8

THE LOGIC OF INFORMATION

Patrick Allo

Introduction

The combination of *logic* and *information* is popular as we as controversial. It is, in fact, not even clear what their juxtaposition, for instance in the title of this chapter, should mean, and indeed different authors have a given a different interpretation to what *a* or *the logic* of information might be. Throughout this chapter, I will embrace the plurality of ways in which logic and information can be related and try to individuate a number of fruitful lines of research. In doing so, I want to explain why we should care about the combination, where the controversy comes from, and how certain common themes emerge in different settings.

Logic, in its most reductive sense, is the study of good and bad arguments. Here, the term "argument" has a rather narrow meaning: it is a relation between a set of expressions called "premises" and a single expression called the "conclusion." As a formal field of study, the aim of logic is to develop precise criteria for telling good and bad arguments apart. The class of good arguments is traditionally identified with the class of *deductively valid* arguments: arguments where the conclusion *follows from* the premises, or where the conclusion merely makes explicit what was already implicit in the premises.

In view of this informal description, it is natural to claim that good (or deductively valid) arguments are exactly those arguments where:

> **(CN)** The content of the conclusion does not exceed the combined content of the premises.

And this establishes nothing short of a deep conceptual connection between the core task of logic and the notion of informational content.[1] In practice, however, the above principle does not play a major theoretical role in the development of logic. The idea that good arguments can be used to extract information from premises is not used to characterise the class of good arguments: the principle expressed by (CN) is the *explanandum* (what needs to be clarified) rather than the *explanans* (the clarification itself). As van Benthem and Martinez remark:

> [Logic] has official definitions for its concepts of proof, computation, truth, or definability, but not of information!
>
> *(van Benthem and Martinez 2008)*

The class of deductively valid arguments is standardly characterised by model-theoretic means, which (i) yields a formal explication of the idea that deductively valid arguments should preserve truth (if all the premises are true, then the conclusion must be true as well), that (ii) is provably equivalent to a characterisation that refers to the existence of a proof that starts with the premises and ends with the conclusion. It is thus no surprise that modern logic includes model-theory and proof-theory as two of its main pillars (computability or recursion-theory is often included as a third pillar), but not information theory.

While the mainstream approach in logic treats information as redundant for serious formal work, the suggestion that logic and information are intertwined is too deeply ingrained in the history of logic as well as in our informal talk about the subject to disappear altogether.

Hintikka, a major figure in twentieth-century logic, even went as far as claiming that the absence of a logical analysis of the notion of information was one of the scandals of the development of modern logic:

> Logicians have apparently failed to relate their subject to the most pervasive and potentially most important concept of information.
>
> *(Hintikka 1973)*

Two general principles underlie many natural connections between logic and information.

(LB) The logical is the lower-bound of the informative.

If, for instance, Alice tells Bob something he could have figured out for himself on the basis of what he already knew, she didn't tell him anything genuinely informative – Bob didn't have to ask additional questions to Nature to figure it out for himself.

(UB) The counter-logical (or absurd) is the upper-bound of the informative.

If Bob tells Alice something that he could not even in principle figure out by asking additional questions to Nature, Alice told him something that could not possibly be true. It is, to use a formulation from Carnap and Bar-Hillel (1952, 8), "too informative to be true".

In the fourth section we shall see how these two principles arise in specific formal settings, and relate to (CN).

The rise of the philosophy of information and the broadening of the scope of logic initiated by the dynamic and interactive turn in logic (see also Chapter 12) have led to a renewed interest in the conceptual connection between logic and information. And indeed, there are several theoretical reasons why the question deserves our attention as well.

The historical connection: In Medieval logic principles like (CN) are explicitly included as a necessary condition for valid consequence.[2] By extending information-theoretic views to modern logic, we can emphasise the continuity between modern and traditional conceptions of validity.

The deductive/inductive gap: A characterisation of valid arguments as truth-preserving arguments suggests a deep gap between deductive arguments and inductive arguments that merely make the conclusion more plausible. On an information-theoretic characterisation we have a natural progression: in deductive arguments the premises provide all the required information, whereas in inductive arguments they provide some but not all the information in favour of the conclusion.

*Problems with truth-talk***:** The description of certain formal enterprises in logic in terms of the traditional concepts of truth and truth-preservation does not always do justice to what they try to achieve. Descriptions of non-classical logics, for instance, seem to imply radical changes to what we mean by truth. Informational descriptions of the same enterprises often lead to a more conservative picture, and overall facilitates a pluralistic outlook on logic (Allo and Mares 2012).

*Attention for conceptual problems***:** Logic often clashes with the use of intensional idioms, and leads to unintuitive results when used within the scope of intensional operators like knowledge or belief. These issues surface in many different contexts, including the problems of logical omniscience and hyperintensionality in logics of knowledge and belief, and the problem of granularity in natural language semantics. The common trait of these problems is that logical equivalence is too coarse to be used as an account of sameness of meaning or sameness of content. To get a better grip on the general structure of these problems, it is advisable to study the relation between logic and information.

Overall, it seems that the connection between logic and information is not just intuitive, but that if we cannot conceive of logic as a means for (or as a model of) information manipulation, the study of logic itself loses much of its appeal.

A logical background

In this preliminary section, I give a brief overview of the basic building-blocks of formal logic, and make the claim that logic is the study of good deductive arguments more precise.

Logic is a formal discipline, but our use of the term "logic" itself is, even in scientific and scholarly contexts, often surprisingly sloppy. The description in the introduction is no exception to this rule. In particular, even though I related logic to a particular class of good arguments (the deductively valid ones), I did not for instance specify whether "good" was meant as a normative (how we should argue, reason, etc.) or as a descriptive (how we actually argue, reason, etc.) delimitation of this class of arguments. In addition, while I clearly distinguished between information and information theory, I did not draw a similar line between logic as a subject-matter and logic as a field of study. If we want to relate logic to information, we'd better be clear about this, for it is an entirely different question whether information (or information theory) plays a role in our best logical theories, or whether the facts of validity are just facts of information-containment. Clearly, the first question is empirical in the sense that we can just examine the existing theories and, as previously indicated, note that information theory is virtually absent. The second question, by contrast, cannot be addressed directly – we have no direct access to facts about validity – but is best recast in terms of how thinking about facts of validity as facts of information-containment can lead to better theories of validity.

Talking about logical theories requires a lot of care, but is overall easier than talking about logic. To do so, we first need to introduce the idea of a formal language.

Formal languages

By a formal language, we mean a schematic language that is introduced by, first, specifying what the logical and non-logical symbols of our language are, and, second, by listing formation-rules or ways in which symbols are combined in admissible expressions or well-formed formulae of our language. Logicians often study different types of languages, and it is easier to grasp the idea of a formal language by considering a specific example.

The language of propositional logic: In a propositional language, the non-logical symbols are called atomic expressions, denoted by the letters *p*, *q*, *r*, … The standard – so-called Boolean – operators: & (and), ∨ (or), ⊃ (implies), and ¬ (not) are its logical symbols. Using these building-blocks, we say that:

1 All atomic expressions are well-formed formulae;
2 If "*A*" and "*B*" are well-formed formulae, then "*A* & *B*", "*A* ∨ *B*", and "*A* ⊃ *B*" are well-formed formulae as well;
3 If "*A*" is a well-formed formula, then "¬ *A*" is also a well-formed formula; and
4 Nothing else is a well-formed formula.

With these guidelines, we can always find out whether or not a given string of logical and non-logical symbols is a well-formed formula of the language of propositional logic.

Model-theory

Once we have a formal language, we can really start to develop the model and proof-theory of a given logical system.

In model-theory we develop an account of what it means for a *case* to support a formula, and use this to characterise the class of valid arguments. The former is done by exploiting the systematic structure of our formal language. Taking the language of propositional logic as an example, we stipulate that:[3]

1 a case *c* supports "*A* & *B*" if and only if *c* supports "*A*" *and* supports "*B*",
2 a case *c* supports "*A* ∨ *B*" if and only if *c* supports "*A*" *or* supports "*B*",
3 a case *c* supports "*A* ⊃ *B*" if and only if *c* supports "*B*" *whenever* it supports "*A*",
D a case *c* supports "¬ *A*" if and only if *c* undermines "*A*".

An argument with $A_1, \ldots, A_n$ as premises and *B* as conclusion is valid if and only if every case that supports $A_1, \ldots, A_n$ is also a case that supports *B*.

What actually follows from what is then made entirely dependent on what *cases* are. When we require that cases are such that for every case *c* every atomic expression *p* is either supported or undermined by *c* (but never both supported and undermined!), and furthermore agree with the requirements 1–4 above, we have characterised the class of valid arguments of *classical propositional logic*: the logic based on the presuppositions that a formula *A* and its negation ¬*A* are jointly exhaustive and mutually exclusive, and where as a consequence every formula of the form "*A* ∨ ¬*A*" is supported by all cases (we call such formulae *tautological*), and every formula of the form "*A* & ¬*A*" is undermined by all cases. We call such cases *complete* and *consistent*.

Different classes of good arguments can be characterised by weakening these requirements, but also by adding further structure to the nature of cases. This additional structure is typically coupled with the introduction of additional logical symbols, as in the case of first order classical logic, but also the modal logics of knowledge and belief we will encounter later on.

Amidst these alternatives, classical logic has a special status. This can be seen from the fact that (1) the uninterpreted "is supported by" relation can naturally be understood as "is true at", and given this assumption (2) the resulting class of valid arguments will coincide with the truth-preserving arguments. I elaborate further on this second point in the section on Soundness and Completeness. For the first point, it suffices to observe that if we only consider complete and consistent cases, the resulting identity between "not supporting" and "undermining", and between "supporting" and "not undermining", coincides with the orthodox identities between "absence of truth" and "falsity" (completeness), and between "truth" and "absence of falsity" (consistency). In other words, relative to complete and consistent cases, we can uncontroversially equate support with truth, and undermining with falsity.

Proof-theory

In many ways, the existence and construction of proofs form the focal point of logic. By a proof, we mean a set T of formulae $B_1,\dots, B_n$ that (a) are organised in a list, tree, or other type of structure, where (b) a possibly empty subset *Prem* of T is taken to be given (the premises of the proof); (c) all other formulae in T are obtained by applying certain rules to the formulae that are also in that list; and (d) a single formula A in T is called the conclusion. One standard form for a proof is just an ordered list of formulae, where the first n formulae are the premises, all other formulae are obtained by applying rules to the formulae higher up in the list, and the final formula is the conclusion.

When we have such a list, we say that A can be deduced from *Prem*. Furthermore, when *Prem* is empty, we say that A is a theorem (it can be deduced from zero premises).

Completeness and truth-preservation

Good deductive arguments are standardly understood as truth-preserving arguments. Yet, so far we have only explicitly identified the class of valid arguments in terms of the preservation of support, and characterised a second class of good arguments in terms of what can be deduced. Using a simplified version of Kreisel's *Squeezing argument* (Kreisel 1967),[4] we can elucidate how these different notions hang together, and indeed agree on a single notion of consequence.

Consider the following two claims:

Claim 1: Validity is a necessary condition for the preservation of truth.
Claim 2: Deducibility is a sufficient condition for the preservation of truth.

Jointly, these amount to the thesis that at least all correct deductions are truth-preserving, and that at most all support-preserving arguments are truth-preserving. Given a set of sensible and intuitively correct rules of proof, we should be confident in the second claim (and indeed, we can examine each rule to confirm this). When it comes to the first claim, we just need to repeat our prior observation that the mutually exclusive and jointly exhaustive notions of support and undermining are just as fine-grained as the concepts of truth and falsity. Consequently, every truth-preserving argument will also preserve support. The support for the last insight is usually phrased in terms of the contrapositive claim: if an argument does not preserve support it will not preserve truth either. More exactly, we can easily transform a model that supports all the premises and undermines the conclusion into a description of how the world should be for the premises to be all true, and the conclusion to be false.

At this point, a formal result can be used to complete the argument. Since the notions of support and proof are entirely formally specified (proofs as well as models are mathematical structures), it can be rigorously proved that for every support-preserving argument there should also be a proof. Such results are called completeness-theorems. When added to our initial two claims, we thus extend our argument:

> *Proof:* Validity is a sufficient condition for deducibility.
> *Conclusion:* Validity and deducibility are necessary and sufficient conditions for truth-preservation.

In summary: given that truth-preservation lies between deducibility and validity, the completeness-theorem shows that truth-preservation is squeezed between the provably unique characterisations of good arguments in terms of deducibility and validity. These features of our formal concepts of validity and deducibility make the connection between logical consequence and truth-preservation particularly attractive.

Structural properties of consequence-relations

Logicians are not only interested in the equivalence of model and proof-theoretic characterisations of consequence-relations, but they also study the properties of the consequence-relations independently of how these are characterised. In particular, they investigate the so-called *structural* properties of (abstract) consequence-relations. If we use the turnstile (⊢) to refer to the consequence relation of classical logic, we can say that it has the following structural properties:

> **Reflexivity:** $\Gamma \vdash A$ whenever $A \in \Gamma$.
> **Cumulative transitivity:** If $\Gamma \vdash A_i$ for all $A_i \in \Delta$, and $\Gamma \cup \Delta \vdash B$, then $\Gamma \vdash B$.
> **Monotony:** If $\Gamma \vdash A$ and $\Gamma \subseteq$ ”, then ” $\vdash A$.

Information as constraint, as resource, and as goal

The notion of information does not only appear to be redundant, but is also a source of further confusion. The latter arises because the notion of information can be used twice in our informal descriptions of what logic is about. It can be used as a constraint on what counts as a good argument – as in (CN) – but also to refer to the information we extract from our premises. As such, information is both the *goal* of logical reasoning and a *constraint* on good arguments. As has often been pointed out, these roles conflict. The so-called paradox of inference attributed to Cohen and Nagel is one among many examples of this tension:[5]

> If in an inference the conclusion is not contained in the premises, it cannot be valid; and if the conclusion is not different from the premises, it is useless; but the conclusion cannot be contained in the premises and also possess novelty; hence inferences cannot be both valid and useful.
>
> *(Cohen and Nagel 1972, 173)*

This is surely disconcerting, and only adds up to the initial impression that characterisations of logical consequence in terms of truth-preservation are not only technically convenient, but also philosophically more satisfactory. Truth serves as a constraint by ruling out arguments that would allow us to step from truth to falsity, and is also the aim of deductive inference: we want to derive new truths from previously accepted truths. Clearly, no conflict arises in this case, and this may suggest that we should not think of logic as something that regulates our use of cognitive resources like information. When we use logic to characterise the class of truth-preserving arguments, we choose our formal notions of support and undermining so that they are jointly exhaustive and mutually exclusive. Given these formal properties, the corresponding notion of information inevitably fudges the distinction between the absence of positive information and the presence of negative information. This obscures the partial nature of holding information, and makes it harder to use logic to reason about information as a cognitive resource.

Moving to a stronger notion of information doesn't help much here. If we adopt Floridi's suggestion that the informativeness of *A* is not only inversely proportional to the probability of *A*, but also proportional to how accurately and precisely *A* describes the truth (Floridi 2004), the problem remains. As illustrated by how scientists treat significant digits in calculations, no valid argument will ever increase the accuracy of our premises.

When two intuitively plausible principles are in conflict, this often indicates that they need to be formulated more precisely. The conflict exposed in Cohen and Nagel's paradox of inference (and many other similar insights) is the result of equivocation, and if we want to develop logic as a theory of how we should use cognitive resources like information, we need to refine the building-blocks of our theory, and pay more attention to the double role of information in logic.

One obvious way to deal with such conflicts is to look for potential equivocal uses of the technical terms that are involved. And since information is a notorious multifaceted concept, it is first in line for a closer examination. With our initial characterisation of deductive arguments we can easily resolve the conflict between the two uses of information. Deductions yield useful information because they make explicit what was already implicit in the premises, and valid because the only information that can be made explicit is the information that was already implicitly available from the premises. In loose terms: deduction transforms an available resource into a readily accessible resource.

The emphasis on the difference between implicit and explicit information is correct, but not entirely satisfactory. Our intuitive feel that information is both a constraint on and a goal of inference is merely transferred on the implicit/explicit divide, but the distinction itself remains to be explained. A good example of why the distinction is less clear than we might think is due to Gilbert Harman and Robert Stalnaker, who point out that it is regularly used to refer to two different distinctions.[6] Explicit information can mean information that has actually been derived, or it can mean information that is readily (or easily) accessible. Yet, since information can be explicit in one sense, but not in the other (think for instance of a complex formula that is listed somewhere in a very large unsorted list), logical deduction is not the only way to make explicit what was merely implicit (Harman 1986; Stalnaker 1991). Surely, crossing the border between implicit and explicit information requires computation, but not all computation is deduction.[7] Search and retrieval is just as much a computational process.

Logic and information systems

As an alternative to the refinement of the different senses of information, we could also explore different perspectives on the nature of logic, and the design of logical theories. My suggestion is that, on a narrow conception of logic, information can only act as a constraint because it fails to register the fine distinctions that are needed to reason about information as a partial and distributed cognitive resource.

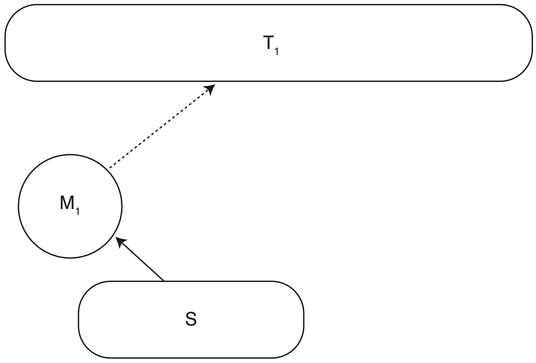

Figure 8.1 System, model, theory

The above insights can be made more precise by introducing a rudimentary notion of a system, and our means for describing and modelling such systems as information systems (see Figure 8.1).

By a *system* I mean any part of reality that can be the subject of further examination. On this account, any part of reality can be thought of as a system. This includes the whole universe, some well-defined spatiotemporal fragment of the universe like the state of a coin on the table, a multi-agent system (e.g. what different agents know and believe about the state of a coin on the table), or even a highly organised repository of information like an archive or a library.

A *theory* about a system S is a set of expressions that gives the best possible description of S given the language we use as well as (in a yet undefined sense) our knowledge of the system. A *formal theory* T is a theory that is formulated in a formal language, and a *logical theory* T' is a formal theory T that has been extended with all the logical consequences of T. We say that T' is the deductive closure of T,[8] and use it as a highly idealised description of our best knowledge of the system under consideration.

A *model* of a system is a set of observables (see Chapter 7 of this Handbook) that can be thought of as a set of facts about the system under consideration. A *formal model* is a mathematical construction (for instance based on set theory) that is used to determine which expressions of a formal language L are supported or undermined by a model.[9] A *total model* is maximally specific in the sense that it either supports or undermines every formula in L.

Formal languages like the propositional language described earlier in this chapter can be used to formulate theories about a given system, and the models of propositional logic (the *cases* from the second section) can be understood as formal models in the above sense. There are two important differences between total models and theories.

First, since models are mathematical structures, they are described in the language of mathematics, and this language is often richer than the formal languages we use to construct our theories. Consequently, it can happen that two distinct models M_1 and M_2 cannot be told apart with the limited resources of a certain formal language L. No theory that is formulated in L will be supported by M_1, but not by M_2. Put loosely: M_1 and M_2 contain more information than can be revealed with the limited resources of L. For the language and models of propositional logic, there is no such principled gap, but this is a rather exceptional situation. The study of the expressive limits of languages is an important field of formal study that is not without significance for the philosophy of information (van Benthem and Martinez 2008, 227–8), but lies outside the scope of this chapter.

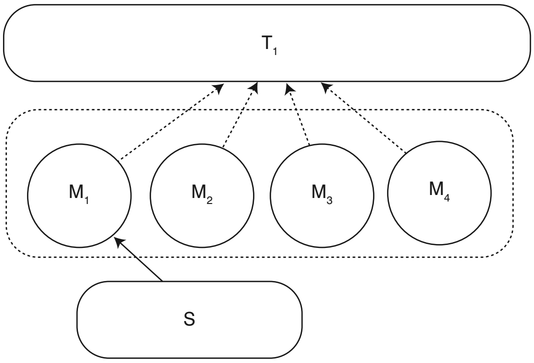

Figure 8.2 System, models, theory

Second, and more importantly, a theory T can be partial in the sense that (even when it is deductively closed) there can be formulae A such that neither A nor $\neg A$ is in T. For every sufficiently complex system there will be features of that system of which we are ignorant and (assuming that these features can be expressed in the formal language we use) our best knowledge of the system – our theories – will have gaps. To recapture this idea at the level of (total) models, we have to resort to the use of sets of models (as depicted in Figure 8.2), and stipulate that our theory T about a given system S is the set of all formulae of L that are supported by all models that, as far as we are concerned (this qualification is important: some of these models are in fact not models of S!), are equally good models of S. Conversely, we can say that the set of all equally good models of S are just those models that support all formulae in our theory T. The underlying idea that our ignorance about certain features of the system under consideration is reflected in the number of models that are equally good models of that system is further emphasised by calling these models *possibilities* (this can be compared to the sample space of possible outcomes in probability theory). Depending on how we conceptualise our information about the system, we shall call such possibilities epistemic

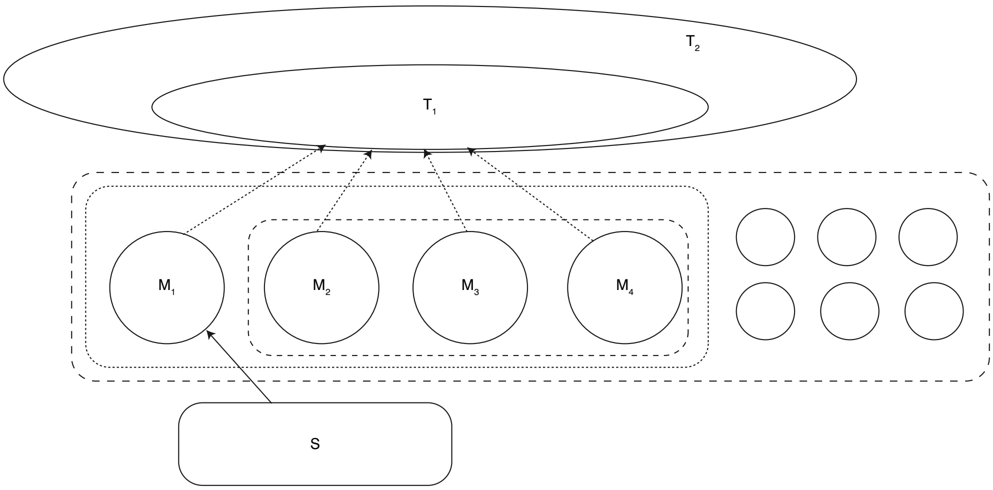

Figure 8.3 System, models, theory (2)

(our knowledge of the system), doxastic (our beliefs about the system), or informational (our information about the system).

At the most basic level, we can distinguish between these by contrasting the two situations depicted in Figure 8.3. We consider two sets of models

M1 = {M1, M2, M3, M4} (dashed line)
M2 = {M2, M3, M4}, (dotted line)

and note that only M_1 is really a model of *S*.[10] As a consequence, since every formula in T_1 must be supported by each model enclosed in the dotted line, it must also be supported by M_1, and since M_1 is an actual model of *S*, every formula in T_1 is a correct claim about *S*. As such, we can think of the models $\mathbf{M_1}$ as a set of epistemic possibilities: the information in $\mathbf{M_1}$ is truthful. In the second case, however, the real model of *S* is not included in $\mathbf{M_2}$, and we have no guarantee that T_2 will be entirely correct. In this case, we can think of the models $\mathbf{M_2}$ as a set of doxastic possibilities. Finally, on a weak notion of semantic information (see Chapter 6), both $\mathbf{M_1}$ and $\mathbf{M_2}$ will count as sets of informational possibilities. On a stronger veridical account, only $\mathbf{M_1}$ will count as a set of informational possibilities. The weak notion is the standard in logical theorising, but the stronger notion has its place as well.

The above description of systems, models and theories gives rise to two equivalent notions of information about a system, which coincide with the qualitative approach that underlies the classical accounts of Carnap and Bar-Hillel (1952) and Kemeny (1953). Here, I presented it explicitly as a means of individuating the modeller's information about a system.

The first notion is related to the theory *T* about the system *S*, and identifies the informational content about the system with the subset of non-tautological formulae in *T*. The second notion is related to the set of (total) models of *S*.

If, as in Figure 8.3, we use $\mathbf{M_1}$ and $\mathbf{M_2}$ to denote two sets of models of *S*, and T_1 and T_2 to denote the two corresponding theories about *S* (*A* is in *T* if and only if *A* is true in all models), we can say that:

- T_2 is at least as informative about *S* as T_1 if and only if T_1 is a subset of T_2.[11]
- $\mathbf{M_2}$ is at least as informative about *S* as $\mathbf{M_1}$ if and only if $\mathbf{M_2}$ is a subset of $\mathbf{M_1}$.

And this reveals a familiar inverse-relation between the size of theories and the size of sets of models: a larger theory means more information, but a larger set of models means less information. This is what van Benthem and Martinez (2008) call *information as range*.

Parenthetical remark The reference to the size of sets in the above principles should not be understood quantitatively (counting models or formulae), but qualitatively in terms of set-inclusion. This has two important consequences. First, it means that two sets of models or two theories (e.g. $\{M_1, M_2, M_3\}$ and $\{M_2, M_3, M_4\}$ in Figure 8.3) can be incomparable in the sense that each contains information that is not included in the other (figuratively speaking, they both could learn from each other). This qualitative aspect is what distinguishes logical approaches from probabilistic approaches. Second, it allows us to bypass the problem that, except in some borderline cases, sets of models as well as deductively closed theories are countably infinite and have thus the same cardinality. *End of parenthetical remark*

Using just the slightest bit of the language of set theory, this gives us:

$$\mathbf{M_2} \subseteq \mathbf{M_1} \text{ if and only if } T_1 \subseteq T_2$$

Which, if we talk about possibilities instead of models, expresses the idea that more information means fewer possibilities and *vice-versa* (Barwise 1997).

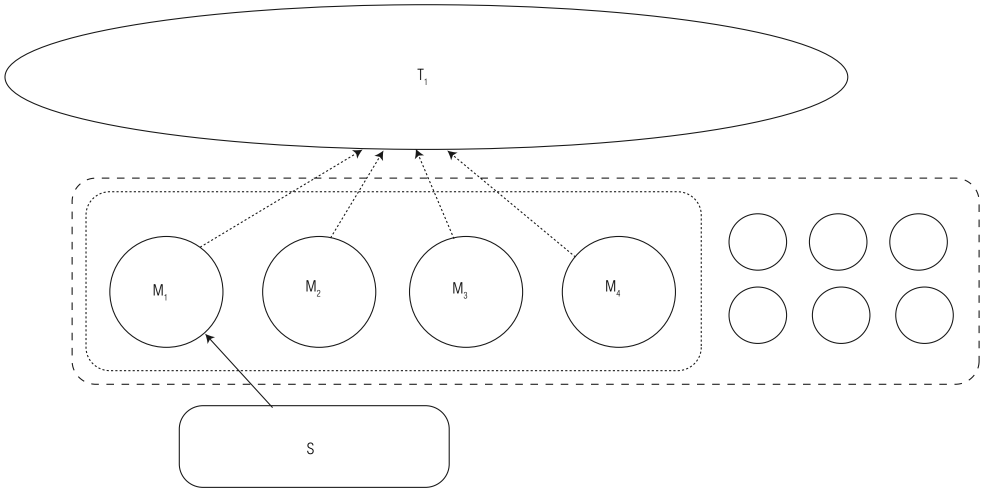

Figure 8.4 System, logical space, theory

Statically, this inverse relationship principle identifies having information about a system with being able to exclude certain possible states of the system. From a dynamic perspective it identifies receiving or obtaining information about a system with the exclusion of possible states of the system.

While one's information about a system *S* can be identified with the set of possible states of *S* one can exclude, this account remains useless without an explicit account of the *total space of possibilities*. Given what has already been said, this space has to be identified with the set of all logical possibilities: new information is information that cannot be logically derived from one's prior information, but this also means that the only possibilities that can be excluded are logical possibilities.

The overall picture becomes even clearer if we look at the extremes, as expressed by the earlier suggestions that informativeness is bounded from below by the logical and from above by the counter-logical.

- If *A* is uninformative, it does not require the exclusion of any possibility, so the logical must be supported by any possibility.
- If *A* is over-informative, it requires the exclusion of every possibility, so the counter-logical must be undermined by every possibility (equivalently: supported by the empty set of possibilities).

More generally, we will then identify the content of T_1 with the logical possibilities that are excluded by T_1 (i.e. the logical possibilities that undermine some or all of the formulae in T_1). In Figure 8.4 this coincides with the possibilities ~~in the grey area~~. If, in addition to T_1, one learns that some formula *A* is true as well, and *A* is supported by $\{M_1, M_2, M_3\}$, but not by M_4, the new information that is obtained is $\{M_4\}$, namely the one possibility that is excluded from $\mathbf{M_1}$ by learning that *A* is true.

Hence, we can see that if we take information about a system as our starting point, informativeness is constrained by logic: only logically contingent expressions can be non-trivially informative (if it can be true at all, then it is only informative if it is logically contingent). If, as we did in the previous section, we take valid arguments as our starting point, it is informativeness (or rather the lack thereof) that constrains logic.

When we reason about a system from an external perspective (the viewpoint of the modeller), logic is used in its constraint role. The idea is that we do not want to add anything to our theory that does not follow from that theory (and if theories are deductively closed, this means we do not want to change anything at all), for this comes down to excluding a possibility, and perhaps even the exclusion of the one possibility that is the actual model of the system under consideration. In other words: going beyond one's information – as would be the case if one would erroneously derive a formula that is included in T_2 but not in T_1 (see Figure 8.3) – is one way to step from truth to falsehood. Here, information acts as a constraint because it precludes erroneous or fallacious reasoning about the system.[12]

Sub-systems and distributed information

The description of systems in the previous section has two important features.

> *System*: the total information about the system is identified with the information in the correct model of the system. At the level of the model-theory this is the relative complement of the correct model in the total logical space; at the level of the formal language this is the set of formulae that are supported by the correct model.

> *Modeller*: there is one agent with partial information about the system. The agent's information can again be modelled as a set of formulae (a theory) or as a set of models.

In a more traditional setting we would identify the total information about a system with what is true *of* the system, or with the facts about the system. Here, we make it explicit that information is relative to the relevant level of abstraction (see Chapter 7). This means that what counts as a true claim about the system is mediated by the correct model of that system as well as by the language we use to describe that model. Apart from this relativisation, our thinking about the total information about a system is rather restrictive. The models of the system are unstructured repositories of total information, and the modeller is an unstructured repository of partial information.

These restrictions are quite useful to clarify the constraint role of information in a modeller's reasoning about a system, which is why I associated it with an external perspective on the system. In a sense, we adopted a relatively high level of abstraction: we ignored the deeper structure of the system, and only considered the information of a single modeller. From a formal point of view (and slightly simplifying matters), the adoption of an external perspective on logic and information is closely related to seeing consequence-relations (i.e. the "turnstile" $\vdash$) as informational or information-containment relations.

If we want to clarify the goal role of information in logic, we need to take into account the underlying structure of a system. Since almost anything can be a system, it can have many underlying structures as well. If we slightly abuse our terminology and provisionally ignore the distinction between a system and models of a system, we can just say that systems have sub-systems, and that the information about the system as a whole is distributed between its sub-systems. If a system is some mechanical artefact like a machine, then some piece of information about the state of that machine can be specifically located in one of its components (information is situated), while this component need not be the locus of all the information about the state of the machine (information is partial).[13]

The suggestion that information could be located in a physical object should not be taken literally. What we mean is that we can obtain information about a system by examining some of its sub-systems. To make this insight precise, we need to add structure to our models and speak of the information supported by a model and its sub-models. We then can say that a sub-model supports the information we can obtain by examining the sub-system it is a model of in exactly the same way – namely relative to a certain level of abstraction – as the correct model of a system supports the information we can obtain by reliably examining a system. The only relevant difference between a model and a sub-model is that sub-models do not satisfy the principle that allows us to infer the presence of negative information from the absence of positive information. The notions of support and undermining are not in general exclusive and exhaustive relative to sub-models.

From a formal point of view, the internal perspective on how logic and information are related is best approached by looking at how implication-relations and other conditional expressions are informational relations, and how the logic of these relations becomes a logic of information flow.

To see how partial and distributed information are deeply intertwined, we can look at a number of specific types of systems and their sub-systems.

Moments in time: A fairly intuitive type of system that can be modelled as a set of sub-models is a history that consists of a succession of different stages in time. Every time-stage contains all the information about itself (it's own "now"), but need not hold all (or even any) information about the past and the future. As a whole, the history contains all the information about every possible moment in time, but specific moments only hold information about moments that are informationally accessible. Thus, the past can be accessible thanks to all kinds of records (we have information about the past because we have informational technologies (Floridi 2014)), and the future can be accessible because some future events are entirely determined by the past and present (and we may know about these connections). Moments in time can be seen as a variant or special case of the situations described below.

Situations: If the system we consider is a part or even the whole world, we can add structure by distinguishing the situations that are part of it. This is the approach from *situation theory*, one of the pioneering frameworks of the modern connection between logic and information (Barwise and Perry 1999; Devlin 1991; Israel and Perry 1990). On this account, a system is (or can be modelled as) a structure $(S, \leq\_)$, with S a set of situations, and $\leq\_$ a relation between three situations, say s, t, v, such that $s \leq_t v$ expresses that by combining information from s with information from t, we obtain information about v (Mares 2010; Allo and Mares 2012). This way of structuring systems allows us to model the following informational phenomena:[14]

1 If c is a situation that contains information about certain regularities in the world (Barwise (1993) calls such situation *channels*), then $s \leq_c t$ means that the information about regularities that is available in c, provides information that can be used to infer the presence of information in t from information that is available in s. That is, s carries information about t in virtue of c, and if an agent is attuned to the information in c, she will be able to exploit information in s to obtain information about t. Thus, X-rays carry information about such-and-so's bone being broken (Israel and Perry 1990), or smoke seen from the bottom of the mountain tells us that there must be a fire on the mountain (Barwise and Perry 1999).

2 If we posit the existence of a set of logical situations Log $\subset S$, then we can define an information-inclusion relation $\leq$ between situations:

$s \leq t$ iff $s \leq_l t$ for some l in Log,

which holds if and only if all the information available in s is also available in t. It is, in other words, a relation that signals that – as an informational, but not necessarily as a concrete physical entity – one situation is part of another situation.

Such relations between situations can be used to model different types of implication-relations (Barwise 1993). For the general case, we have an implication relation $\rightarrow$ that can express informational connections between situations. Formally: the information that $A \rightarrow B$ will be available in a situation s if and only if for all situations t and u such that $t \leq_s$ u it is the case that if t has the information that A, then u has the information that B.

Other types of implication-relations can be seen as special cases. For instance, $\rightharpoonup$ can express connections between situations that are part of each other. Formally: the information that $A \rightharpoonup B$ will be available in a situation s if and only if for all situations t such that $s \leq t$ it is the case that if t has the information that A, then t has the information that B as well. In a sense, this *intuitionistic implication*, expresses connections between stages of information-accumulation. Finally, our implication relation from classical logic $\supset$ (see second section) expresses the most trivial kind of informational connection: the information that $A \supset B$ is available in a situation s if and only if it contains information that $\neg A$ (negative information that A) or that B.

Agents: If the system we are modelling is a multi-agent system with agents that can adopt various attitudes towards the current state of the system – such as knowing that A, believing that B, being ignorant with respect to C – different formal perspectives can be adopted.[15] We can model systems as so-called *interpreted systems* that treat global states of a system as the Cartesian product of the local states of the environments as well as of the agents (Fagin *et al.* 1995). Or, we can model systems as multi-modal Kripke-models where each possible state (including the actual state) of the system is included in this model, and a relation between these states models the information available to each agent. This can be seen as an elaboration of how we modelled the information available to the modeller of a system, and will receive a more detailed treatment in Chapter 12.

In either system, we can explicitly describe how the information in the system is distributed. We can for instance say that Alice knows that A, but that Bob does not know this.

$$K_a A \,\&\, \neg K_b A$$

By extending the language of propositional logic with modal operators for knowledge, we make room for partial information ($K_a A \vee K_a \neg A$ is not a logical truth) without having to invalidate $A \vee \neg A$.

Additionally, we can say that Bob knows that B, and that Alice knows that Bob knows whether B

$$K_b B \,\&\, K_a(K_b B \vee K_b \neg B),$$

but that Bob believes that Alice ~~knows~~ that he is ignorant with respect to B

$$B_b K_a(\neg K_b B \,\&\, \neg K_b \neg B).$$

This ability of expressing higher-order attitudes (one agent having information about the information of another agent) allows us to express a first type of informational connection between agents: knowing, believing or otherwise being informed of the state of other agents.

A second type of informational connections can be expressed by assigning information to groups rather than to individual agents. Thus, we can say that Alice and Bob both know that either *A* or *B* is true (with E for "everybody knows"):

$$E_{\{a,b\}} (A \vee B),$$

or that by pooling their knowledge they know *A* & *B* (with D for "distributed knowledge"):

$$D_{\{a, b\}} (A \& B).$$

In either case, these connections bear on how (and how much) information is shared between agents.

A third type of informational ~~connections~~ does not only relate to how information *is* shared, but also to how it *can be* shared. That is, related to how agents can pass on information. This is the topic of dynamic epistemic logic (Baltag and Moss 2004; Baltag, van Ditmarsch and Moss 2008; Van Ditmarsch, van der Hoek and Kooi 2007), and one of the more active areas of development related to the dynamic and interactive turn in logic. Its primary concern is the individuation of several ways in which information can be communicated through public or private announcements, and how such announcements can modify the distribution of information. For instance, if neither Alice nor Bob knows which side of the coin lies face up, the public announcement that it lies *HEADS* up will lead to the common knowledge that it lies heads up, and hence make this information freely and transparently available to all agents in the system: the agents went from common ignorance to common knowledge. If, by contrast, it is privately announced to Alice that it lies *HEADS* up, she will come to know that it lies *HEADS* up, while Bob will remain ignorant, and will now falsely believe that they both are still ignorant: the agents went from common ignorance to unevenly distributed knowledge and even error. Not unlike the channels and constraints of situation semantics that express natural regularities, different types of announcements act as regularities of social interaction.

Conclusion: local changes versus global invariance

The place of information in logic is confusing because it serves two different and even opposed roles: it can serve as a constraint (the content of a valid argument should not exceed the combined content of its premises), and as a goal (we use logic to extract information from our premises, or more generally from our environment). This double role can be further clarified by looking at information systems. If we adopt an external perspective on such systems, the emphasis is on the fact that the total information in the system cannot increase,[16] and that our reasoning about the system should be constrained by the information we have about the system. Formally, this is associated with logic as a consequence relation: sets of premises are an unstructured repository of information, and all that can be extracted was already there from the start.

If, by contrast, we adopt an internal perspective on such systems by no longer treating them as unstructured information-repositories, we put more emphasis on how information is distributed within the system. While the total information in the system can still not increase, the distribution and thus the information available in certain sub-systems can indeed increase. Information, in this case, becomes the goal of logic, and logic becomes a logic of information flow. Formally, this is associated with various modal and conditional operators that allow us to describe features of informational connections and distributions of information.[17]

Notes

1 I here use the term "informational content" to speak of the semantic content or semantic information (as described in Chapter 6) we measure (quantitative) or individuate (qualitative).
2 See the principle "(Co) In a valid consequence, the conclusion is contained/understood in the premises." mentioned in Dutilh-Novaes (2012), and attributed to Abelard amongst others.
3 The terms "support" and "undermine" are non-standard, but perfectly suit our aims: we can be undecided between supporting and undermining a given proposition, and although it seems problematic to both support and undermine the same proposition, it is not entirely inconceivable either.
4 The full strength of Kreisel's argument surfaces only relative to full first-order logic, but the basic idea of looking at necessary and sufficient formal conditions for a precise though informal concept can be used for didactic purposes as well.
5 Chapter 15 deals explicitly with the question of how we should think about the information contained in logical and mathematical theorems, and how we can characterise the information obtained through deductive inference. Here, I only rely on these problems as a means to clarify the relation between logic and information.
6 Stalnaker and Harman focus on the concept of belief, but their diagnosis can be straightforwardly transferred to our thinking about information.
7 If we take complexity concerns into account, the converse is false as well: some deductions are not even feasible computations.
8 Given a theory *T*, we obtain the deductive closure of *T* by adding to *T* all the (and indeed, infinitely many) logical consequences of *T*. Analogously, we say that *T'* is deductively closed if and only if it already contains everything that follows from it.
9 Logicians often say that a model *M* is an interpretation of the language *L*. Here, we stick to our prior use of "support" and "undermine" instead of the more traditional notion of "truth in a model".
10 I leave aside the question of whether there can really be more than one genuine model of reality, or even whether there can ever be exactly one genuine model of reality. If one takes the method of abstraction seriously, it should already be clear that what counts as a real model of reality depends on the level of abstraction one adopts.
11 The subtraction of tautological formulae does not make a difference here.
12 Here too, what counts as an error is itself relative to the level of abstraction we adopt in our reasoning about the system. In particular, it depends on what can and cannot be said in and discerned with the formal language we use, and how we structure the logical space in which we situate our models of the system. These choices are themselves not constrained by logic.
13 This description still allows for information about a system that is not located in one of its strict sub-systems, but only in the total system.
14 When compared to moments in time or agents, situations have the peculiar feature that as sub-models or sub-systems they are of the same type as the system or model itself. The total system or model is just the largest situation.
15 The former is common in computer science, whereas the latter has become the standard in philosophical logic.
16 Abramsky (2008) relates this to Shannon's identification of information with negative entropy.
17 Similar features can also be recaptured by consequence relations, which then no longer satisfy all the structural rules of classical logic. This topic falls outside the scope of the current chapter.

Further reading

This chapter barely covers all of the pertinent ways in which logic and information can be related. In addition to the works cited in the text, the interested reader is invited to consult the following works:

1 The traditional connection between possibilities and content: Stalnaker (1984) and Rayo (2012).
2 The connection between non-classical logic, situation semantics, and information: Mares (1997, 2009, 2010); Restall (2005); Wansing (1993).
3 Informational semantics as a philosophical account of logical consequence: Allo and Mares (2012); Saguillo (2009)
4 The specificity of *being informed* in comparison with knowing and believing: Floridi (2006); Allo (2011).
5 The semantics of informative and inquisitive content as a common ground in conversations: Ciardelli *et al.* (2013).
6 An account of the role of logic within the philosophy of information: Chapter 12 of the e-Textbook *The Philosophy of Information: A Simple Introduction* available at http://www.socphilinfo.org/teaching/book-pi-intro.

References

Abramsky, Samson. 2008. Information, processes and games. In *Handbook on the Philosophy of Information.* Ed. Johan Van Benthem and Pieter Adriaans. Elsevier.

Allo, Patrick. 2011. The logic of 'being informed' revisited and revised. *Philosophical Studies* 153 (3): 417–434.

Allo, Patrick and Edwin Mares. 2012. Informational semantics as a third alternative? *Erkenntnis* 77 (2): 167–185.

Baltag, Alexandru, Hans P van Ditmarsch, and Lawrence S Moss. 2008. Epistemic logic and information update. In *Philosophy of Information (Handbook of the Philosophy of Science).* Ed. Pieter Adriaans and Johan van Benthem. Amsterdam: North-Holland.

Baltag, Alexandru and Lawrence S Moss. 2004. Logics for epistemic programs. *Synthese* 139 (2): 165–224.

Barwise, Jon. 1993. Constraints, channels, and the flow of information. In *Situation Theory and Its Applications Vol. 3.* Ed. Peter Aczel, David Israel, Stanley Peters, and Yasuhiro Katagiri. Stanford, CA: CSLI.

Barwise, Jon. 1997. Information and impossibilities. *Notre Dame Journal of Formal Logic* 38 (4): 488–515.

Barwise, Jon and John Perry. 1999. *Situation and Attitudes.* Stanford, CA: CSLI Publications.

Carnap, Rudolf and Yehoshua Bar-Hillel. 1952. An outline of a theory of semantic information. In *Language and Information: Selected essays on their theory and application.* 221-74. Reading, Mass; London: Adidison-Wesley.

Ciardelli, Ivano, Jeroen Groenendijk, and Floris Roelofsen. 2013. Inquisitive semantics: A new notion of meaning. *Language and Linguistics Compass* 7 (9): 459–476.

Cohen, Morris R. and Ernest Nagel. 1972. *An Introduction to Logic and Scientific Method.* London: Routledge & Kegan Paul.

Devlin, Keith. 1991. *Logic and Information.* Cambridge: Cambridge University Press.

Dutilh-Novaes, Catarina. 2012. Medieval theories of consequence. In *The Stanford Encyclopedia of Philosophy.* Summer 2012 ed. Ed. Edward N Zalta.

Fagin, Ronald, Joseph Y Halpern, Yoram Moses, and Moshe Y Vardi. 1995. *Reasoning About Knowledge.* Cambridge, MA and London: MIT Press.

Floridi, Luciano. 2004. Outline of a theory of strongly semantic information. *Minds & Machines* 14 (2): 197–222.

Floridi, Luciano. 2006. The logic of 'being informed'. *Logique & Analyse* 49 (196): 433–460.

Floridi, Luciano. 2014. *The Fourth Revolution How the Infosphere Is Reshaping Human Reality.* Oxford and New York: Oxford University Press.
Harman, Gilbert. 1986. *Change in View. Principles of Reasoning.* Cambridge, MA: MIT.
Hintikka, Jaakko. 1973. *Logic, Language-Games and Information. Kantian Themes in the Philosophy of Logic.* Oxford: Oxford University Press.
Israel, David and John Perry. 1990. What is information. In *Information, Language and Cognition.* Ed. Philip Hanson. Vancouver: University of British Columbia Press.
Kemeny, John G. 1953. A logical measure function. *The Journal of Symbolic Logic* 18 (4): 289–308.
Kreisel, Georg. 1967. Informal rigour and completeness proofs. *Studies in Logic and the Foundations of Mathematics* 47: 138–186.
Mares, Edwin. 1997. Relevant logic and the theory of information. *Synthese* 109 (3): 345–360.
Mares, Edwin. 2009. General information in relevant logic. *Synthese* 167 (2): 343–362.
Mares, Edwin. 2010. The nature of information: A relevant approach. *Synthese* 175 (1): 111–132.
Rayo, A. 2012. *The Construction of Logical Space.* Oxford: Oxford University Press.
Restall, Greg. 2005. Logics, situations and channels. *Journal of Cognitive Science* 6: 125–150.
Saguillo, José M. 2009. Methodological practice and complementary concepts of logical consequence: Tarski's model–theoretic consequence and Corcoran's information-theoretic consequence. *History and Philosophy of Logic* 30 (1): 21–48.
Stalnaker, Robert. 1984. *Inquiry.* Cambridge MA: MIT Press.
Stalnaker, Robert. 1991. The problem of logical omniscience, I. *Synthese* 89 (3): 425–440.
van Benthem, Johan and Maricarmen Martinez. 2008. The stories of logic and information. In *The Stories of Logic and Information.* Ed. Pieter Adriaans and Johan van Benthem. Amsterdam: North-Holland.
Van Ditmarsch, Hans, W. van der Hoek, and Barteld Kooi. 2007. *Dynamic Epistemic Logic.* Dordrecht: Springer.
Wansing, H. 1993. *The Logic of Information Structures.* Berlin: Springer-Verlag.