Chapter 1

Introduction

Physicists are currently searching for what they call a "Theory of Everything." However, it turns out that the "everything" they have in mind falls far short of every thing. The physicists' theory of everything has nothing to say about mental phenomena, agency, values, norms, teleology, intentionality, to mention but a few. In fact, physicists rarely have much to say about the natures of the fundamental elements of their theory: particles, fields, space and time.

None of this is surprising, and none of it is a criticism of current physics as such. When physicists refer to a coming "theory of everything", they do so (or, at least, the sophisticated ones do so) with tongue in cheek. It is metaphysics, and not physics, whose province it is to fashion a theory of everything. This book is a work in real, honest-to-God, no-apologies-given metaphysics, but metaphysics conducted in a thoroughly scientific spirit. My hope is that it will help to stimulate a return to the perennial concerns of philosophy.

1.1 A Comprehensive Realism

A class of propositions can be interpreted *realistically* when two conditions are met:

1. Some of the propositions are evaluated as true or false.

2. The truth or falsity of the propositions in the class is determined by some set of facts, and this set of facts plays an indispensable role in explaining our knowledge of the truth or falsity of the propositions in the class.

The first condition is not sufficient, since the truth-values of the propositions could be determined by facts about our collective acts of affirmation or projection, in which case the propositions could not be interpreted realistically. The causal element introduced by the second condition is critical, because it specifies a direction of asymmetric dependence: our knowledge depends causally on the fact establishing the truth or falsity of the corresponding propositions. This entails that the facts determining these truth conditions do not include facts about our attitude toward those very propositions, since causal dependency cannot be circular.

I will argue that propositions involving reference to the following things can and should be interpreted realistically:

- Natural properties and relations.
- Situation and event tokens.
- Modality and objective probability.
- Causal connections.
- Numbers.
- Proper functions (teleofunctions).
- Mental states.
- Secondary qualities.
- Enduring substances.
- Values and norms.

Although my position is one of a comprehensive realism, I give a relatively simple and unified picture of the world. The first three items on the list above are treated as primitives, but all of the others are explicated in terms of these more fundamental entities, properties and relations. Everything that is posited to exist is posited to exist because of some role it plays in the causal network of the world. My approach is resolutely non-dualistic: I reject any sort of Cartesian or neo-Cartesian postulation of a scientifically inaccessible realm of subjectivity.

At the same time, I do not start with any apriori or dogmatic requirement. My aim has not been to build a theory of the mind that is materialistic or physicalistic or naturalistic. To begin one's metaphysical inquiry with such dogmatic commitments is methodologically irresponsible. We must simply follow the evidence where it leads. If it leads to materialism,

well and good, but if it leads away from it (as my own account does in several respects), we must be willing to be accountable to the facts, not to philosophical fashion.

Theories of content, meaning and representation in terms of causal connection have become very prevalent. A number of philosophers have taken causal theories of content as reason to be anti-realist about values (Mackie, Harman), numbers (Field), and minds (the Churchlands). In my view, the burden of such anti-realism is too great for a theory of content to bear. However, if a causal theory of content could be devised which vindicated realism about values, numbers and minds, such a theory would give us the best of both: a plausible, informative and simple account of content, and the accommodation of much of our commonsense view of the world. In this book, I will try to develop such a theory.

1.2 Metaphysical Method

This book is unapologetically a work of substantive metaphysical theory. Fortunately, blind anti-metaphysical prejudice is not as common as it once was. Nonetheless, many may legitimately ask for the ground rules of the enterprise. In a recent book on causation, Daniel Hausman[72] proposed five criteria for evaluating metaphysical theories:

1. Intuitive fit.
2. Empirical adequacy, consistency with what we know about the world, including our best scientific knowledge.
3. Epistemic access. The theory should include some account of how we could come to know its truth.
4. Superseding competitors. The theory should incorporate the successes of its predecessors.
5. Metaphysical fecundity. The theory should shed light on a variety of metaphysical issues.

The only criterion that I would add to the list is that of simplicity or elegance. A good metaphysical theory should not be in need of ad hoc rescues or endless epicyclic tinkering.

The principal motivation of my work is that of *unification*. I aim to provide a unified account of intentionality and knowledge, one in which we give exactly the same kind of account both for our thought about and knowledge of objects and events in space and time, and for our thought

about and knowledge of the facts of logic, mathematics, laws of nature, and objective chance. We should not accept a bifurcated, disjunctive account of thought and of knowledge so long as a unified account is possible. The theoretical cost of postulating genuine modal facts (as I do) is small in comparison to the benefits of unification.

1.3 A Third Way: Causalism as an Alternative to Both Physicalism and Mysterianism

Since causal relations play the fundamental role in my metaphysics, the term "causalism" might be an appropriate term for my approach. In recent years, others have taken what could be described as an essentially causalist approach to the metaphysics of mind, namely Armstrong, Millikan, Dretske, Papineau and Lycan. A causalist theory of mind identifies intentionality with a certain kind of causal property (perhaps involving higher-order causal connections), and the peculiar qualities of conscious experience are taken to be explicable in terms of their intentionality. In all of these cases, causalism is seen as a strategy for defending materialism against various objections concerning intentionality and consciousness. The opponents of these approaches, including Searle and McGinn, have been labeled the "mysterians", since they hold that we can expect to find no informative account of the nature of intrinsic intentionality or consciousness.

Unfortunately, those participating in these controversies have overlooked the fact that causalism is separable from a commitment to physicalism or materialism. A non-physicalist causalism would include an informative account of the nature of mental states without insisting that everything can ultimately be explained in terms of atoms and the void. I will argue that all of the extant objections to causalist theories of mind are in reality objections to the *conjunction* of causalism with physicalism. A non-physicalist causalism provides the resources for an adequate answer to these objections. In addition, I will argue that there are independent grounds, having nothing to do with the philosophy of mind, for rejecting physicalism.

1.4 Causal Internalism and Vertical Causation

The notion of causality is absolutely central to recent philosophical work in semantics, the philosophy of mind and intentionality, epistemology, and philosophy of science. Work by Donnellan, Kripke [100], and Putnam [141] helped to make causal connections an indispensable part of our accounts of reference and signification. This in turn has generated causal theories of information and content ([44] and [60]). The Gettier problem led to the renaissance of causal theories of knowledge by Goldman [66], Armstrong [5] Pollock [138], and Plantinga [137]. Causality is put to much work in recent theories of personal identity and of the nature of mental states (as in the functionalism of Lewis [109] and Putnam [141]). Causation continues to figure prominently in philosophy of science (e.g., Wesley Salmon's causal theory of evidence [152]) and in theoretical science, both within physics and outside.

Additionally, causal reasoning plays a central role in both understanding and predicting events. Recent work in AI has brought causal reasoning into renewed prominence. For example, the much discussed Yale Shooting Problem reveals (according to most diagnoses, see especially [135]) the absolute necessity of recording and using information about the causal links between the bits of information we have about the world.

Attempts to explain away causation or to replace it with some purely statistical regularity (whether or not supplemented by some kind of psychologistic decoration) have proved to be catastrophic failures. Every attempt to explain causal direction (surely one of the most fundamental features of causality) in terms of the nomological-deductive model has failed. Such models of causality have generated paradoxes far more rapidly than ad hoc solutions can be invented for them.

If a robust sense of reality leads us to recognize causal connections as first-class citizens of our ontological inventory, we must also make room for those special kinds of objects that can serve as relata for causal relations, whether we call these objects possible 'facts', 'situations', or 'states of affairs'. These objects must be distinguished from propositions and from quasi-linguistic representations if we are to capture accurately the logical relations governing causal idioms. The restoration of such fact-like entities to respectability has also been a common theme of recent work in philosophy, including philosophical linguistics and the Stanford situation theory of Barwise and Perry.[10]

The project of building a unified theory of intentionality and knowledge in causal (or teleo-causal) terms faces a major obstacle: accounting for our

knowledge of *modal* facts, i.e., facts about necessity and possibility (including logical and mathematical modality), about counterfactual conditionals, about objective chance or propensity (as a generalization of objective modality), and about physical or natural necessity as embodied in *natural laws*. This obstacle is a generalization of the problem Paul Benacerraf [14] [15] has raised in the case of mathematics: how is definite reference to and substantive knowledge about mathematical objects possible, given that our best theories of reference and knowledge involve causal connections between our thoughts and their targeted aspects of reality? Benacerraf's problem generalizes to our thought about the laws of nature, about the objective chances of certain kinds of events in certain situations, and about various kinds of possibility and necessity. In each case, we seem to have intentional reference and knowledge to things that the philosophical tradition has long considered to be causally inert.

Overcoming this obstacle calls for a revolutionary rethinking of our standard picture of causation. This standard picture I call the *horizontal* or *externalist* model of causation. The alternative I am proposing is the thesis of *causal internalism*, which countenances the reality of *vertical causation*.

On the standard, horizontal model, causes and effects are, exclusively, physical, spatio-temporally local states and occurrences. The causal nexus, whether it consists in a kind of necessary, stochastic, or nomic connection, stands outside of both the cause and the effect. This is why I call it *causal externalism*: the causal nexus is wholly external to both the cause and the effect. The horizontal/externalist model can account for our knowledge of occurrent properties realized in spatiotemporal locations, but it leaves the entire realm of modality causally, and, therefore, cognitively and epistemically, inaccessible.

My alternative proposal is that we consider the modal (or nomic or stochastic) facts that tie the cause to the effect to be *internal* to the cause or to the effect. Depending on the details of one's account of causation, causes necessitate or probabilify or possibilify their effects. On an internalist model, the fact that a given cause necessitates its effect is itself an integral part of the total cause, not something that stands outside or above the cause-effect pair. Consequently, modal facts are every bit as causally efficacious as are occurrent physical facts, and so there is no barrier to providing a unified, causal theory of all of human thought and knowledge. For instance, we can think about and gain knowledge of natural laws by virtue of the fact that each of these laws enter into some, but not all, causal connections. When we observe a regularity (like the elliptical orbits of the planets) that is really caused by a particular nomic fact (like the law of gravitation), then our observations provide us with intentional and epistemic contact with that nomic fact.

Here are some of the more significant claims that I make in Part I concerning the nature of causation:

1. The causal nexus is not something above and outside the cause and effect but consists of facts wholly internal to the cause and the effect. This thesis of causal internalism commits me to the existence of *vertical causation* from modal and nomic facts to ordinary spatiotemporal ones, crucial to giving a unified, causal account of intentionality and knowledge.

2. Modal facts exist, including facts of logical and mathematical necessity, and these facts are not reducible to or supervenient on the occurrent facts of the world (including its merely actual regularities). The existence of logical types (negations, conjunctions, disjunctions, etc.) of arbitrary complexity is a substantive fact about the world.

3. There are compelling reasons for rejecting a strong version of determinism, reasons that are independent of the problem of free will (chapters 4 and 5).

4. Only actual situations exist, but in constructing models for modal logic, it is convenient to introduce the fiction of merely possible and even impossible situations.

5. I propose a new solution to the problem of the scope or extent of causation, namely, that every wholly contingent state has a cause. On the basis of this principle, I demonstrate the existence of a necessary First Cause (chapter 8).

6. It is possible to give a principled basis for a defeasible or nonmonotonic logic that incorporates causal information. This logical calculus (developed in appendix B) generates rich and plausible conclusions about probable consequences of known or hypothesized states.

My theory of causation is designed to provide an exact, mathematical model that satisfies the following aims:

1. Causal connections and order should be defined without reference to space and time, permitting the construction of a non-circular, causal theory of spacetime.

2. It should permit the possibility of higher-order or *vertical* causal connections, in order to explain logical and mathematical knowledge, mind/body interaction, and the nature of teleofunctions.

3. It should provide natural explanations of the formal properties of causation and causal explanation, including transitivity, asymmetry, and veridicality.

4. It should match the data provided by intuitions about the validity and invalidity of various forms of causal reasoning. In particular, it should explain the failure of substitution of classical equivalents in causal contexts (see Chapter 3), and our default assumption of the universality of causal explanation (Chapter 8).

5. It should be able to navigate successfully through the complexities of the relationship between causality on the one hand, and modal and statistical relations on the other. It should not treat causation as a primitive, with no intrinsic relationship to correlation or necessity, but it must avoid the paradoxes that have resulted from attempts to reduce causality to statistical relations.

6. It should be compatible with indeterminism, and with merely probabilistic connections between cause and effect (Chapters 5 and 6).

7. It should provide an account of the modularity (or locality) of causal reasoning: the role (recently much investigated by researchers in the field of artificial intelligence) of causation in enabling us to draw correct default conclusions in the presence of irrelevant information (Appendix B).

The last desideratum is especially important, since any theory of causation that does not account for the special virtues of causal reasoning is seriously incomplete. Researchers in logic and artificial intelligence, such as Judea Pearl[135], have discovered that reference to causal relations plays an indispensable role in our commonsense reasoning about the world. The Yale Shooting Problem of McDermott and Doyle (which I discuss in Appendix B) is an excellent example of the sort of problem of reasoning about prospective change that requires a causally-informed description of the situation. I argue that the fundamental characteristic of causality that explains its importance in commonsense reasoning is the Markov property: when one fact is causally screened off from a second by one of its causes, then the conditional probability of the second on the cause is independent of the first fact. This justifies our exclusion of causally irrelevant information (information that is causally screened off from our prospective conclusions by our premises) in reasoning defeasibly.

1.5 The Ontology of Causation

In order to make sense of causal relations, we must be able to apply the part-of relation (and the associated machinery of mereology) to the causal relata. This means that we must acknowledge the reality of concrete existences, tokens, that can play the role of concrete events and states (or "situations"). In addition to these situation-tokens, we will need abstract, repeatable situation-types. The situation-types represent intrinsic qualities or characters of situation-tokens. This choice of primitives is drawn from the work of Barwise, Perry and Etchemendy (Stanford situation theory).

The situation-tokens can serve as the truth-makers for propositions, playing the role that "facts" play in the philosophies of Austin, Bergmann, and Hochberg. When it is true that the cat is on the mat, there is a concrete cat-on-the-mat situation-token s that makes it true. This token s is of the cat-on-the-mat type.

Complex situation-types can be constructed from simpler ones by means of logical operators, such as negation and disjunction. These operators should be interpreted by means of the strong Kleene three-valued truth tables or the four-valued Dunn tables (as explained in Appendix A).

In addition to tokens and types, there is a causal priority relation $\prec$, a strict partial ordering (transitive, irreflexive and asymmetric) of situation-tokens. If $s \prec s'$, then s is qualified to act as part of a cause of s'. Intuitively, we can think of $s \prec s'$ as meaning that s is wholly in the backward time cone of s'.

In Chapters 5 and 8, I advocate the thesis that all of the causal antecedents of a token are essential to its identity: if any of them had failed to exist, the token itself could not have existed. If we accept this thesis, then we can define the causal priority relation in this way: $s \prec s'$ if and only if s and s' do not overlap mereologically (that is, they have no part in common), and no part of s' could exist unless s existed.

There are two notions of causation that I define: (1) total causation (s is a total cause of s') and (2) INUS causation. INUS causation refers to J. L. Mackie's account of a cause as an insufficient but necessary part of an unnecessary but sufficient condition for the effect [114]. Both total causation and INUS causation introduce a modal or statistical element: a total cause must makes its effect conditionally necessary, or at least, conditionally much more probable than it would otherwise be. An INUS cause is an indispensable part of some total cause: s is an INUS cause of s' just in case there is a total cause s'' of s', s is a part of s'', and any part of s'' that does not contain s as a part is no longer a total cause of s'.

1.6 The Need for an Indeterministic Model

In Chapter 4, I develop a deterministic model of causation, one in which a total cause necessitates its effect. However, I discover a number of independent reasons for being dissatisfied with such a model:

1. We have clear intuitions that causation should be possible in an indeterministic world.

2. If causes necessitate their effects, and effects necessitate their causes (since the identities of their causes are essential to their own identities), then causes and effects would be modally inseparable.

3. When applied to specific examples, the necessitation model overgenerates causal connections, and inflates the minimal content of causal explanations.

There are several difficulties that pose serious problems for building an indeterministic model of causation, however. First of all, verifying the transitivity of causation is no longer trivial, once we abandon strict necessitation as the standard. Verifying the veridicality of causation is also non-trivial. In addition, mere probabilistic relevance is neither necessary nor sufficient, as is demonstrated by two kinds of cases: (1) causes with no or even with negative statistical relevance to their effects, and (2) preempted causes, pre-conditions with positive statistical relevance that are nonetheless not causes, because some independent factor preempts their operation. Finally, there is the Markovian independence principle that I mentioned above, which is critical to explaining the modularity of causal reasoning, but which is also difficult to secure in an indeterministic setting. In Chapter 6, I use Lewis/Stalnaker conditionals in a novel way to overcome these difficulties.

1.7 A Causal-Probabilistic Theory of Information

My teleological account of mental representation depends crucially on being able to define information without reference to mentality or teleofunctionality. In order to do this, I borrow heavily from the work of Fred Dretske[44], in which information is defined by means of objective probabilities. According to Dretske, a fact p carries the information q just in case the conditional probability of q on p is equal to 1, which Dretske interprets as meaning that p necessitates q.

The principal difficulty with such an account is that of accounting for the possibility of error or misinformation. If p carries the information that q, then it is impossible for p to be true and q false. There are two popular solutions to this difficulty, neither of which is really satisfactory. We could require only that the conditional probability of q on p be within some small, finite interval of 1, or we could require only that the conditional probability of q on p be higher than that of q on $\neg p$. However, if we do either of these, we lose the validity of the Xerox principle, the principle that information is transitive: if p carries the information that q, and q carries the information that r,then p carries the information that r.

A second popular strategy (adopted by Dretske himself), is to add some condition N, representing normal or canonical training conditions, and require that the conditional probability of q on the conjunction $p \,\&\, N$ be equal to 1. These normal conditions are usually specified retrospectively, be reference to some salient, historical facts. In Chapter 9, I argue that these retrospective strategies are inadequate, and I propose two alternative solutions, one using infinitesimal probabilities, and the other conditional functions.

A token s carries the information that p *robustly* in world w just in case every part s' of w that contains s as a part carries the information that p. This means that s carries the information that p, and every extension of s in w also carries this information. Robust information is the pre-cognitive analogue of knowledge. When one knows something on the basis of robust information, one is immune to Gettier-like counterexamples.

1.8 Why an Exact Theory?

A formal or exact theory is an attempt to use logic and mathematics to represent a conception (or family of conceptions) of a particular subject matter. For example, Newtonian mathematics involved the use of the calculus to represent a conception of the physics of motion.

Formal or exact metaphysics should not be thought of as the analysis of concepts, or as a branch of pure logic. Nor should it be identified with the articulation of our commonsense worldview (the conception of the world ensconced in ordinary language and everyday practice), although metaphysics typically begins with this task.

An exact theory of a metaphysical subject, such as causation, is an attempt to express our best, most educated guesses about the truth of the matter in a form that is as falsifiable and corrigible as possible. The alternative to developing an exact theory is operating with an undisciplined miscellany of hunches and intuitions, poorly defined and changing unsys-

tematically as one moves from one sphere of application to another. Without an exact theory, inconsistency is very difficult to detect. Unanticipated consequences are rarely discovered, and one's reasoning is often afflicted with non sequiturs and unintended equivocation.

The task of defining and investigating an adequate formal language for representing causal reasoning remains unfinished. Recent work by Pearl and Verma ([135][136]), and by Spirtes, Glymour, and Spirtes ([166]), is suggestive but limited, in that all this work takes the relation of causation to hold among a fixed enumeration of *dynamic variables*. However, in ordinary causal reasoning, we often take complex facts and events to be causal factors. In Part I, I define a formal language for causal reasoning that is capable of treating facts of arbitrary complexity as causes and effects, and of resolving many of the outstanding logical puzzles.

I am confident that the theory of causation that I develop in Part I is clear and precise enough to be falsifiable. Where it goes wrong (as I'm sure in many places it does), it should be possible to construct clear counterexamples, either from real life, or from imagination, accompanied by strong intuitions of real possibility.

The subject of causation has experienced a renaissance in analytic philosophy over the last generation. Theories and arguments involving causation proliferate, in epistemology, philosophy of mind, philosophy of science, and philosophy of language. However, few working in these areas have attempted systematic and exact accounts of causation, and no such an account, to my knowledge, is directly relevant to as broad a range of outstanding philosophical problems as is the account presented here.

1.9 The Big Picture: Preview of Part II

In this book, I develop a theory of causation, and I apply this theory to a large number of outstanding problems in philosophy, including such topics as:

- The definition of proper function (teleology)
- The semantics of mental representations
- The mind/body problem (including free will)
- The causal basis for logical and mathematical knowledge and cognition
- The problem of induction (including Goodman's puzzle)

- Enduring substances and their identity-conditions

- The construction of space and time

- The objectivity of values and moral norms

Obviously, I cannot do justice to the vast literature on any one of these topics. However, in each of these topics, the concept of causation plays a central role, and I cannot claim to have developed an adequate theory of causation without at least beginning the task of testing my theory against the data provided by each of these problem areas. For this reason, I have been forced to cast my net very broadly.

I do not pretend to have said anything dispositive on any of these subjects in this book, but I do believe that the novel account of causation that I develop here enables me to make a genuinely original contribution in each case, one that I hope will stimulate further discussion. In each case, confusion about the nature and conditions of causation have produced an impasse. The introduction of an exact account of causation, together with the development of some novel proposals, may help to move the discussion to more fruitful ground.

The overall structure of the project goes something like this. The theory of causation and information (developed in Part I) is used to construct a theory of teleofunctionality as a form of higher-order causation (Chapter 12), and an account of the causal efficacy of logical and mathematical facts (Chapter 15). After a survey of recent accounts of mental representation, I combine my theories of information and teleology, resulting in an account of the semantics of mental representations (Chapter 14): a mental representation carries the content p just in case it has the teleofunction of carrying the information that p. The theory of mental representation is then used in developing theories of mind/body interaction, qualia and free will (Chapter 16), and knowledge and induction (Chapter 17) . I develop a causal/teleological theory of enduring substances and their identities through time in Chapter 18. Both the theory of teleology and that of mental representations are used in the development of a eudaemonistic theory of ethics (Chapter 19), which in turn is used in sketching an account of moral realism (Chapter 20).

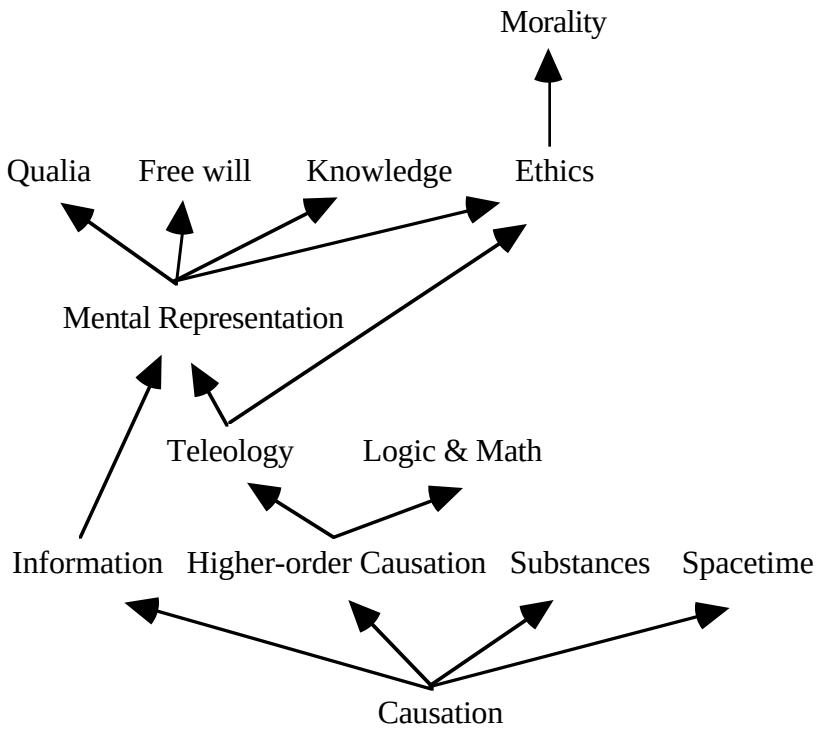

Here are some of the more significant claims that I make in Part II:

1. There is a tight connection between the semantics of belief and epistemology: once we have the semantics right, the theory of knowledge is merely a corollary (Chapters 14 and 17).
2. There are powerful reasons for rejecting materialism (which I take to include, at a bare minimum, the limitation of causal relations to spatiotemporal items), reasons that are independent of the well-known problems in the philosophy of mind (see Chapter 21 for a summary of these reasons).
3. A simple, causal theory of mathematical thought and knowledge is possible, one that unifies the theory of mathematical knowledge with that of empirical and scientific knowledge (Chapter 15).
4. Taking functions seriously leads to a very robust form of ethical realism, one that does not identify objectivity with some sort of idealized

subjectivity but instead revives the eudaemonism of Plato and Aristotle (Chapters 19 and 20).

5. The use of the mereology of events and of non-classical (three- and four-valued) interpretations leads to more sophisticated conceptions of supervenience, type identity, and token identity than were available heretofore. These more sophisticated conceptions enable us to solve the problem of mental causation (Chapter 16).

My aim in this book is to bring an end to the dualism that has dogged philosophy since the downfall of Aristotle's metaphysics (including his "metaphysical biology") at the beginning of the modern era. Commentators such as Leo Strauss, Alisdair McIntyre and John McDowell have all located the roots of the dualisms of mind and body, of fact and value, and of objectivity and subjectivity, in that early modern separation of scientific fact and normativity. In my view, the early modern turn away from Aristotle has been both unnecessary and disastrous. Aristotle's "metaphysical biology" is more viable in light of modern knowledge than it has ever been, and the recognition of this fact can bring about a great re-unification of our view of the world.

At the same time, I will argue staunchly against a false re-unification, built upon a narrow physicalism. Physicalists have been right to insist that our knowledge of the real cannot extend beyond the network of causation. They were right, therefore, to challenge the viability of positing a subjective and normative realm beyond the reach of science. However, they were wrong to think that science teaches us that only physical states, states located within the framework of space and time, can be causally efficacious. In fact, science provides abundant evidence, albeit implicitly, of the causal efficacy of physical, mathematical and logical modality.

There is no need to read the chapters of this book in strictly sequential order. In fact, I expect few readers to be interested in all of the topics covered. For example, if you have little interest in logic or in formal theories of causal reasoning, you can skip Appendices A and B altogether, without doing damage to your comprehension of the rest of the book. If you don't care about learning the ins and outs of the metaphysics of causation, then I would recommend giving Part I only a cursory reading and getting into the applications in Part II as quickly as possible. You could go directly to Part II, referring back to Part I only as needed (I hope the cross-references, the index and the table of contents will give you all the guidance you need). If you would like to read just enough of Part I to grasp the outlines of my account of causation, I would suggest reading Chapters 3, 4 (especially 4.1 through 4.8), and 9, while skipping the technical material, like the proofs and detailed examples.

Alternatively, if your interests lie exclusively in the field of philosophical logic or theories of causation, there is no reason for you to read Part II at all. In addition, you should feel free to jump around within Part II all you wish: the order of the chapters is not essential. My only recommendation would be for you to read Chapter 12 before reading Chapters 14, 16, 17, 19 or 20, and to read Chapter 14 before 16 or 17.

1.10 A Glossary of Symbols

Although this book contains a considerable number of formulas of symbolic logic, the meanings of the formulas is nearly always spelled out in plain English. There are a few logical and mathematical symbols that the reader must be familiar with:

Logical Symbols

- $\neg$ represents negation, "it is not the case that..."
- $\vee$ represents inclusive disjunction, "either ... or ... (or both)".
- & represents conjunction, "both ... and ..."
- $\rightarrow$ represents a conditional, "if ..., then ..."
- $\leftrightarrow$ represents the biconditional, "... if and only if ..."
- $\forall x$ represents universal quantification, "every object x is of such a kind that ..."
- $\exists x$ represents existential quantification, "there is at least one object x of such a kind that ..."
- $\Box$ represents necessity, and $\Diamond$ represents possibility.
- $\Box\!\!\rightarrow$ represents a non-truth-functional conditional: $(\phi \Box\!\!\rightarrow \psi)$ means that ψ is extremely probable (objectively speaking), conditional on ϕ. These conditionals warrant defeasible inferences.
- $\phi[t/x]$ represents the substitution of x by t throughout formula ϕ.
- $Pr(A/B)$ represents the conditional probability of A on B.

Metalinguistic Symbols

- $\models$ represents the relation between a token (or a token in a model) and a type relative to a model, where $\mathcal{M}, s \models \phi$ is true just in case s supports type ϕ (according to model $\mathcal{M}$. In accordance with standard mathematical practice, I also sometimes use the $\models$ symbol to represent the relation of logical consequence or implication between formulas or propositions (especially in Appendix A).
- $|\approx$ represents the relation of nonmonotonic or defeasible consequence, defined in Appendix B.
- $\vdash$ is used in representing the inference rules of a logical system. The symbol $\dashv\vdash$ represents a two-way, or reversible, inference rule.
- $\|t\|$, $\|\phi\|$ represent the interpretations of symbols t and ϕ in the model under consideration.

Set Theoretic Symbols

- $\in$, $\subseteq$ represent membership and subset, respectively.
- $\cup$, $\cap$ represent union and intersection.
- $\emptyset$ is the empty set.
- $R[\{A\}]$ is the image of A under relation R, that is,the set of all of the objects that are related by R to something in A.

In addition to these familiar symbols, I will make use of a significant number of special symbols. These are all introduced at appropriate places in the text, but I have assembled them all here as well, for the sake of later reference.

Symbols of Mereology

- $\sqsubseteq$ represents the non-strict part-to-whole relation (everything bears this relation to itself.
- $\sqsubset$ is the symbol for proper parthood (asymmetric).
- $\sqcup$ and $\sqcap$ represent mereological union and intersection, respectively.
- $\bigcirc$ represents mereological overlap (having a part in common).
- $\hat{x}\phi$ represents the mereological sum of all the things that satisfy the open formula ϕ.

Special Primitive Symbols

- As represents the actuality of situation s (its being part of the actual world).
- $|\equiv$ is used to form a higher-order type by conjoining a situation-token and a type, i.e., the expression $(s|\equiv \phi)$ represents the type that is realized by any token s' whenever s supports the type ϕ. $|\equiv$ is an object-language counterpart to the metalinguistic $\models$.
- $\prec$ represents the relation of causal priority. (This is primitive in Chapter 5, but definable according to the model built in Chapter 6.)

Defined Symbols

- $\prec_0$ represents immediate causal priority: $s \prec_0 s'$ just in case s is prior to s', and there is nothing intermediate between any part of s and any part of s'.
- $\rhd$ represents the total cause relation: $s \rhd s'$ if s is a total cause of s'.
- $\rightsquigarrow$ stands for causation in the sense of Mackie's INUS condition: an insufficient but necessary part of an unnecessary but sufficient condition. I also use this symbol to represent the closely related idea of causal relevance of one fact to another.
- $|\sim$ represents the relation of causal constraint between types.
- N stands for the immediate causal succession relation: sNs' means that s' is the mereological sum of all the situations immediately posterior to s.
- $\approx\!\!>$ and $\mapsto$ represent the simple and robust carrying of information.
- The expression $(s : \phi)$ is used to represent an ordered pair consisting of a situation s and a type ϕ. These ordered pairs are typically used to represent actual or possible facts.