1 Logic and Information

1.1 Speed Limits and Firefighters

One day, a few years ago, my partner and I were driving along New Zealand's Highway 1 near Wellington when we were confronted with contradictory speed limit signs, as depicted in Figure 1.[1]

It seemed as if we and the other motorists were being told two incompatible things: the unique maximum speed limit on that stretch of road was 50 kilometres per hour and that the unique maximum speed limit was 70 kph.

I suggest that in this circumstance there is conflicting information. Conflicting information is a significant phenomenon in our everyday lives, and it presents the theoretical logician with interesting difficulties. These difficulties concern both how to treat the logical consequences of the information that one has and what to advise agents about how to change their information states in response to such conflicts. Sections 3 and 4 of this *Element* are about these problems, respectively.

Cases like the one given previously also illustrate important issues regarding the nature of information. First, I argue in Section 2 that having a piece of information is different from having a belief. To foreshadow this argument, in the speed limit case, we did not believe that the intended speed limit was 70 kph or that it was 50 kph. We had the information, say, that the speed limit was 50 kph, but neither of us believed that it was 50 kph. Thus, having a piece of information and believing it are distinct.

Second, the speed limit case raises an interesting normative issue. Whereas the prudential thing may have been to travel at the lower speed limit, the presence of the higher one might *entitle* my partner to drive at 70 kph. At least if she were caught going over 50 kph but below 70 kph on this stretch of road, she could argue that the presence of the 70 sign did entitle her to do so. Consider the following rather similar case. Suppose that you are shopping and find a product with two price stickers. You take it to the counter and argue that you should pay the lower of the two prices. The presence of the lower price in this case (given certain consumer laws and conventions) entitles you to pay the lower price. You might even believe that the lower sticker was put on by mistake, but that does not remove the entitlement. The information that the product has this lower price generates this entitlement.

I do not give a theory of informational entitlements of this sort. Doing so would force us to take detours into the philosophy of law and ethics and away from logic. What I do look at is two ways in which the topics of information

[1] As photographed on a phone through a dirty windscreen. Photo by the author.

Figure 1 Highway 1 outside Wellington

and logic interact. I look at ways in which logic can help us treat information. Sections 3 and 5, in particular, are about what logic can do for information. They talk about how different logical systems can help us organise information and how logical techniques can tell us how to update our information states in response to new information. In addition, theories of information can help us understand problems in the philosophy of logic. In Section 3, I look at using ideas from the theory of information to interpret logical systems, such as Kleene's 3-valued logics, the logic of first-degree entailments (FDE), and relevant logic. In Section 4, I use the theory of information states developed in this Element to treat the way in which we learn (or information is conveyed) through logical derivations.

Here is a similar sort of case due to Michael Dunn [28]. Suppose that you are in a hotel room and you hear the fire alarm sounding. You exit your room to see two firefighters one closer to you than the other. Firefighter 1 is pointing to your left and says 'Go this way (down a staircase). It is the only way out'.

Firefighter 2 is pointing to a staircase on your right and says the same thing. The firefighters case raises other interesting issues.

Like the speed limit case, you have to make a decision concerning on what information you are going to use as a basis for your actions. In the speed limit case, one would usually merely travel at 50 kph or less and know that they were avoiding trouble in so doing. In firefighters case, however, one does not have a natural solution of this sort. Suppose, however, you see that one of the firefighters has a book in his pocket and the title of this book is *Terrorist Manual, Volume II: The Use of Disinformation to Increase Harm*. Seeing this would prompt you to **update** your information state and remove the direction that this alleged firefighter is pointing towards as a means of escape. This shows that information, in the sense that I am using it in this Element, is *defeasible*. Section 5 treats updating information both defeasibly and non-defeasibly.

1.2 What Is Information? Preliminaries

Consider the difference between the case of the two speed limit signs and the following scenario. One night, some drunken students painted a crossing in the middle of a tunnel in Wellington. One end of the crossing ran up against a wall and so it was a crossing to nothing. But there was no real sense in which the painting presented us with the information that there was an official crossing in this tunnel. The painting work, as a consequence of the students' state of inebriation, was very sloppy and would fool no one. Unlike the case with the speed limit signs, which looked official, there was no appearance that the crossing was *grounded* in the right sort of facts that make a series of white stripes on a road into an official crossing. This *appearance of grounding*, on my view, is what is essential to the notion of information.

A piece of information is grounded if it has the relationship to a fact (a 'truthmaker') such that it is made true by that fact. One might wonder why I require that pieces of information merely appear to be grounded rather than be actually grounded. I admit that there is an important problem with taking the salient property of information that it appear to be grounded. What appears to be grounded to one person, might not appear to be grounded to another. This sort of subjectivity can cause problems in determining what is a particular agent's information state at a particular time. But getting rid of this subjectivity comes with other costs. In Sections 2 and 6 I discuss the grounding relation.

One way to make the notion of information more objective is to require as Luciano Floridi [32] does that information always be true. Then what counts as information is not agent relative in the way that I suggest. But Floridi's notion of information eschews disinformation and misinformation. I think that some of the most important reasons for analysing information using logic are

to show how logical methods can be used to help eliminate misinformation and disinformation from one's information state.[2]

The issue of the subjectivity of information is rather complex. Some authors think of information as something that we possess in a rather literal sense. Information might be something contained in a database, in a notebook, or even in someone's mind. Others, such as Jon Barwise and John Perry [9], think of the information that is available to an agent as contained in his or her environment.

I agree with Barwise and Perry that the notion of an information state has to include one's environment as well as what is in his or her mind or in a particular database. Many of us keep information in various locations in our environments. I keep all my appointments (or try to) in the calendar in my university email program. I keep the receipts I use for tax purposes in an envelop or in my email inbox. I keep lists of family birthdays in a notebook. And so on. All of these elements of my information state are representational. But, in a very straightforward sense, I also possess a lot of information in a non-representational manner. Facts themselves, if they are available to me, can count as information. As I write this, I see my dog in her bed taking a nap. The fact that I perceive is information that is available to me. As J.J. Gibson [37] says, this sort of information is available to us perceptually in our environment.

1.3 Information Conditions versus Truth Conditions

Throughout this Element, I talk about the *information conditions* for statements. The notion of an information condition is very closely related to the truth condition of a statement but is importantly different.

This distinction is clearest perhaps when we talk about universally quantified statements [57]. In a Tarskian model for classical logic, a universal statement such as 'All the neighbourhood dogs are in the park' is true if and only if in the model every object o in the domain satisfies the open formula $N(x) \supset P(x)$. In other words, the Tarskian truth condition of 'All the neighbourhood dogs are in the park' is true if and only if every neighbourhood dog is in the park. This is a clear instance of the disquotational nature of truth conditions.

General information is very different from a truth condition of this sort. One might see each dog – Sadie, Nova, Milo, and so on – in the park and not realise that every neighbourhood dog is in the park. Something else is required. As Bertrand Russell [79] points out, what one needs to know is that this list of dogs is the list of all the dogs from the neighbourhood.

In order to capture this aspect of universal statements, Russell adds general facts to his ontology. I doubt that general facts are needed as truthmakers

[2] I think this now. In my book [55], I accepted the view that all information is true.

for statements. General information of this kind, however, is needed in order to explain how individuals have access to general truths about their environments. Such general information, however, is not mysterious (as the notion of a general fact may be). I have lived in this neighbourhood for years, and am there much of my time. I see people walking their pets and have gotten to know their dogs quite well. I bring this background with me and it is part of the environment and, together with each individual dog's being there, constitutes the information condition for 'All the dogs in he neighbourhood are in the park'.

Throughout this Element I discuss the information conditions associated with the various logical operators.

1.4 Information States

Another key concept in this Element is of an *information state*. Information states are states of agents – an information state is the information that the agent has in their possession at a given time. What is meant by 'possession' here is rather loose. I follow most of the philosophers of information in taking the way in which one holds their information to be either internal or external to their brains. The view that an information state can be (at least in part) external to the agent is akin to the *extended mind* hypothesis. The extended mind hypothesis of Andy Clark and David Chalmers [18] claims that the mind does not exclusively reside in the body, that one's memories, say, can reside in computers, note pads, and so on, outside of his or her body. One can have just as ready access to information that is represented in a computer's database as information that is remembered.

Our information states can change radically over time. Our environment changes and we sometimes we forget or lose old information. In Section 5, I discuss the application of dynamical logical methods to track the way in which information states change and are updated.

Modelling information states is crucial to the understanding of the relationship between logic and information, but the nature of such states is far from straightforward. One problem is about the coherence of information states. Suppose that an information state i contains two pieces of information. Does it contain their conjunction? For example, when I saw the two speed limit signs did I have the information that the speed limit was both 70 kph and 50 kph? Or did I merely have the information that it was 70 kph and the information that it was 50 kph? The principle of inference being questioned here is the rule of adjunction, namely,

$$\frac{i \vDash A \quad i \vDash B}{i \vDash A \wedge B}.$$

The rule of adjunction is quite important. We use adjunction to conglomerate information to make inferences. Getting rid of adjunction altogether produces a logic that is extremely weak and of little use. There are systems, such as Peter Schotch and Ray Jennings's forcing logic [81], that allow one to make conjunctive inferences from consistent sets of information. Although these systems are formally interesting, they do share a problem. We often infer information from conflicting sources and then forget where pieces of information come from. Forcing allows us to conglomerate such pieces of information as long as they do not conflict with one another. Thus, it allows us at times to use information from conflicting sources. It seems just as reasonable to use a logic that allows us to conjoin information that has explicit conflicts. (In Section 3, I examine discussive logic, which is non-adjunctive, and construct a simple modification of its semantics that makes it adjunctive.)

Another issue has to do with the metaphysics of information states. Information is propositional – it tells us something about the world. Are information states merely sets of propositions? In Section 4 I argue that what are called 'structured propositions' are involved in information states, but in Section 3 I give reasons why an information state should largely be thought of as a set of *indices*. An index, in the sense used in semantics, is a point at which a statement can be evaluated as true or false. Some examples of indices as used in semantic theories are possible worlds, times, events, and situations. A semantics that uses indices is called an 'indexical semantics' or a 'pointed semantics'. I discuss various choices of indices for a semantics of information in Section 3.

There are two aspects to the logical analysis of information states. The first one might be called the *static* aspect. The static logic is a theory of the logical principles under which the information contained in a state is closed. For example, if a state contains the information that A and the information that $A \rightarrow B$ ('A implies B'), then the logic might dictate that the state also contains the information that B. As I explain in Section 3, this closure principle is natural under certain interpretations of implication (and hence in certain logical theories) but not on other interpretations. The choice of a static logic is important for a theory of information because it, in part, determines what information is considered to be available to agents.

The second aspect is the *dynamic* aspect. The dynamic logic of information states determines the way in which states should be updated with respect to new information. The dynamic theory that I investigate in Section 5 is a modification of the dynamic epistemic logics of Johan van Benthem, Alexandru Baltag, and Sonja Smets. It combines elements from standard epistemic logics with elements from the sort of dynamic logic used to analyse computer programs. On the semantical theory, an information state is a set of indices. When a piece of new

information that is consistent with the existing state is added, this set shrinks to include only those indices in the old state that make true the new information. What is more interesting is what happens when information that is inconsistent with the old state is added. I discuss this issue at length in Section 5.

1.5 Plan of the Element

The purpose of this Element is both to introduce readers to the field of the logic of information and to pursue a particular theory in this field. In this book, I try to cover the mainstream theories: the Carnap-Bar Hillel theory of the quantity of information, Dretske's theory of information content, Floridi's definition of information, situation semantics, Barwise and Seligman's channel theory, the use of dynamic logic to understand updating and revising information states, and the use of complexity theory to understand information in logical proofs. But, I have used being asked to write this book also as an opportunity to revise my old views on logic and information in my book, *Relevant Logic* [55], and produce a theory based on relevant logic and its semantics. This attempt to construct a theory is a theme that ties this Element together (especially Sections 3, 5, and 6).

In Section 2, I attempt to define information as propositions that appear true in a context. I use Luciano Floridi's definition of information as true, well-formed data as a foil, criticising the various aspects of it to motivate the definition I take. In particular I claim that there is false information. In this, I side with Michael Dunn and oppose Floridi, Grice, Barwise and Seligman, and many others. In Section 3, I look at ways in which to model information (in particular its static aspects), and pay special attention to models that allow inconsistent information. I examine possible world semantics, closed set semantics, the four-valued Dunn-Belnap semantics, and the ternary relation Routley-Meyer semantics for relevant logic.

In Section 4, I look at ways to understand the quantity of information in logical truths and increases in information that occur when proving theorems. In order to understand the nature of the information involved, I appeal to the distinction between implicit and explicit information. Implicit information is understood in terms of unstructured propositions – propositions as sets of situations. Explicit information is understood in terms of structured propositions – propositions that have individuals and properties as constituents. I think of a proof as the extraction of explicit information from premises. The difficulty of the proof (as measured by the complexity of the proof) is a measure of the amount of information in the proof. This is what I call an 'upstream' measure of information. A downstream measure is one that measures the difficulty of

integrating a proposition into one's beliefs – the degree to which an agent would have to change their beliefs or other mental states rationally to accept that proposition as a belief.

Section 5 examines the dynamic aspect of information. When new information is integrated into an agent's information state, sometimes changes need to be made. I develop semantics of straightforward updating and defeasible updating that are based on the ternary relation semantics discussed in Section 3. In Section 6, I look at the Carnap-Bar Hillel (CBH) theory of information. CBH is based on a set of worlds and a probability function on those worlds. The virtues of this theory are discussed, and it is given an interpretation in terms of physical probability. CBH is also generalised to fit with almost any non-classical logic, including relevant logic. I discuss the prospect of making the connection between probability and information stronger, that is, by taking probabilities as guides to information. This idea, at least in its simplest form, is rejected.

2 What Is Information?

2.1 Introduction

In order to understand how information theory and logical theory can be made to help one another, we first need to have a workable concept of information. Although I have discussed the notion of information in the introductory section, a lot more work needs to be done before we have a useful concept. In characterising the notion of information, Michael Dunn says:

> The sense of information I have in mind can be nicely contrasted to knowledge. Plato [thought of] knowledge to be something like 'justified true belief.' Then information is what is left of knowledge when one takes away belief, justification, and truth. Of course we commonly speak of someone having information to mean it as a kind of success (truth) but this can be shown to be a mere conversational implication by the fact that we can say of a person that she has false information. Information in the sense I mean it is not even mere opinion, for that requires at least some weak sense of belief. 'She had the information all the time, but she didn't believe it at all.' Information is meant here as a kind of semantic content – the kind of thing that can be expressed by language. [25, p 423]

I have in mind a notion of information something like Dunn's, although I think there is a notion of justification at play with information. I have already said that what counts as information always appears to be true. The notion of appearance that is involved here is somewhat tricky, and is discussed at length in Section 2.7.

I discuss the relationship between information and truth in Section 2.6. I make more of an argument there concerning the separation of information