equivalently) classes and categories are said to have **members**. Equivalently, if an object exemplifies a property used as a type or category or class, then the object is said to be **of** the type or category or class, or to be **in** or to **belong to** the category or class.
3. In some cases more specific words are used. When the type is linguistic or intentional (Part II), for example, one talks of **tokens**, especially when they are written, and of **uses** or **utterances**, when they are mental or oral. Thus whereas it would be natural to refer to an “instance” of the type *chair*, it would be more common to speak of a “token” (rather than instance) of the name ‘Jehoshaphat,’ and to a “use” or “utterance” (rather than an instance) of the sentence “Time to split!”

## II • Composites

### A. States of affairs (SOAs)

1. **States of affairs** (or **situations**) are essentially *ways that chunks of the world are (or might be)*.
2. Using states of affairs as technical entities may initially seem odd—because it is not a “cut” on the world that is normally referred to as such<sup>10</sup> in informal language. But “reifying” (i.e., treating as objects or entities) states of affairs is hugely useful, theoretically, not least because they are often the best way to understand what it is that a representation is representing.
3. There are various ways of classifying states of affairs. One obvious one:
  - a. **Simple**: An object exemplifying a property; or two or more objects standing in a relation;
  - b. **Complex**: Two or more simple states affairs connected by logical operators (Part I §II.C).
4. An important (and orthogonal) way of classifying SOAs has to do with what is called their **modal** status:
  - a. **Actual** (equivalently: the SOA **obtains**)—in which case any sentence expressing the state of affairs is **true** (Part II §II.G.2.b);
  - b. **Possible**—not actual, but could be, if things had been arranged differently, or depending on how this come out in the future (i.e., their being actual wouldn’t violate logic, or any other fundamental laws);
  - c. **Impossible**—they violate laws of logic, reason, or perhaps physics (e.g., a situation in which one equals two, or backwards causation is possible [affecting the past], etc.).

### B. Combinations and collections

1. Things come not only singly but in various combinations and collections.
2. **Sets**, **groups**, and **sequences** are common collective forms, all of which are said to have **elements** or **members** (written ‘ $\in$ ’, as in ‘ $x \in \alpha$ ’ to mean that  $x$  is a member of the set  $\alpha$ ).
3. When properties are attributed to collectives, the relation between the property, the

---

<sup>10</sup>Registered; see Part III.

collective as a whole, and the members of the collective, can vary widely. For example, the property may have to do with:

- a. The whole as a unity (an *important set*);
  - b. The collectivity of the members, rather than anything about them individually (a *large group*);
  - c. A relation between or among the elements (a *diverse set*; a *fractious committee*);
  - d. A property of the elements (“the *white set* belongs in that box”)
4. **Wholes** commonly have **parts**. The study of the relations among parts and wholes is called **mereology**. Mereological issues permeate many real-world classifications.

### C. Logic

1. Most so-called **logical properties** are best understood as epistemological (Part II §II.G), but it is common, depending on one’s metaphysical perspective, to take some to be ontological as well. Thus many people assume that the following “logical operators” combine, ontologically, with properties and relations, yielding logically composite versions:<sup>11</sup>
  - a. **Negation** (‘not’, also written ‘ $\neg$ ’)—such as the property of *not being a teacher*;
  - b. **Conjunction** (‘and’ or ‘ $\wedge$ ’)—such as the property of *being tall and handsome*; and
  - c. **Disjunction** (‘or’ or ‘ $\vee$ ’)—such as the relational property of *hiding in either Chicago or Dubuque*.
2. Logically composite relations and states of affairs would also be licensed, on such a view (such as *not being Trevor’s barber*, and *Frank and Bobbie’s not having seen each other for a year*).

### III • Other issues

A. **Formality**: As repeatedly stated, nothing in this characterisation requires that properties (and classes) be well-behaved, “objective,” hierarchical, unique, exclusive, orderly, lawful—or exhibit any other property of being “well-behaved.” For some purposes and in some contexts, however, strict constraints on properties and their associated classes are stipulated, assumed, required, thought convenient, etc. (For example, there are those, not including the current author, who hold that these properties are required in order for an organization to be a *categorization*, as opposed to a mere *classification*.)

We will therefore call a property, relation, collection, class, or other ontological category **formal** just in case the following interrelated facts are true:

1. For any object, there is a definite (*yes/no*) fact of the matter as to whether it does or does not exemplify the property in question, or does or does not belong to the class—i.e. without ambiguity.
2. Exemplification of the property (or membership in the class) is uniform or homogeneous, in the sense that no objects are better exemplars of the property than others,

---

<sup>11</sup>There is considerable debate about the metaphysical “reasonableness” of such logical composites. Conjunctive versions seem the most obviously OK; negative properties are perhaps next in line. Ontological disjunction (it actually being the case, ontologically, that a person can exemplify the disjunctive property of either living in Chicago or living in New York) seems more problematic than the seemingly more reasonable idea that they *either* exemplify the (non-composite) property of living in Chicago *or* exemplify the (similarly non-composite) property of living in New York. But opinions vary.



more central or paradigmatic, exemplify the property to a greater or lesser extent, etc.—i.e., without matter of degree.

3. There is a well-defined boundary to the property or class, implying that if  $\phi$  is a formal property or class, then so is “not  $\phi$ ” (written ‘ $\neg\phi$ ’ in logic).

On this definition *being friendly* and *bread* are thus pretty clearly not formal properties, whereas the mathematical property of *being a rational number* evidently is.

- B. **Reification:** Any entity (including the above, including properties, relations, SOAs, etc.) can be **reified**, by being treated as an object—and thus can exemplifying further properties, can standing in relations, etc. (This document, in particular, attempts to reify just about everything relevant to the MOE landscape.)

## II • Epistemology — Language, Representation, and Mind

Consciousness, being towards death, the achievement of “no-self”—mysteries of mind remain among the most challenging for contemporary intellectual inquiry. Here we will focus on a much more mundane fact about human cognition, but one still surpassingly impressive: our ability to talk, think, imagine, hope, and reason *about the world around us*—i.e., to engage in what philosophers call **intentional** behaviour.<sup>12</sup>

For present purposes it is most helpful to consider thought, language, and other forms of human intentional behaviour under the general rubric of **representation**.<sup>13</sup>

## IV • Representation

- A. Three facts about representation are most important:

1. **Directedness:** As a species of intentionality, representations are **directed towards** objects, states of affairs, and other phenomena in the world—typically, towards objects, states of affairs, and phenomena external to those in which the systems themselves play a causal or constitutive role.
2. **Aspect:** Representation is what is sometimes called **aspectual**: representations do not simply represent parts or chunks of the world *in and of themselves*, but *in a certain way*.
  - a. Thus a photograph cannot represent a table without representing it *as standing*, as

---

<sup>12</sup>Three homophonous words need to be distinguished. (i) As a word of lay English, the word ‘intentional’ means something like deliberate or done on purpose—i.e., something arising from an act of will or *intention* on the part of the actor. (ii) As indicated, the same term (or at least a word of the same spelling) is used technically in philosophy to denote those phenomena that are characterised by what in the text is called the “arrow of directedness,” that involve issues of *meaning, interpretation, representation, semantics*, etc. —i.e., any phenomena characterized by what is informally called “aboutness.” The two primary categories of intentional phenomena in this sense (i.e., in the sense being characterised in the text) are *mental* and *linguistic*; but a spectacular number of more specific familiar intentional types can be identified, as suggested in Figure 1 on the next page. (iii) Finally, as suggested in Part II §II.E, the term ‘intensional’ (i.e., a version spelled with ‘s’) is used in logic, linguistics and philosophy of language, in contrast with ‘extensional,’ to indicate contexts in which nouns phrases, primarily, are *used*, in which case it is the sense or meaning of the phrase that contributes to the meaning of the whole sentence, rather than the reference or interpretation.

The three senses are related; but an explanation of how goes beyond the scope of this document.

<sup>13</sup>There are those who think of all representation talk as intrinsically committed to a particularly strong (and in their view pernicious) form of naïve realism and pure objectivism. This is not true; but see the over-all introduction, and Part III.

1. **General semantic (intentional) phenomena**
  - a. Signs, symbols, names, pictures, icons, images, schemata...
  - b. Scores, codes, blueprints...
  - c. Words, phrases, sentences, expressions, utterances...
  - d. Statements, assertions, questions, stories...
  - e. Representations, information...
  - f. Descriptions, depiction, specifications, illustration...
  - g. Analysis, simulation, account, theory, explanation, model...
  - h. Pointing, showing, indicating, directing, interpreting...
  - i. Languages, communication, discourse, conversation...
  - j. Proofs, arguments, demonstrations, rebuttals...
2. **Cognitive correlates**
  - a. Thought, knowledge, recognition, realisation, deliberation...
  - b. Belief, desire, concepts, ideas, decisions, reasons...
  - c. Action, behavior, perception, sensation, interaction...
3. **Computational versions**
  - a. Programming languages, specification languages...
  - b. Symbol manipulation, information processing, data...
  - c. Data bases, knowledge representation, meta-level systems...
  - d. Programs, instructions, codes, memory, data structures...
  - e. Pointers, identifiers, names, variables...
  - f. Call by value, call by name, referential transparency...
  - g. Syntax, semantics, interpretation, models...
  - h. Correctness...
4. **Theoretical notions used to analyze semantic phenomena**
  - a. Syntax, semantics, meaning, content, significance...
  - b. Reference, interpretation, designation, denotation...
  - c. Truth, falsity, certainty, likelihood, intension, extension...
  - d. Inference, consequence, entailment...
  - e. Induction, deduction, abduction...
  - f. Logic, consistency, soundness, completeness, validity...
  - g. ... etc.

Figure 1 — The realm of the intentional

- round* or *as square*, etc. Similarly, a diagram of a triangle cannot avoid taking a stand on whether it is acute, right, or oblique.
- b. Even the phrase “a table” represents its referent [below] *as a table*, rather than *as a theatrical prop*, or *as an item of indoor furniture*, etc.
  - c. It is because of the ineliminably aspectual nature of representation that it is essential to understand the (ontological) world in terms of *properties, relations, and states of affairs*, not merely in terms of individual and *sets* (and other collections) of *objects*.
3. **Function:** From the point of view of any given creature, system or community, the

world is in general wildly diverse and widely distributed, with most of it (especially at any given instant) causally inaccessible. Against this background, it is helpful to see representational systems, including language and thought, as methods of organizing and presenting, in a **local, constrained, effective** way, a range of typically **distal** and/or **non-effective** represented phenomena or states of affairs or subject matters of interest.<sup>14</sup>

- B. It follows from the foregoing that representations is governed by a fundamental dialectic.<sup>15</sup>
1. **Physically** or **mechanically**, representations must work in virtue of what is *immediately causally available*: the capacities deriving from their concrete physical embodiment (how they are built), and resources to which they are immediately connected.
  2. **Normatively**, they are accountable to the semantic relations they bear to what they represent.
  3. **Dialectic**: What representational systems *should* do, in other words, is defined in terms of relations in which they stand to situations or states of affairs that are in general remote, non-effective, or for some other reason cannot play a direct causal role. What they *do* do stems from their physical constitution. They are thus inexorably riven by a particularly potent version of an *is/ought* dialectic.
  4. Because of this dialectic, representations cannot in general be guaranteed to be correct or to “do the right thing.”
  5. Representational designers and users—including everyone who thinks and talks!—do their best to ensure that the system *exploits whatever local, effective properties are available* (including the states of our brains and our immediate perceptions) so that they can “stand in for” or “serve in place of” *the remote or distal or other non-effective situations to which they are accountable*.
- C. A huge variety of linguistic, cognitive, artistic, and other textual and documentary forms are representational. This document focuses primarily on:

---

<sup>14</sup>The term ‘effective,’ plus or minus, applies to properties and states of affairs that are *causally potent*—that can do mechanical work. ‘Non-effective’ properties and states of affairs cannot do such work—at least not in virtue of being those properties or being those states of affairs. Three notes. (i) Whether something is effective is thus a perspectival fact; a magnetic field may be effective within an appropriate range (falling off at  $1/r^2$ ); rotational motion is likely to be effective only for devices physically connected to it. (ii) Some properties, such as *having been purchased at Indigo*, are not effective, but may be correlated with other properties (such as squiggles of ink on a small card that constitute an Indigo price tag) may be both effective (will reflect incident light in a determinate pattern) and correlated with having been bought at Indigo. Thus a mechanical agent may be able to determine whether the item was purchased at Indigo by responding to the effective property (appropriately configured squiggles of ink) that correlates with the property in which it is interested. This sort of behaviour is the very essence of representation. (iii) Virtually all properties of interest in ordinary human affairs are, in fact, in-effective (such as a person’s identity). We are so used to playing this game of correlating properties that are effective, that we don’t intrinsically care about, with properties we do care about, which aren’t effective, that we barely notice how much it permeates our every action and reaction.

If the phenomena of interest are close by and effectively accessible, it is sometimes easier to use them directly (e.g., the distance to a wall, or the status of the power cord).

<sup>15</sup>Although rarely understood in these terms, formal logic is best understood as an in-depth analysis, within a mathematical context, of the nature of this dialectic.

1. **Linguistic:** sentences, words, paragraphs, conversations, etc.; and
2. **Cognitive:** thoughts, ideas, mental images, and so on.

A set of unanswered questions arise, in the foregoing characterisations, regarding the relation between the two phenomena of *representation* and *classification*. Classification was presented in Part I *ontologically* (the identification of properties used to create types or classes of objects, where the types and classes are taken to be abstract). Salient examples of classification are all epistemic, and thus *representational*, activities: thinking, speaking, organizing supporting intentional structures (catalogues, data bases, etc.). Would we need to classify, if everything in the world were causally accessible to us, at all times, and therefore we didn't need to represent? Are there any non-representational uses of classification?

## V • Language

- A. Human language is taken to be the prime vehicle for expressing or describing the world. There is therefore a (preternatural?) correspondence between the (epistemic) structure of language and what was claimed in Part I to be the (ontological) structure of the world. That said, it is misleading to assume any strict “type-type” correspondence between: (i) type (or kind) of *representation* (noun, verb, sentence, etc.) and (ii) type (or kind) of *entity represented* (object, property, state of affairs, etc.) But the following can be taken as an initial guide:
1. Objects  $\Leftrightarrow$  singular terms (i.e., count noun, or noun phrase that serves the role of a count noun)
  2. Properties and relations  $\Leftrightarrow$  adjectives, verbs, and common (non-count) nouns
  3. States of affairs  $\Leftrightarrow$  clauses, gerundial phrases (*your arriving late*, *Secretariat's winning the Belmont Stakes*), and nominative absolutes<sup>16</sup>
- B. **Speech acts:** There is more to a sentence or thought than what it represents. Uttering sentences and thinking thoughts are *actions*, which *do* various things—called, in the linguistic case, **speech acts**. **Indicative** sentences **assert** (or **deny**) **claims** (i.e. assert that represented SOAs are actual); **interrogative** utterances **ask questions**; **imperative** utterances **command**, etc. There are also implied actions (literally, “is the salt next to you?” asks a question, but what is called its illocutionary effect is to convey a request).
- C. **Types:** What we normally take to be “a word”, “a sentence”, etc., is a *type* with multiple tokens or instances (or, as one often says in the spoken case, **utterances**). It is only because the (single) name ‘Harriet’ supports an arbitrary number of (distinct) utterances that we say that it is “her name.”
- D. **Compositionality and Productivity**
1. Language (and, it is claimed, thought) is said to be **compositional**—in the sense that the *meaning of the whole* is systematically constructed from the *meaning of its parts*. This is how we can construct and understand an indefinite number of entirely new sentences that we have never before encountered (e.g., “My aardvark’s favourite grapefruit was damaged in Florence,” which has surely never before been written, said, or thought).

---

<sup>16</sup>Having the same form as gerundial terms, but used as absolutes, such as ‘the day’s being sunny’ in the sentence “The day’s being sunny, we took the Ferrari.”

2. The compositionality of language is the property that allows it to be **productive**: to be a finite system that nevertheless allows for the generation and production of an indefinitely large number of sentences, expressing an indefinitely large number of meanings, most of which have never before been uttered or encountered, but that can nevertheless be comprehended by finite speakers and hearers.
- E. Internal contexts within a (composite) expression are called **transparent** (or **extensional**) just in case the reference or denotation (or truth, in the case of sentences) depends only on the reference or denotation of the item that fills the context; **opaque** (or **intensional**) otherwise.
1. Thus the underlined phrase in “The first winner of the Fields medal was Danish” occurs in a transparent context, since “Ahlfors was Danish” has the same truth-value (it is false; Ahlfors was a Finn). However the same underlined phrase occurs in an opaque context in “Amiya wants the first winner of the Fields medal to come for dinner,” since this sentence could be true even if “Amiya wants Ahlfors to come to dinner” is false (perhaps because Amiya knows Ahlfors and doesn’t like him, but also doesn’t realize that he won the Fields medal).
- F. A related (but different) distinction involves the question of whether a noun phrase refers at all, or merely identifies a property. “John wants to marry a Norwegian” could be true either of two senses:
1. **Referentially**: There is a person, who is Norwegian, whom John wants to marry, whether or not John knows that this person he is so fond of is Norwegian; and
  2. **Attributively**: John would like whomever he marries to be Norwegian, though as yet he has no particular person in mind.
- G. **Logic**
1. As well as being understood ontologically, the ordinary **logical connectives** can be analysed as representational or sentential operators, which affect the meaning (and reference) of expressions in the obvious way:
    - a. **Negation** (‘not’ or ‘ $\neg$ ’)
    - b. **Conjunction** (‘and’ or ‘ $\wedge$ ’)
    - c. **Disjunction** (‘or’ or ‘ $\vee$ ’)
  2. There are many other logical operators as well, which receive more complex analysis. Among the most common
    - a. **Implication** (‘ $\supset$ ’, ‘ $\rightarrow$ ’, or ‘ $\Rightarrow$ ’): to say that one situation implies or entails another.<sup>17</sup>

---

<sup>17</sup>In point of fact, implication (or derivability) and entailment are different—and in a fundamentally important way. To put it bluntly, implication and derivability are essentially *epistemic* notions, having to do with language and thought—i.e., with relations between and among *representational expressions*. Entailment, in contrast, is fundamentally *semantic*, and thus ultimately rests on *ontological* relations between and among *states of affairs*.

To say that situation  $\phi$  *implies*  $\psi$  (written ‘ $\phi \rightarrow \psi$ ’ or ‘ $\phi \supset \psi$ ’, or sometimes ‘ $\phi \vdash \psi$ ’, though the latter connotes a different logical framing) is to say that, given statement or representation  $\phi$ , it is possible to derive or prove  $\psi$ . That is, implication is a *representational* (or *epistemological*) notion. To say that  $\phi$  *entails*  $\psi$ , on the other hand (written ‘ $\phi \Rightarrow \psi$ ’, or ‘ $\phi \models \psi$ ’, again in a different logical framing), is to make a semantic or ontological claim: that in situations or states of the world in which  $\phi$  obtains,  $\psi$  will obtain as well.

One way to understand the difference between implication (derivability) and entailment is by analogy with the notions of prediction (epistemic) and determinism (ontological). Saying that  $\phi$  *implies*  $\psi$  is some-

- b. **Equality** ('='), viewed as a verb designating a relation between two entities (objects). The state of affairs designated by the sentence is actual, and hence the sentence is true, just in case the two entities are the same (i.e., are in fact one entity).
- 3. There are also a number of **quantifiers**, which affect the meaning of a sentence in complex ways. The three most common:
  - a. **For all** (' $\forall$ '): used in sentences to say, e.g., that "all birds fly" (' $\forall x \text{ bird}(x) \rightarrow \text{fly}(x)$ ').
  - b. **There exists** (' $\exists$ '): to say, e.g., that "there is (at least one) person who is 7' tall: ' $\exists x [\text{person}(x) \wedge [\text{height}(x) > 7 \text{ feet}]]$ '
  - c. **Numbers**: Ordinary numbers are treated as quantifiers in logic.

The foregoing discussions assume that the only ontological structures in the world that can be thought about or described in language are those that fit into the object-property-relation-collection (OPRC) framework outlined in Part I. In fact, however, the expressive potential of thought and language are almost surely much wider—even if the OPRC conception is hugely common. Some simple examples, such as "It is raining," show that even when the structure of language suggests a commitment to an OPRC view, no such commitment is actually being made (since there is no referent to the use of 'it' in that phrase, no referent or object of which it can be said that *that object* is raining). Similarly, if one recalls standing at the beech and feeling a wind flow over one's face for minutes on end, it is not clear that the way that the world is being remembered involves the exemplification of any properties by any (especially distinct) objects. And so on.

The content and behaviour of such examples are studied within philosophy under the label **non-conceptual content**, under a conception of 'conceptual content' that takes it to have, necessarily, an OPRC structure. Discussions of the meaning and semantics of non-conceptual thought and language outstrip the mandate of this guide, but the existence of such inquiries provides an avenue to those who find the cartography of this MOE landscape excessively limiting.

## VI • Meaning, Interpretation, and Semantics

Intentional systems in general, and representations in particular, are characterised in terms of such familiar but demanding notions as *meaning*, *interpretation*, *content*, *reference*, *semantics*, and *truth*. None of these words will be new to anyone reading this text; but the concepts they name are among the most difficult and contested in the entire MOE pantheon.

Fundamentally, these notions—and a variety of other theoretical entities allied with them—have been developed in order for us, as theorists, to *characterise* the relations that connect: (i) the concrete mental act of thinking a thought or uttering a sentence, or a material list, or a specific classificatory typology, on the one hand; and (ii) the objects or entities or states of affairs towards which those acts or representational structures are directed (II.A.1, above), on the other—i.e., as we say, the objects or entities or states of affairs that are *being represented* or *discussed* or *referred to*. In coming to understand meaning and its ilk, the most important preparatory step for theoretical analysis is to recognize that, with just the two concrete "ends" of

---

what analogous to saying that, given knowledge of  $\phi$ , one could predict or figure out  $\psi$ . Saying that  $\phi$  *entails*  $\psi$ , in contrast, is analogous to saying that  $\phi$  *determines*  $\psi$ —whether or not, from knowing  $\phi$ , one is able to figure out  $\psi$ .

such an intentional “**arrow of directedness**,” we as theorists do not have the conceptual wherewithal to make sense of the regularities that govern this ubiquitous and mundane behaviour. Ordinary speakers and thinkers can perfectly well speak and think without trafficking in these notions *per se* (in exactly the same way that a child can perfectly well use nouns and verbs without, at least in any explicit way, possessing the concepts of ‘noun’ and ‘verb’). But to analyse even relatively simple thinking and discourse requires the deployment of this, as it is called, “semantical” machinery.

- A. **Semantics**, most generally, is (the study of) the full suite of relations connecting a representational entity or system or phenomenon, on the one hand, with that towards which it is directed. For the moment, I will treat all of the following phrases as equivalent ways of identifying that “object” of representational or intentional directedness:
1. That which it is **about**;
  2. That towards which it is **directed**;
  3. That which it **denotes** (in which case the object is called the **denotation**);
  4. That to which it **refers** (in which case the object is called the **referent**).
- B. The properties *human* and *featherless biped* may apply to the same things (people), but the phrases ‘human’ and ‘featherless biped’ are taken to **mean** different things—since someone could use one, fully competently, without knowing that they are referring to the same things as the other. Similarly (and famously), ‘the morning star’ and ‘the evening star’ have different meanings, but the same **referent** (which we also refer to with the proper name ‘Venus’). Similarly, the police could be looking for “the people who stole Munch’s *The Scream*,” without knowing that they were looking for Bjørn Hoen and Petter Tharaldsen.
- C. The best way to understand the term ‘**meaning**’ is as designating that which is in common (shared) between and among those who use the same or equivalent phrases—where by *same* or *equivalent* is intended something roughly like this: that they are or would be interchangeable, without noticeable differences, for most (all?) epistemic purposes, such as predicting people’s behaviour, causing the same responses or results, “understanding” the cognitive state of the person who has made the claim, etc.
- D. In the normal case meaning is a finer-grained (more finely-individuated) phenomenon than the objects or situations in the world to which it refers (though see §VI.E, below). But the issue of how to individuate (identify) meanings is a hugely contested. Various criteria have been suggested for claiming that two expressions *mean* the same thing, including:
1. *They necessarily imply each other* (according to which ‘il pleut’ and ‘it is raining’ mean the same thing, but this criterion also implies that ‘ $1+1=2$ ’ and ‘ $2+2=4$ ’ mean the same thing, which is counterintuitive, and violates suggestion #3);
  2. *They cannot be differentially believed* (i.e., one cannot rationally believe one and not the other—so ‘is the son of Mark’ and ‘has Mark as father’ would mean the same thing);
  3. *They designate the same states of affairs* (same objects and properties).
  4. *They are written in the same way* (according to which “il pleut” and “it is raining” would not be counted as having the same meaning).

What is most notable about all of these suggestions is that they defer questions of **identity** to other cases—which *thoughts* cannot be differentially believed; when are two *properties* in fact one and the same property (such as *being human* and *being a featherless biped*), when

two *inscriptions* can be counted as the same, etc. The issues are fraught, however, and disputed. Caveat emptor.

- E. The foregoing considerations all apply to so-called **context-independent** expressions, which refer to or designate the same thing in all situations of use.
1. Whether any natural language expressions are truly context-independent is not clear ('37' denotes the number thirty-seven if taken as a decimal numeral, and thirty-one if taken as octal; there are thousands if not millions of people with the (same?) name 'John', etc.).
  2. A particularly important form of **context-dependent** expression are those variously called **indexical** or **deictic**—*here, today, I, you, over there*, etc. In spite of the manifest different *referents* on different occasions of use, however, it is usually said that different utterances of the word 'I', for example, have the same *meaning*. (You know what someone "means" when they use the term, even if you've never met them—because, it is said, you *know what the word means*. But the *referent* of course differs, depending on the speaker.)

### III • Metaphysics

A minister moves into town, and wants to come to know his community. "That's a wonderful garden that you and God have there," he says to a neighbour, leaning over his backyard fence. Straightaway the neighbour replied, "It sure is—but you should have seen it when God had it alone!"<sup>18</sup>

As stated in the introduction, people differ, sometimes quite radically, on what we can call the "division of labour" for determining how the world is—between and among at least: (i) individual consciousnesses, (ii) cultures, communities, and local histories; and (iii) "reality out there," beyond or otherwise independent of how it is that humans take it. It is characteristic of what is sometimes called 'naïve realism' to place all the responsibility on the third item in the list—to assume that the world *is as it is*, and perhaps also that science is determined to describe this, independent of all facts about human activity or consciousness. (In informal parlance this view is often simply called 'realism,' but it is too strong a version to warrant that label.)

As a way of framing the issue in order so as to allow the deepest and more interesting metaphysical positions to be described, and also in aid of avoiding unhelpful disagreements, this section introduces a family of related terms (*register, registration, etc.*) in terms of which to describe how it is that we take the world. The aim is: (i) to incorporate what is sensible and ordinary in common (mundane) discourse; (ii) to avoid making either inadvertent realist or irrealist commitments when not helpful or necessary; (iii) to serve a workably shareable language that can be used by adherents of different metaphysical views; and (iv) notwithstanding the foregoing, to accommodate what in point of fact I myself take to be the "right answer" to the overall issue.<sup>19</sup> In particular, the view is designed to do justice to what I take to be two fundamental **humilities**:

<sup>18</sup>I included this story in a footnote in *On the Origin of Objects*, MIT Press, 1996; p. 300, fn. 14.

<sup>19</sup>*Ibid.*



1. The **humility underlying realism**: that there is more to the world than we, singly or in community, are able to grasp or understand—an ultimate “more” to which we must always defer; and
2. The **humility underlying constructivism** and related views (especially as adumbrated in feminist philosophy and so-called “science and technology studies” or STS): that we humans, in our thoughts and languages and practices and projects and interventions, inevitably and inexorably shape and alter and adjust the realities in which we live, in ways ranging from the specific, locally contingent, and historically particular to broad biases and perspectives and traditions of culture, history, and community.

In contrast to Parts I and II, which aim to present concepts and terminology roughly in line with the mainstream of contemporary academic practice, the views expressed in this section are in effect my own proposal. I have therefore written it in the first person.

Overall, I say that we **register** the world in terms of the objects, properties, situations, states of affairs, features, etc., that we thereby take it (the world, that is) to consist in. Thus: *I register a building across the street, she registers a lake on the horizon, they registered their university as a petty and hapless institution mired in scholastic bureaucracy*, etc. Some comments:

1. The term ‘registration’ is intended to index the aspectuality of registration (Part II §I.A.2)—i.e., the fact that all representational systems, including human thought, perception and understanding, as well as language, inexorably take the world *as being a certain way*, rather than representing it, as it were, simpliciter.
2. The term is intended to be neutral as to any distinction between or among *sense, perception, thought, judgment*, etc.
3. To the extent that the registration is systematic, or made against a background of concepts that form a world view, one can say that the **registrar** has registered in the world in terms of a **registration scheme**.
3. Unlike ‘conceive’ or ‘cognize,’ but like ‘see’ and ‘perceive,’ the term ‘register’ is intended to be a “**success**” verb. Thus if, in ordinary circumstances, a person registers—i.e., takes there to be—a tree, then it is fair to assume (i.e., it is a legitimate default assumption) that there *was a tree there to be so taken*, and that the person *did so take it*, in whatever is the applicable full semantically and normatively appropriate way.<sup>20</sup> (In contrast, to say that someone “conceives” a large house, there is almost the reverse implication: not that they are thinking about a house that exists, but rather than they have imagined one that does not.)
4. As should be evident from the above (including uses of the qualifier ‘successful’), registration is taken to be **normatively laden**, in the philosophical sense of serving as the subject of such issues as truth, objectivity, worth, etc.
5. By following the verb’s direct object with ‘as...’, the construction facilitates at least a first step towards distinguishing how we, as theoreticians (philosophers, cognitive scientists,

---

<sup>20</sup>This is a rather realist characterisation of “success”; it should be replaced as appropriate for other metaphysical views. The point is simply that “ $\alpha$  registered  $\beta$ ” should be considered “true”, or otherwise worthy or worthwhile, just in case something roughly of the form “There is  $\beta$  and  $\alpha$  took it to be  $\beta$ ” is similarly true, worthy, or worthwhile.

epistemologists, knowledge representers, etc.), register situations or phenomena in the world, and *how we take them to be registered by other subjects or people or agents, of whom we may be speaking*. Thus we might say, of an infant, “She registers her mother’s coming to the door not as the re-appearance of a recognized individual object, but more as *re- or repeating placement*, in Strawson’s sense, of the feature *Mama*.”<sup>21</sup>

6. If not used with an explicit ‘as’ construction—i.e., notwithstanding (5)—it can be assumed that the direct object of the verb ‘register’ includes *the aspectual nature of the way in which that phenomenon or entity is registered by the (individual designated by) the sentence’s subject*. In this way ‘register’ differs from at least common uses of such perceptual verbs as ‘see.’ Thus while some would claim that it is possible for the sentence “Randy saw the Northern Lights” to be true even if Randy did not recognize them *as* the Northern Lights, it is considered to be a default implication of the sentence “Randy registered the Northern Lights” that Randy did so take them (providing there is no following ‘as’ clause, and providing as well that nothing is said to counter the implication).<sup>22</sup>
7. The term ‘register’ is explicitly intended to be usefully neutral on the division of responsibility between person and world for the resulting ontological “take”—i.e., is neutral as between realism (I successfully register a table *as* a table because it *is* a table<sup>23</sup>), strong forms of constructivism (I successfully register it as a table because of the contingent and historical forces constituting the social community of which I am a member, or even due to the particular exigencies of my own individual history), idealism, solipsism, and many other epistemic, ontological, and metaphysical proposals.



<sup>21</sup>It is only a first step, because of the evident but fraught issue of *how we register the subject’s registration* (e.g., in the ‘ $\gamma$ ’ part of the sentence ‘ $\alpha$  registers  $\beta$  as  $\gamma$ ’)—including whether we can, and if so how much, and in what respects.

<sup>22</sup>I.e., the direct object position of the verb ‘register’ is thus not assumed to be referentially transparent.

<sup>23</sup>Note that this statement “because it *is* a table” does not require that the applicable form of realism be naïve.