THE REVIEW OF SYMBOLIC LOGIC
Volume 11, Number 4, December 2018

COMPLIANCE AND COMMAND I—CATEGORICAL IMPERATIVES

KIT FINE
Department of Philosophy, New York University

Abstract. I develop a semantics for imperatives within the truthmaker framework by taking the meaning of an imperative to be given by the actions that are in compliance with or in contravention to the imperative.

The main aim of this series of three articles is to develop a truthmaker semantics for the logic of imperative and deontic sentences. The first part deals with categorical imperative sentences, the second with deontic sentences and their interplay with categorical imperative sentences, and the third with the interplay between indicative, imperative and deontic sentences and with the conditional forms of imperative and deontic sentences in particular. It would be helpful, though not strictly necessary, to have some standard exposition of truthmaker semantics at hand (such as in Fine (2016a)); and I have, for the most part, been content with informal exposition, although the reader may consult the appendix for technical detail.[1]

There are two central ideas behind the application of the truthmaker approach to the semantics and logic of imperatives. The first is that the semantics is to be understood in terms of the actions which are in compliance with or in contravention to a given imperative sentence. However, in line with the general spirit of the truthmaker approach, compliance and contravention are taken to be *exact*. The compliant or contravening actions must be wholly relevant to the imperative with which it complies or to which it contravenes. Thus shutting the door is compliant with the imperative 'Shut the door!' while shutting the door and opening the window is not.

The second idea is that the validity of an imperative inference is to be understood in terms of the mereological relationship between the actions in compliance with the premiss and conclusion of the inference. More particularly, we should be able to see the content of the conclusion as part of the content of the premiss, where this is a matter of every action in compliance with the conclusion being part of an action in compliance with the premiss and every action in compliance with the premiss having as a part an action in compliance with the conclusion.

Received: May 7, 2016.
2010 *Mathematics Subject Classification*: 03B45.
Key words and phrases: possible worlds semantics, truthmaker semantics, imperative logic, Ross's paradox, entailment, partial content, free choice.

[1] I developed the central ideas of the article around 2010 and have since lectured on them in a number of venues and presented them in a number of seminars. I should like to thank the participants at these lectures and seminars for their many helpful comments; and I am particularly indebted to Maria Aloni, Ivano Ciardelli, Ian Rumfitt, and Peter Vranas. Two referees for the journal also made many very helpful comments on both parts of the article.

I might add, on a more personal note, that I first became interested in imperative inference as an undergraduate around 1966, when I read the articles of Kenny (1966) and Geach (1966). I came up with the idea of truth maker semantics around 1969 and had even considered the application to imperative logic, but it took me almost half a century to appreciate the full scope and significance of the approach.

doi:10.1017/S175502031700020X

It is worth mentioning some key features of the account which set it apart from some other, more prominent, accounts in the literature. In the first place, the account is hyper-intensional. Indeed, certain truth-functionally equivalent sentences will not be assigned the same semantic or logical role. Thus even though p is truth-functionally equivalent to $p \vee (p \wedge q)$, the account will distinguish, as it should, between 'Raise an arm!' versus 'Raise one or both arms', since the second imperative allows one to raise both arms while the first does not; and the account will not sanction, as it should not, the inference from the first imperative to the second.[2] Indeed, I believe that the adoption of an intensional stance (and the adherence to the possible worlds framework, in particular) has been one of the main obstacles to developing a satisfactory semantics and logic for imperatives.

In the second place, the account is action- rather than outcome-oriented. The semantics for imperative is not given in terms of the *result* of complying with the imperative but in terms of the *action* by which compliance is achieved. In this regard, the account differs from the 'static' approach (as in Chellas (1971), for example), under which the content of an imperative is represented by a set of worlds, those which might result from its having been obeyed (or appropriately complied with); and it also differs from the 'dynamic' approach (as in Segerberg (1990), for example), under which the content of an imperative is represented by a 'transition' relation on worlds, relating a given state of the world to any of the states which might result from the imperative having been obeyed.

One of the main advantages of an action-oriented approach, to my mind, is that it relates the semantic import of imperatives directly to their implementation. We may imagine a robot which has a certain repertoire of actions at its disposal—moving its left or right leg forward x inches, swivelling x degrees clockwise or counterclockwise, sounding the alarm etc. etc. If we wish to tell the robot what to do, then it is clearly desirable if our instructions can relate directly to the actions in its repertoire. If we merely say, make it the case that the vase is on the table, then this is bad for the robot, since it has to go through the complicated task of working out how this is to be done, and it is bad for us, since we have no assurance that, in complying with our instruction, the robot will not end up doing something we would prefer it not to do. Although there are other models of how we might interact with the robot, there is something especially attractive about a model in which the connection between an imperative and the actions is requires is made completely explicit.[3]

Of course, there is a way in which the dynamic approach is also action-oriented. For the transition relations correspond to the actions by which they are induced. But the actions themselves are represented in terms of the result of performing those actions, whereas we leave them alone, so to speak, and simply regard them as being among the states or events in the world. It is possible to connect actions with transition relations on our approach, but they are not conceived as such and nor do these relations play any role in specifying the basic semantics. This has the great advantage that we do not need to face the problem of which transition relations should be associated with negative actions (not doing something) or concurrent actions (doing one thing at the same time as another). These problems do not go away, but they do not prove integral to the development of the basic semantics.

---

2 A more general argument against the intensional approach to statements of permission, which also applies to imperatives, is given in Fine (2014a).

3 We have here something akin to the distinction between declarative programming (as in Prolog) and what, significantly, is called 'imperative' programming.

In the third place, we achieve a kind of unity in the semantical and logical treatment of imperatives and indicatives. For the content of an indicative sentence is given by the states that verify it while the content of an imperative sentence is given by the actions that comply with it. Thus verification is to indicatives as compliance is to imperatives; and the semantical treatment of the two notions can be seen to run on parallel tracks, with corresponding clauses for the connectives and with a corresponding account of valid inference. It is clearly desirable to be able to suppose that the connectives have essentially the same meaning in application to the two kinds of sentence.

This unity is lost on the dynamic approach since the content of an indicative sentence (a set of worlds) is essentially different from the content of an imperative (a relation on worlds); and so no uniform semantic treatment of the two cases is possible. Some unity can be achieved on the outcome-oriented approach (which no doubt has been regarded as a strong point in its favor) but without any of the benefits that accrue from giving the semantics of imperatives in terms of actions rather than worlds. Thus the unity is achieved on the present approach by thinking of a state as verifying an indicative in much the same way as an action complies with an imperative. From this point of view, the apparent inapplicability of a truth-conditional model of meaning to imperatives lies not in an inadequacy in the truth-conditional model as such but in an inadequate conception of what it is.

I should perhaps make a few remarks about my general approach to the topic. My interest is in what one might call the underlying semantics and logic for imperatives. The linguistic evidence in this area, as in any other, is incredibly complex. There are many different kinds of factors at work; and if one is to make any progress in understanding the evidence at hand, then one must isolate a set of factors, which work together in explaining some significant aspect of the evidence. In this regard, it is unfortunate that there is no language of the gods, in analogy with the celestial motion of the planets, that is relatively free from the kind of interference that prevents one from clearly discerning what general principles might be in play.

This means that I have not attempted to take every aspect of our actual use into account. I have been as much guided by a sense of what might be of logical significance in our usage than by a sensitivity to the actual form that it takes. It is perhaps worth recalling, in this connection, that the development of classical logic—one of the great intellectual achievements of the early twentieth century—was achieved neither by close attention to ordinary usage nor by ignoring it altogether but by attempting to isolate a significant aspect of that usage, especially as it relates to mathematical discourse. And it is much the same spirit that has guided my own, more modest, efforts in the present article, with the more formal presentation of programs or algorithms playing somewhat the same role as mathematical discourse.

Of course, the project of getting at what I have called the underlying logic and semantics is not irrelevant to the understanding of natural language. It serves as a first step and as a useful foil against which other more realistic attempts may be made and evaluated; and I have made a number of suggestions throughout the article as to how this might be done. But the project also serves a further important purpose that is to some extent at odds with the attempt to understand natural language. For it provides a basis for systematically regimenting how we might state imperatives and how we should reason with them; and often the needs of such regimentation will run counter to the demand for descriptive adequacy. My view is that the current 'linguistics turn' in the philosophy of language, while beneficial in many respects, has also been detrimental in deflecting attention away from some of the more traditional concerns of conceptual and logical analysis.

The plan of the present part of the article is as follows. I begin by outlining truthmaker semantics for indicative logic (§1); I show how the same form of semantics applies to imperatives (§2); I turn to the concept of validity for imperative inference, giving both a brief informal account (§3) and a formal account (§4); I then look at the logic of imperatives—considering the connection between various logical notions (§5), setting up a formal system of reasoning (§6), and discussing some inferences of special interest (§7); and I conclude with a comparison between my account and some related accounts in the literature (§8).

The reader should bear in mind that several other philosophers and linguists have developed very similar ideas over the last decade or so. Despite the similarities, my account differs in a number of significant respects from theirs. (1) It is far more abstract and, in particular, is not in any way tied to the representation of states or actions as sets of worlds or as sets of atomic sentences and their negations or as sets of atomic states, as is so commonly assumed. Just as Kripke's treatment of worlds as unstructured points was important for the technical development of modal logic, so I believe, the treatment of states or actions as unstructured points is important for the proper technical development of the current approach. (2) It countenances a rich ontology of impossible actions, which play an important role in the development of the theory and enable one, in particular, to distinguish between different inconsistent imperatives, which could not otherwise be distinguished. (3) It comes with a precisely formulated logic, something notably absent in many other accounts. (4) As will become apparent from the subsequent parts, it provides a natural basis for extending the logic to deontic, conditional and quantificational constructions.

§1. Truthmaker semantics for indicative sentential logic. Indicative formulas are constructed in the usual way from indicative sentential atoms $p_1, p_2, \ldots$ by means of negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$) and the verum constant $\top$.

We take each sentence to be *verified* (*made true*) or *falsified* (*made false*) by a state. It is integral to our understanding of verification and falsification that they be exact, i.e., that the verifier or falsifier should be wholly relevant to the sentence that is verifies or falsifies. Thus a state which contains an exact verifier as a proper part will not itself, as a rule, be an exact verifier.

An informal, and somewhat rough, test for when a state is an exact verifier for a sentence is that the sentence should be true *by way of* the state's obtaining. Thus we may say that the ball's being red is a truth-maker for the sentence 'The ball is colored' since the ball is colored by way of its being red. By contrast, we cannot properly say that the ball's being red and round is a truth-maker for the sentence, since it is not true that the ball is colored by way of being red and round.

In more formal terms, we suppose given a *state space*, which is a set of candidate truth-makers and falsity-makers ordered by part-whole; and we call the candidate truth- and falsity-makers 'states', without any implication that they are states in any ordinary sense of the term. We do not presuppose that there is single state space that is in play regardless of the context but allow the state space to vary with the discourse. Thus in a context in which we are not interested in different shades of red, we may have a single state of an object's being red, whereas, in another context, we may wish to discriminate between the object's having different shades of red. We assume that, in a state space, the fusion of any number of states is also a state. This means, in particular, that there will exist the fusion of no states, the *null* state $\square$, and the fusion of all states, the *full* state $\blacksquare$. The null state $\square$ will be a part of every state and nothing is required for it to obtain; and the full state $\blacksquare$ will contain every state and everything is required for it to obtain.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

Of course, given what would appear to be a well-defined notion of fusion on states, there will be a well-defined notion of part, for we may say that state s is a part of state t when the fusion of s with t is identical to t. However, in particular cases it may be hard to say when one state is a part of another. Is John's turning the knob, for example, a part of his opening the door? We may again use the by-construction to provide an informal test for when one state is a proper part of another. For s will be a (proper) part of t just in case s obtains in part by t obtaining. Thus John's turning the knob will be a part of his opening the door since it is in part by turning the knob that he opens the door. On the other hand, his whistling is not a part of his opening the door, even if he whistles as he opens the door, since it is not in part by whistling that he opens the door.[4]

We distinguish between *possible* and *impossible* states and take two states to be *compatible* when their fusion is a possible state. Normally, a state space will contain impossible states since the fusion of two possible but incompatible states will be an impossible state; and so ■, in particular, will be an impossible state. However, a state space may contain many different impossible states besides ■, the most impossible of the impossible states. The state of an object's being red and green, for example, might be distinguished from the state of its being round and square or from the state of its being round and square and red in terms of its composition from the component possible states.[5] In the special case in which ■ is the sole impossible state, we shall say that the state space is *topsy*; and, similarly, in the special case in which □ is the sole necessary state, one compatible with every possible state, we shall say that the state space is *turvy*. And, of course, a space will be *topsy-turvy* when it is both topsy and turvy, with a single impossible state on top and a single necessary state at the bottom.

A *state space model* is a state space accompanied by an assignment of a set of verifiers and a set of falsifiers to each indicative atom. We shall assume that the two sets are nonempty, i.e., that each atom has both a verifier and a falsifier. But we should note that the semantic values of atoms are in no way special; the content of any formula, no matter how complex, is capable of serving as the semantic value of an atom. This is a desirable feature of a semantics if we want to adopt a general interpretation of the sentential atoms, under which they are capable of standing in for any sentence whatever; and under a semantics of this kind, the set of valid formulas will automatically be closed under substitution.

We have the following clauses for when a conjunction or a disjunction or verum is verified by a state within a model:

(1) the state s verifies the conjunction $T \wedge U$ iff it is the fusion $t \sqcup u$ of states t and u which, respectively, verify its conjuncts T and U;
(2) s verifies the disjunction $T \vee U$ iff s it verifies one of its disjuncts, T or U;
(3) s verifies $\top$ iff s is identical to the null state □.

For negation, we make use of the notion of (exact) falsification:

(4) s verifies a negation $\neg S$ iff it falsifies the negated indicative S.

We then introduce 'dual' clauses for when either of the previous complex sentences is falsified:

[4] The by-constructions are further discussed in Fine (2018) and there has been extensive discussion of the construction in the philosophical literature on the identity of actions.

[5] Fine (2016b) shows how a multitude of impossible states can naturally be introduced into a state space containing only possible states.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

(1′) s falsifies the conjunction T ∧ U iff s falsifies one of its conjuncts, T or U;
(2′) s falsifies the disjunction T ∨ U iff it is the fusion $t \sqcup u$ of states t and u which, respectively, falsify its disjuncts T and U;
(3′) s falsifies ⊤ iff s is identical to the full state ■;
(4′) s falsifies the negation ¬S iff it verifies the negated indicative S.

We might define ⊥ as ¬⊤. The clauses for ⊥ are then the reverse of the clauses for ⊤:

(5) s verifies ⊥ iff s is identical to the full state ■;
(5′) s falsifies ⊥ iff s is identical to the null state □.

Within the present framework, we might identify a (*propositional*) *content* or *proposition* either with a set of states (the *unilateral* content) or with a pair of sets of states (the *bilateral* content) and we might take the (*positive*) *content* $[S]^+$ *of* a sentence S to be the set V of its verifiers, its *negative content* $[S]^-$ to be the set F of its falsifiers, and its *bilateral* or *full content* [S] to be the pairing $<V, F>$ of its positive and negative content. It is readily shown that every formula, like every atom, will have both a verifier and a falsifier, although the verifier—as with p ∧ ¬p—may be an impossible state.

We may introduce two notions of entailment—the familiar notion of classical entailment and the less familiar notion of analytic entailment. Say that a model is *classical* if (i) any verifier of an indicative atom is incompatible with any of its falsifiers and if (ii) for any atom, any possible state is compatible with a verifier of the atom or with a falsifier of the atom. Sentence S then *classically entails* sentence T if no verifier for S in any classical model is compatible with a falsifier for T. On the other hand, S will *analytically entail* T if (i) every verifier for S contains a verifier for T and (ii) every verifier for T is contained in a verifier for S.[6] Similar relations, here and elsewhere, can be defined to hold between the contents themselves, rather than between the sentences with those contents.

Classical entailment corresponds, as one might expect, to the notion of classical tautological implication, while analytic entailment corresponds to a slight variant of Angell's notion of analytic implication.[7] It might plausibly be argued that the notion of analytic entailment captures the intuitive idea of the content of one sentence being *part* of the content of another (Yablo, 2014; Fine, 2016a).

§2. Truthmaker semantics for imperative sentential logic. We turn to the corresponding form of truthmaker semantics for imperatives. We may suppose that imperative formulas are constructed in the usual way from imperative sentential atoms $\alpha_1, \alpha_2, \ldots$ by means of negation (¬), conjunction (∧), disjunction (∨) and the verum constant ⊤.

We might think of the atoms as corresponding to simple imperative sentences such as 'shut the door', 'turn on the light' or 'open the window'. We are then able to form negations, such as 'do not shut the door' or 'do not open the window', conjunctions, such as 'shut the door and open the window', and disjunctions, such as 'shut the door or open the window'. More complicated constructions are also possible, of course, as with 'do not both shut the door and open the window' or 'shut the door and either open the window or turn on the light'.

6 We might also add the condition that every falsifier for T should be a falsifier for S. This results in a somewhat different notion of entailment. For p ∧ (q ∨ r) will entail (p ∧ q) ∨ (p ∧ r) without the addition, though not with the addition.

7 To be exact, it is the first variant system of §11 of Fine (2016a).

In the case of indicatives, we have taken the basic semantic notions to be ones in which a state verifies or falsifies a sentence. Similarly, in the case of imperative sentences, we take the basic semantic notions to be ones in which an action is in *compliance with* or is in *contravention to* an imperative. We might also talk in this connection of the imperative *allowing* or *disallowing* the action. Thus the action of shutting the door will be in compliance with the imperative 'shut the door', while the action of leaving it open will be in contravention to the action.

We might well regard an action as a special case of a state, i.e., as being a candidate verifier or falsifier though, for certain purposes, we might wish to distinguish an action from the performance of the action, which would then be the state proper. Actions might be construed as action types not tied to any particular agent (shutting the door), or as action types tied to a particular agent or set of agents (your shutting the door, our building a house), or as action tokens tied to a particular agent or set of agents (your shutting the door or our building a house on such and such an occasion). How we take them will not much matter for present purposes, although it is, of course, important that the interpretation of the imperatives should appropriately correspond with our understanding of the actions.

We regard an action space as a special case of a state space, one in which all the states are actions; and we suppose that it is subject to the standard conditions on a state space. In particular, actions will be closed under arbitrary fusion, so that any sequence or combination of actions, no matter how large or gerrymandered, will also constitute an action.

Just as with verification or falsification, compliance or contravention is understood to be *exact*; the action must be in compliance or contravention with the imperative as a whole. So whereas the action of shutting the door will be in compliance with the imperative 'shut the door', the action of shutting the door and opening the window will not, in the relevant sense, be in compliance with the imperative; and similarly, whereas the action of leaving the door open will be in contravention to the imperative 'shut the door', the action of leaving the door open and closing the window will not, in the relevant sense, be in contravention to the imperative.

There remains a question of which action or actions should be taken to be in contravention to an atomic imperative such as 'shut the door' or, equivalently, in compliance with the negated imperative 'do not shut the door'. There are at least three possible views on this question. One is that it is the negative action of not shutting the door; another is that it is a positive action, such as latching the door, that prevents one from shutting the door; and a third is that it is a total action of performing certain positive actions, which do not include shutting the door, and of doing nothing else.

Part of the difficulty here is that the 'operational' import of a positive imperative is quite different from that of its negation. A positive imperative is obeyed by straightforwardly doing something, but the way in which one obeys a negative imperative is less straightforward and may, to some extent, depend upon the means for action available to the agent. It is possible that the agent may be able to place a block on performing a certain action—the "module", so to speak, by which this action is performed can be rendered inoperative. As long as this action cannot be performed in other ways, it is then appropriate to take an action in contravention to a positive imperative to be the negative action of placing a block on the corresponding positive action. This is in line with the first option and provides a natural sense in which the agent may be said to 'refrain' from performing the action. It is also possible that the agent may be able to perform a positive action which prevents him from performing the positive action in question: and, in such a case, we may take the actions in contravention to the positive imperative to be those which prevent him from performing that action. This is in line with the second option and provides a sense in which the agent might be said to 'prevent' a certain action. Finally, it is possible that neither of these means of complying with the negative imperative action is available to the agent; and, in this case, the agent should take care that the sum-total of the actions she has performed does not include the given action. This is in line with the third option and provides a sense in which the agent might be said to 'exclude' a certain action.

Thus each of the three interpretative options corresponds to different ways of complying with a negative imperative; and, of the three, the first would appear to require least of an agent, in the form of a mere block on action, the second somewhat more, in the form of an alternative positive action, and the third most of all, in the form of a continual monitoring of which actions have been performed. However, it should be noted that the first option does not require that, for *each* action, there should be another action which is its negation. Even even though the actions of shutting the door and of opening the window may have negations, for example, we need not assume that the action of shutting the door and opening the window also has a negation. Thus the action space need only be partially closed under negation on this view.

It should also be noted that, under any of the options, there may well be a number of actions in compliance with or in contravention to a given imperative. Thus two actions, shutting the door or opening the window, will comply with the imperative 'shut the door or open the window', and two actions, not shutting the door open or not opening the window, will be in contravention to the imperative 'shut the door and open the window'. This point is of some importance for, under an action-oriented semantics for imperatives, it is often supposed that there is but one action in compliance with an imperative—as given perhaps by some transition relation on worlds—and no meaningful sense in which we are provided with a *choice* of actions, even though the outcome of a given action may itself be indeterminate.

Among the actions to be taken into account will be the null action, the fusion of zero actions, which we also designate as $\square$. It is an action that will be part of any other action and which one does (as part of) whatever one does. Among the actions may also be impossible actions, such as shutting the door and not shutting the door, which results from the fusing of possible actions.

When it comes to the compliance and contravention conditions for arbitrary imperative formulas, we have exact analogues of the clauses for indicatives:

(1) the action a complies with the conjunction $X \wedge Y$ iff it is the fusion $b \sqcup c$ of actions b and c which, respectively, comply with its conjuncts X and Y;
(2) a complies with the disjunction $X \vee Y$ iff it complies with one of its disjuncts, X or Y;
(3) a complies with $\top$ iff a is identical to the null action $\square$;
(4) a complies with the negation $\neg X$ iff it contravenes the negated imperative X;
(1′) a contravenes a conjunction $X \wedge Y$ iff it contravenes one of its conjuncts, X or Y;
(2′) a contravenes the disjunction $X \vee Y$ iff it is the fusion $b \sqcup c$ of actions b and c which, respectively, contravene its disjuncts X and Y;
(3′) a contravenes $\top$ iff a is identical to the full action $\blacksquare$;
(4′) a contravenes the negation $\neg X$ iff it complies with the negated imperative X.

Of special interest is clause (3) for compliance with the null imperative $\top$; the one and only way to comply with $\top$ is to perform the null action $\Box$. Accordingly, we might read $\top$ as 'do nothing'. But this has to be appropriately understood. It does not mean 'perform no action at all', nor 'do nothing in particular, i.e., anything you like', but rather 'perform that particular action which requires nothing of you'. Thus the nothing here is not a nothing at all, nor an anything whatever, but a something that amounts to nothing (my poor attempt to outdo Heidegger); and compliance with the imperative $\top$ merely requires null compliance, i.e., the performance of the null action $\Box$.

As we have seen, three kinds of content may be associated with an indicative sentence—the (*positive*) *content*, the set of its verifiers, the *negative content*, the set of its falsifiers, and the *full content*, the positive and negative content taken together; and we might identify a *propositional content* or *proposition* with a set of states (its verifiers). Likewise, three kinds content may be associated with an imperative sentence—the (positive) content, the set of actions it allows, a negative content, the set of actions it disallows, and the full content, the two taken together; and we might identify a *prescription* or *prescriptive content* with a set of actions (those that are allowed).

There is clearly a close parallelism in our semantic treatment of indicative and imperative sentences. There is, first of all, a parallelism in our choice of semantical primitives, with verification corresponding to compliance and falsification to contravention. There is, in the second place, a parallelism in the semantical clauses, with the clauses for the one being the exact counterpart of the clauses for the other. Thus when S is an indicative sentence, such as 'you shut the door and open the window' and X = S! is the corresponding imperative, 'you, shut the door and open the window', then an action/state will verify or falsify S just in case it is in compliance with or in contravention to X.

There is a question as to what accounts for this parallelism, both at the syntactic and at the semantic level.[8] On one view, the syntax and semantics for imperatives *derives* from the syntax and semantics for indicatives (as suggested by McGinn (1977), for example). Thus any imperative X is syntactically derived from a corresponding indicative sentence S and the semantics for X is derived from the semantics for S by transforming its verification and falsification conditions into compliance and contravention conditions.

Another view, loosely deriving from Frege (1879) and one which I am inclined to adopt, is that there is a common element to an indicative S and the corresponding imperative S!, what Hare (1952) calls the 'phrastic'. These can then be taken to be subject to positive and negative 'realization conditions', which are the neutral counterpart of the verification and falsification conditions for indicatives and of the compliance and contravention conditions for imperatives. The syntax and semantics for indicatives and imperatives alike is then derived—by means of appropriate mood-indicators or 'neustics' in Hare's terminology—from the syntax and semantics for the phrastics, with realization being given a 'world-to-word' spin in the case of indicatives and a 'word-to-world' spin in the case of imperatives.

There are perhaps other ways of accounting for the parallelism. Fortunately, there will be no need for us to take a stand on this question or even on whether it is meaningful; and although we shall later appeal to the correspondence between indicatives and imperatives, it can be understood without regard to which, if either, should be given priority.

Finally, let us note that the use of $\top$ enables us to make a remarkable reduction. For as many writers have observed, imperatives can come both with an obligatory and with a permissive force, as when one says 'take a cake' (if the commanding officer says 'take a

[8] There is an extended discussion of this question in Chapter 1 of Mastop (2005).

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

cake', the cadet can appropriately say 'no, thanks', but not if he says 'salute!'). One might have thought that if the imperatives of our system are taken to have obligatory force then we would be unable to capture the permissive use of imperatives. But this is not so. For we can equate the permissive use of the imperative X with $X^{\top} = X \vee \top$. The actions in compliance with $X \vee \top$ will be the actions in compliance with X or the null action $\Box$; and so an alternative to complying with X will be to do nothing. This reduction might also perhaps be of help in explaining how the simple imperative in natural language is capable of being used with permissive force.[9]

§3. The intuitive concept of imperative validity. There would appear to be an intuitive sense in which one imperative may follow from others or from a combination of imperatives and indicatives. Imagine that A is an agent, C an authority who issues commands to the agent, and B an intermediary who attempts to interpret what C says for the benefit of A.[10] Suppose that C says 'it is raining' and 'if it is raining then take an umbrella'. Then it is perfectly appropriate for B to say to A, 'so take an umbrella'. Similarly, if C says 'take an umbrella' and 'put on a rain coat', then it is perfectly appropriate for B to say to A, 'so take an umbrella and put on a rain coat'.

In both cases, B appears, on the face of it, to make an inferential use of the word 'so', one which indicates that the ensuing sentence (the conclusion) does indeed follow from the preceding sentences (the premisses). This suggests that B has performed two inferences; and the fact that it is appropriate for B to say what he does suggests that the inferences are valid, that the imperative conclusion in each case does indeed follow from its premisses. But whether or not there is a genuine inference here, it seems clear that the transition from the one imperative sentence to the other is made on the basis of a logical relation between the two sentences; and so we would like to know what that relation is and which other sentences also stand in the relation. In this way, we can avoid the contentious question of whether the relation is genuinely inferential.

There are some related logical notions which also appear to make good intuitive sense in application to imperatives. Suppose C says 'shut the door or open the window'. Then an equivalent imperative is 'open the window or shut the door'; and nothing would be lost if, in reporting what C had told A to do, B had said the one rather than the other. Or again, suppose C says 'shut the door' and 'leave the door open'. Then in a clear sense, the two imperatives are inconsistent; the one is, as a matter of logic, in 'conflict' with the other.

Moreover, the application of these various logical notions to imperatives is of general importance. For instructions or directives are part of the warp and woof of everyday life—whether they issue from ourselves, or from those who have authority over us, or from the institutions or governing bodies to which we belong. For to understand what to do in the face of some instructions is often a matter of knowing what follows from them, to know whether we have correctly interpreted some instructions is a matter of knowing whether one formulation of them is equivalent to another, and to know whether we can even do what we have been instructed to do is, in part, to know whether the instructions are consistent.

These logical notions also apply to indicative sentences, of course. But we should wary of assuming that their application to imperatives is a straightforward analogue of their

[9] Pragmatic explanations of the permissive use of imperatives in natural language are to be found in §5.1.1 of Kaufmann (2012) and in Portner (2007).

[10] B could, in principle, be identical to C but this raises other issues, which I prefer to put on one side.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

application to indicatives; and we should be wary, in particular, of assuming that the validity of an imperative inference will consist in the preservation of some straightforward imperative analogue of truth or that the logic of such inferences will be a straightforward analogue of classical logic.

In this regard, the 'paradox' of Ross (1941) assumes great importance (and similarly for 'free choice' effects). For the main reason, it seems to me, for accepting the inference from 'shut the door' to 'shut the door or burn the house down' is the validity of the corresponding inference in classical logic. But if we no longer take the correspondence with classical logic for granted, then we are under an obligation to find an account of validity for imperative inference in which such inferences will no longer be valid and in which the correspondence with classical validity can no longer be sustained.

Of course, it would be desirable to arrive at a common conception of validity for indicative and imperative inference and a common set of inference rules. This is not ruled out. But such commonality, if it is to be had, will lie in finding an alternative approach to indicative logic rather than in extending the classical approach to imperative logic.

§4. Validity defined. I wish now to propose an account of imperative validity which I believe is largely in conformity with our intuitions and which also makes good sense of the notion. It is evident from the definition why the notion should be as it is.

The proposal is that the imperative X will *entail* the imperative Y—or, equivalently, that the inference from X to Y will be *valid*—just in case (in any model) (i) any action in compliance with X contains an action in compliance with Y and (ii) every action in compliance with Y is contained in an action in compliance with X. When the inference is from several imperatives $X_1, X_2, \ldots$ to a given imperative Y, the inference is taken to be valid just in case the inference from the conjunction $X_1 \wedge X_2 \wedge \ldots$ to Y is valid in the previous sense, i.e., just in case (i) any fusion of actions in compliance with each of $X_1, X_2, \ldots$ contains an action in compliance with Y and (ii) any action in compliance with Y is contained in the fusion of actions in compliance with $X_1, X_2, \ldots$.[11]

Thus the notion of entailment between imperatives is the exact analogue of the notion of analytic entailment between indicatives; and, indeed, we might with some justice claim that for X to entail Y is for the content of Y to be *part* of the content of X. Entailment, in the case of imperatives, is the articulation of prescriptive content into its parts.

The above account involves two requirements: (i) that X *subsumes* Y; and (ii) that Y *subserves* X. Subsumption and subservience can, in their turn, be explained in terms of partial and entire compliance. An action *a partially complies with* an imperative X if it is contained in an action that exactly complies with X and the action *a entirely complies with* X if it contains an action that exactly complies with X. Then X subsumes Y iff:

> any action that entirely complies with X entirely complies with Y

or, equivalently, iff:

> any action that exactly complies with X entirely complies with Y,

and Y subserves X iff:

---

[11] As before, we might add a further clause to the effect that every action in contravention to Y should be contained in an action in contravention to X, although I am inclined to think that our focus in imperative inference is on compliance rather than contravention.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

any action that partially complies with Y partially complies with X

or, equivalently, iff:

any action that exactly complies with Y partially complies with X.

Thus the validity of the imperative inference X/Y amounts to the preservation of entire compliance from left to right and the preservation of partial compliance from right to left. Or to pur it more intuitively, the imperative inference X/Y will be valid just in case:

(i) any action in compliance with the premiss must go all of the way towards complying with the conclusion and
(ii) any action in compliance with the conclusion must go some of the way towards complying with the premiss.

We might think of the validity of the imperative inference in terms of means-end reasoning. A means here is a *partial constitutive* means, i.e., one that goes some way to constituting the end. The validity of an imperative inference then consists in the conclusion being a necessary means to the premiss. The conclusion is a (constitutive) *means* since any action in compliance with the conclusion is contained in an action in compliance with the premiss (the subservience requirement); and it is a (constitutively) *necessary* means since any action in compliance with the premiss guarantees, via some part, compliance with the conclusion (the subsumption requirement).

We might also think in terms of an imperative as having a double role or meaning (as in Aloni (2007) and Aloni and Ciardelli (2013)). Suppose that $a_1, a_2, \ldots$ are the actions in compliance with X. Then we might regard X as making obligatory one of the actions $a_1, a_2, \ldots$ and as making permissible each of the actions $a_1, a_2, \ldots$. Thus the imperative conveys in this way both an obligation and a permission.

We might then think of the validity of the imperative inference X/Y in terms of the preservation of this double aspect: we need some guarantee that whatever is made obligatory or permissible by X is also made obligatory or permissible by Y. But what does such a guarantee amount to? Suppose that the obligation made by X is satisfied, i.e., that one of the actions in compliance with X is performed. Then what serves to guarantee that one of the actions in compliance with Y is obligatory is that the action a in compliance with X, whatever it might be, should contain an action b in compliance with Y—which is Subsumption. Suppose now that each of the actions in compliance with X is permitted. Then what serves to guarantee that each action in compliance with Y is permitted is that each such action b should be part of an action a permitted by X—which is Subservience. (This understanding of imperatives will later be useful in working out their connection with deontic statements and in developing a preservationist conception of validity).

The above account of validity for imperative inference goes some way towards achieving the desired unity with indicative inference. For there is a notion of partial content for indicative sentences, of independent plausibility, under which the indicative inference S/T is only taken to be valid when the content of T is *part* of the content of S. But this then is the very same as the criterion of validity for imperative inference; the inference X/Y will be valid just in case the content of the conclusion Y is part of the content of the premiss X.

However, a disparity remains. For the notion of classical entailment, essentially the preservation of truth, is of great, perhaps pre-eminent, significance to the logic of indicatives. What corresponds to truth for indicatives is satisfaction for imperatives, where an imperative is satisfied in given circumstances if it is obeyed (or suitably complied with) in those circumstances. However, the corresponding notion of validity for imperatives is of relatively little significance.

What accounts for the disparity, I believe, is a difference in our concerns. We have a great interest in the preservation of truth, in which sentences will be true given that other sentences are true. For we have, among other things, an interest in extending our knowledge—in discovering new truths on the basis of what we already know. But we have no general interest in enlarging the set of imperatives which we may take to be satisfied—in coming up with new imperatives which we can take to be satisfied on the basis of imperatives which have already been satisfied. Our interest is in compliance rather than satisfaction. Given certain imperatives, our concern is not so much with which further imperatives would be satisfied if these imperative were satisfied but with *how* I might comply with these imperatives. The notion of validity of primary interest to us is therefore one that reflects this concern.[12]

Thus I do not think it is altogether correct to say that there is one notion of validity—or, better, entailment—for indicative inference and another for imperative notion. There are two notions in each case, but one is of pre-eminent (though not sole) significance for indicatives and the other of pre-eminent (though perhaps not sole) significance for imperatives.

§5. Other logical notions. As we have noted, there are some other logical notions, besides validity, of interest; and in this section, we shall discuss the sense in which some imperatives might be unsatisfiable (or inconsistent) or in which one imperative might be equivalent to another, and we then relate these two notions to the previous notion of validity. We shall see that the logical landscape looks rather different than from the usual classical perspective.

There are two senses in which a set of imperatives $X_1, X_2, \ldots$ might be said to be consistent or satisfiable. They will be *wholly satisfiable* if every action in compliance with their conjunction is possible and *partially satisfiable* if some action is compliance with their conjunction is possible. We may say, correspondingly, that the imperatives X_1, $X_2, \ldots$ are *wholly unsatisfiable* if they are not partially satisfiable, so that no possible action is in compliance with their conjunction, and that they are *partially unsatisfiable* if they are not wholly satisfiable, so that some impossible action is in compliance with their conjunction.[13]

So, for example, the imperative X = 'shut the door and do not shut the door' (of the form $\alpha \wedge \neg\alpha$) will be wholly unsatisfiable, since any action in compliance with X will be the fusion of an action a in compliance with the imperative and an incompatible action a' in contravention to X. On the other hand, the imperative X = 'either shut the door and do not shut the door or open the window' (of the form $(\alpha \wedge \neg\alpha) \vee \beta$) will be partially unsatisfiable though not wholly unsatisfiable, since an impossible action in compliance with the first

[12] The above distinction corresponds to Ross' distinction, discussed in Segerberg (1989), between the logic of satisfaction and the logic of validity.

[13] These notions of satisfiability and unsatisfiability are relative to a model but we might lift the relativity by taking satisfiability or unsatisfiability simpliciter, whether weak or strong, to be satisfiability or unsatisfiability in every model (or every classical model).

We might also take the imperatives $X_1, X_2, \ldots$ to be be unsatisfiable if there is no action whatever in compliance with their conjunction $X_1 \wedge X_2 \wedge \ldots$, although this notion of unsatisfiability will only have application if we allow there to be imperatives which allow no actions.

disjunct will be in compliance with the disjunction while a (presumably possible action) in compliance with the second disjunct will also be in compliance with the disjunction.

We might see the notion of partial satisfiability as arising from the obligatory aspect of imperatives. If the obligation to do one of the actions in conformity with the imperative is to be met, then one of those actions must be possible. Similarly, we might see the notion of whole satisfiability as arising from the permissive aspect of imperatives. If the permission to do each of the actions in compliance with the imperative is to be effective, then each of the actions must be possible.

The distinction is relevant to a puzzle discussed by Hare (1967) and Williams (1963) (and later discussed by Veltman (2009)). For there seems to be some kind of contradiction involved in the imperative 'Help your mother or help your father, but don't help your mother' ($(\alpha \vee \beta) \wedge \neg\alpha$). But it is not the same kind of straight flat out contradiction that is involved in 'Help your mother and don't help your mother' ($\alpha \wedge \neg\alpha$), since it is possible to comply with the first imperative (by helping your father) but not with the second.

The difference, for us, is simply the difference in the two kinds of unsatisfiability, with there being a satisfiable permission in the case of the second imperative though not the first and an unsatisfiable permission in the case of both imperatives (cf. Gombay (1965)). But we should note that it is only by allowing impossible actions to be in compliance with an imperative that we are able to draw the distinction between the two kinds of contradiction. A more standard intensional approach, in which actions are taken to be nonempty sets of worlds, is not able to explain the difference in this way or, perhaps, in any way at all.

We turn to equivalence. We might say that the imperatives X and Y are *exactly equivalent in* a model if the same actions are in compliance with both within the model and that they are *exactly logically equivalent* if they are exactly equivalent in each model (a somewhat stronger definition would also require the same actions to be in contravention to X and Y).

We can provide an account of exact entailment in terms of exact equivalence, where X exactly entails Y if every action in (exact) compliance with X is in (exact) compliance with Y. For X will exactly entail Y just in case Y is exactly equivalent to $X \vee Y$.[14] We can also provide a definition of being wholly unsatisfiable in terms of exact equivalence. For let $*$ be the thoroughly inconsistent imperative which allows all and only impossible actions. Then the imperative X will be wholly unsatisfiable just in case $*$ is exactly equivalent to $X \vee *$.

There is an alternative account of equivalence as two-way analytic entailment. The imperatives X and Y will be equivalent in this sense just in case every action in compliance with X contains and is contained in an action in compliance with Y and if every action in compliance with Y contains and is contained in an action in compliance with X.

Perhaps somewhat surprisingly, the two definitions are not coextensive. For consider the imperatives $X = \alpha \vee (\alpha \wedge \beta \wedge \gamma)$ and $Y = \alpha \vee (\alpha \wedge \beta) \vee (\alpha \wedge \beta \wedge \gamma)$—as in 'Press button A or buttons A, B, and C' and 'Press button A or buttons A and B or buttons A, B, and C'. Then even though each analytically entails the other, according to our definition, they do not allow the same actions. For the action of pressing buttons A and B will be in exact compliance with the second but not the first.

Some other standard connections between the different logical notions will also not hold. For example, the pair of imperatives $X = \alpha$ and $Y = \beta \wedge \neg\beta$ will be unsatisfiable (either wholly or partially). However, X does not entail $\neg Y$ since an action in compliance with or

14 In this respect, there is a difference between exact entailment and analytic entailment since, as we shall see, there is not an analogous definition of analytic entailment in terms of exact equivalence and conjunction.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

contravention to β and hence in compliance with ¬Y need not be contained in an action in compliance with X.

Or again, it does not follow from the fact that X analytically entails Y that X is analytically equivalent to X ∧ Y. For X entails X. However, Idempotence fails, X is not in general analytically equivalent to X ∧ X. For suppose X allows both b and c. Then their fusion will be in compliance with X ∧ X but not, in general, with X itself. To take an actual example 'Shut the door or open the window, and shut the door or open the window' will allow one to both shut the door and open the window, whereas plain 'Shut the door or open the window' will not. But this means that X will entail X even though X is not equivalent to X ∧ X.

We may obtain a more satisfactory account of the relationship between analytic entailment and equivalence by modifying the notion of conjunction. Suppose we are given two imperatives X and Y whose respective contents are $X = \{a_1, a_2, \dots\}$ and $Y = \{b_1\ b_2, \dots\}$. Then we may say that the imperative Z is a *coordinated conjunction* of X and Y if there is a relation R whose domain is X and whose range is Y and which is such that the content Z of Z is the set of all those fusions $a \sqcup b$ for which a stands in the relation R to b. Thus coordinated conjunction only requires pairing every member of X with a member of Y and every member of Y with a member of X, whereas ordinary conjunction requires that every member of X be paired with every member of Y (the notion of coordinated conjunction is further discussed in Part I of Fine (2017)).

It is now easy to show that X will entail Y just in case X is a *coordinated* conjunction of Y and X—or of Y and any other imperative Z, for that matter. We might think of the additional conjunct Z as telling us how to go on once we have complied with Y, where how we go on is a function of how we began (with a particular member of Z being paired with particular members of Y). Thus X will entail Y if there is a way to go on, given compliance with Y, that will then constitute compliance with X.

The failure of Idempotence is also relevant to the attempt to account for entailment in terms of updating. One familiar way to do this is as follows (Veltman, 1996, p. 224; Mastop, 2005). We are given a context c, which may be updated by a sentence X, to give an updated context c[X]. X is then taken to entail Y just in case c[X][Y] = c[X] for any context c; updating with Y, after updating with X, gives nothing new.

Within the present setting, we may naturally take a context c to be given by a (unilateral) content C. Updating is then a form of conjunction: where X is the unilateral content of X, c[X] $= C \wedge X$. But letting C be the 'null' context {□}, c[X] $= X$ and c[X][X] $= X \wedge X$ and, since X is not in general the same as $X \wedge X$, the account will not deliver the result that X entails X. To take an actual example, if I update with 'take an apple or a pear' and then update again with 'take an apple or a pear', then I am offered a choice which I did not have before, which is to take both an apple and a pear.

In fact, the problem is much worse than this example suggests, for it may be shown that *no* inference from $X_1, X_2, \dots, X_n$ to Y will be valid in the proposed sense as long as each of the formulas $X_1, X_2, \dots, X_n$, Y is free of ⊤. Thus in the present context, the standard definition of validity in terms of updating is not acceptable.

We could avoid these various anomalies by closing the content of an imperative under two conditions: Fusion, according to which the fusion of any nonempty set of actions in compliance with an imperative is also in compliance with the imperative; and Convexity, according to which any action which lies between two actions in compliance with an imperative is also in compliance with the imperative. There are contexts in which these conditions (or their alethic counterparts) can reasonably be imposed. If our interest is in

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

partial truth or verisimilitude of a theory, for example, then nothing is perhaps lost by supposing that its content is closed under both of these conditions.

However, in the present context, Fusion is unacceptable. We do not want the imperative ‘press button A or press button B’, for example, to allow one to press both buttons, as would be the case with closure under Fusion. Moreover, if we start off with an inclusive concept of disjunction, under which $\alpha \vee \beta$ is equivalent to $\alpha \vee \beta \vee (\alpha \wedge \beta)$, then it will not be possible to define a noninclusive concept, under which the equivalence fails, in terms of it. But we may, of course, define the inclusive notion in terms of the noninclusive notion as $\alpha \vee \beta \vee (\alpha \wedge \beta)$. Thus considerations of expressivity require that we adopt a noninclusive treatment of disjunction (and of other such connectives).

Convexity is perhaps also unacceptable, although the case is less clear. For the imperative ‘press button A or press buttons A, B, and C’ does allow one, in a sense, to press buttons A and B, even though the pressing of A and B is not in *exact* compliance with the imperative. Thus if we close the content of an imperative under convexity, it will no longer be clear which actions are in exact compliance with the imperative although it will still be clear which actions are in partial compliance and also which actions are in partial compliance while still containing an action in full compliance with the imperative (and hence go ‘far enough’ without going ‘too far’).

It might be thought that the failure of exact equivalence to coincide with two-way analytic entailment points to a shortcoming in our notion of entailment. For our interest in drawing imperative consequences from some imperative premisses is not only in determining the necessary constitutive means by which we might comply with the premiss but also in determining how we might be in exact compliance with the premisses. It might have been thought that this further requirement simply amounts to the conclusion entailing the conjunction of the premisses. But, as is shown by the button pressing example, this is not so; and so we seem to lack the means of expressing this further requirement.

However, the situation is not as dire as it might seem. For the two-way analytic entailment of X and Y *will* be a necessary and sufficient condition for exact equivalence, as long as the contents of X and Y are convex. Thus we may use the criterion in this special case. This appears to be a severe restriction. But it seems to me that whenever we are prepared to state an imperative X, we should also be prepared to state an imperative whose content is convex. Consider the button case, for example, in which the imperative is ‘Press button A or press buttons A, B, and C’. We may now ask if we are willing to allow the pressing of buttons A and B to be a compliant action. If we are, then we should be willing to state ‘Press button A or buttons A and B or buttons A and B and C’. If we are not, then we should be willing to state ‘Press button A without B or press buttons A, B, and C’. By continually expanding the imperative in this way, we will eventually end up with an imperative whose content is convex; and instead of applying the criterion to the original imperatives we may apply it to their expansions.

All the same, we may be interested in the notion of exact equivalence in cases in which the content of the imperatives in question is not convex. An independent treatment of exact equivalence is then called for and I have for this reason presented two sets of axiomatic rules, one for exact equivalence and the other for analytic entailment.

Let us note, finally, that we might introduce modal counterparts of the previous mereological notions of validity. Thus we might say that the content X *modally subsumes* the content Y if each action a in X necessitates an action b in Y, i.e., if is impossible that a is performed and b not performed (or alternatively, to adopt a somewhat weaker condition, if X necessitates Y in the sense that it is impossible that an action in X is performed while

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

none of the actions in Y are performed); and, similarly, we might say that the content Y *modally subserves* the content X if each action a in Y is necessitated by an action b in X. We can then define the imperative inference X/Y to be (*modally*) *valid* if X modally subsumes Y and Y modally subserves X.

The weak version of the modal criterion will lead to different results from the previous mereological criterion. Thus an inference of the form $\alpha/\alpha \vee (\alpha \wedge \beta)$ (as in 'post the letter, so post the letter or post the letter and burn down the building') will be valid under the weak modal criterion but not under the mereological criterion. The strong version of the modal criterion will, of course, lead to the same results as the mereological criterion as long as the relation of whole to part coincides with the relation of necessitation but otherwise they may differ. Thus the imperative 'Square the circle' will entail 'Find a counterexample to Fermat's Last Theorem' under the modal criterion but not under the mereological criterion.

§6. Some systems of imperative logic. We shall present some sound and complete axiom systems for equivalential and implication versions of imperative logic, with or without the constant $\top$. Proofs of soundness and completeness are omitted, though they can be established through a variant of the method of disjunctive normal forms developed in Fine (2016a). The basic idea is to reduce each formula to a normal form and then provide a 'canonical' semantics in which the states are taken to be sets of literals. However, in the present case we need to take "copies" of literals in order to be able to distinguish, for example, between the contents of α and $\alpha \wedge \alpha$.

Take the sequent X ⇔ Y to be *valid* if $X = Y$ in any model. The following axioms and rules then determine the valid $\top$-free sequents of the form X ⇔ Y:

E1 $X \Leftrightarrow \neg\neg X$
E2 $X \wedge Y \Leftrightarrow Y \wedge X$
E3 $(X \wedge Y) \wedge Z \Leftrightarrow X \wedge (Y \wedge Z)$
E4 $X \Leftrightarrow X \vee X$
E5 $X \vee Y \Leftrightarrow Y \vee X$
E6 $(X \vee Y) \vee Z \Leftrightarrow X \vee (Y \vee Z)$
E7 $\neg(X \wedge Y) \Leftrightarrow (\neg X \vee \neg Y)$
E8 $\neg(X \vee Y) \Leftrightarrow (\neg X \wedge \neg Y)$
E9 $X \wedge (Y \vee Z) \Leftrightarrow (X \wedge Y) \vee (X \wedge Z)$
E10 $X \Leftrightarrow Y \,/\, Y \Leftrightarrow X$
E11 $X \Leftrightarrow Y, Y \Leftrightarrow Z \,/\, X \Leftrightarrow Z$
E12 $X \Leftrightarrow Y \,/\, X \wedge Z \Leftrightarrow Y \wedge Z$
E13 $X \Leftrightarrow Y \,/\, X \vee Z \Leftrightarrow Y \vee Z$.

If the verum constant $\top$ is added to the language, we then require the addition of the following two axioms:

E14 $\top \wedge X \Leftrightarrow X$
E15 $\bot \wedge X \Leftrightarrow \bot$.

We might call the resulting logic the *imperative logic* IL(⇔) *of exact equivalence*.

Let us now take the sequent X >Y to be *valid* if X analytically entails Y in any model. The following axioms and rules then determine the valid $\top$-free formulas of the form X >Y (where $X \approx Y$ is used to indicate the pair of sequents X >Y and Y >X):

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

A1 $X \approx \neg\neg X$
A2 $X \wedge Y > Y \wedge X$
A3 $(X \wedge Y) \wedge Z > X \wedge (Y \wedge Z)$
A4 $X \approx X \vee X$
A5 $X \vee Y > Y \vee X$
A6 $(X \vee Y) \vee Z \approx X \vee (Y \vee Z)$
A7 $\neg(X \wedge Y) \approx (\neg X \vee \neg Y)$
A8 $\neg(X \vee Y) \approx (\neg X \wedge \neg Y)$
A9 $X \wedge (Y \vee Z) \approx (X \wedge Y) \vee (X \wedge Z)$
A10 $X \wedge Y > X$
A11 $X > Y, Y > Z / X > Z$
A12 $X > Y / X \wedge Z > Y \wedge Z$
A13 $X > Y / X \vee Z > Y \vee Z$.

If the verum constant $\top$ is added to the language, we then require the addition of the following two axioms:

A14 $X > \top$
A15 $\bot > X$.

According to the first, every imperative entails the null imperative and, according to the second, the full imperative entails every imperative. We might call the resulting logic the *imperative logic* IL(>) *of analytic entailment*.

We have presented two separate systems of imperative logic, IL($\Leftrightarrow$) for exact equivalence and IL(>) for analytic entailment. We might combine the two systems into a single system IL($\Leftrightarrow$, >) which countenances both kinds of sequent. In this case, the axioms and rules can be greatly simplified. Leaving the constant $\top$ out of the picture, we might adopt the axioms and rules, E1–E13, from the system IL(>, $\Leftrightarrow$), A10, A11 and A13 from the system IL(>), and the following 'mixed' rule:

$A \Leftrightarrow B / A < B$.

There would be some interest in developing systems for some of the other logical notions that have been considered and for considering the different systems which would result from imposing further semantic constraints. If we required, for example, that there should be one impossible action (as in a topsy space), then we would also want

$X, \neg X > Y$

to be an axiom.

§7. Some special inferences.

I should like to discuss some cases of valid and invalid imperative inference of special logical or historical interest.

7.1. Ross. The inference from α to $\alpha \vee \beta$ (as with 'post the letter, so post the letter or burn the house down') is not valid under the present semantics, since an action in compliance with β may not be contained in an action in compliance with α. Compliance with the conclusion is *necessary* for complying with the premiss but not a necessary *means* for complying with the premiss.

In this regard, there is a striking connection with the notion of analytic entailment. Angell (1977) and Parry (1933) were concerned to block the argument from $p \wedge \neg p$ to q by blocking the inference from p to $p \vee q$. The definition of analytic entailment appropriate

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

for Angell's system (Fine, 2016a) turns out to be a slight variant of the current notion of entailment for imperatives; and we thereby effect a remarkable synthesis of two distinct traditions within which the principle of disjunctive weakening has been challenged.

We should note that just as the inference from α to $\alpha \vee \beta$ is invalid according to the semantics, so is the inference from $\neg\alpha$ to $\neg(\alpha \wedge \beta)$, given the equivalence between $\neg(\alpha \wedge \beta)$ and $(\neg\alpha \vee \neg\beta)$. This seems very reasonable. For from 'Do not administer the poison!' we do not want to infer 'Do not administer the poison and the antidote'—to forbid the one is not to forbid the other. And it might be regarded as a strong point in favor of the present semantics, in contrast to the more traditional forms of alternative semantics, that it is able to give a parallel treatment of the two cases.

However, for reasons I do not understand, the intuitive data seem somewhat murky. For even though we have no inclination to accept the inference from 'Do not press button A' to 'Do not press button A or do not press button B', we have some inclination to accept the inference from 'Do not press button A' to 'Do not press button A and button B' (contrary to the case above). To make matters worse, we have a strong inclination to accept the inference from 'Do not press button A and button B' to 'Do not press button A or do not press button B', thereby putting our various inclinations in conflict with one another.

A related issue arises in connection with Hare's puzzle. As we have mentioned, imperative premisses of the form $\neg\alpha$ and $\alpha \vee \beta$ (as with 'Don't help your father', 'Help your father or mother') somehow seem contradictory. But premisses of the form α and $\neg(\alpha \wedge \beta)$ (as with 'Press button A', 'Don't press button A and B') do not. And yet $\neg(\alpha \wedge \beta)$ appears to entail $\neg\alpha \vee \neg\beta$ ('Do not press button A and B' entails 'Do not press button A or do not press button B').

7.2. Disjunctive syllogism. The inference from $\neg\alpha$ and $\alpha \vee \beta$ to β (as with 'do not shut the door, shut the door or open the window, so open the window') is not valid under the present semantics, since the fusion of an action a' in compliance with $\neg\alpha$ and of an action a in compliance with α, and hence in compliance with the premisses, may not contain an action in compliance with β.

It is rather hard to say whether the inference is intuitively valid. The premisses are inconsistent in a certain way and our intuitive judgements as to what follows in the face of such inconsistency are somewhat infirm. The inference may be rendered valid if we restrict our attention to topsy models, since then the impossible action which results from fusing a' and a will contain any action in compliance with β. However, under this restriction, the inference from $\neg\alpha$ and $\alpha \vee\beta$ to $\beta \vee \gamma$ (as with 'do not shut the door, shut the door or open the window, so open the window or burn the house down') will be valid. This is not really an instance of the Ross paradox, since it is an inference in which the premisses are already unacceptable. It is perhaps problematic in the same way as the inference from $\neg\alpha$ and α to β and can, with as much justice, be either rejected or retained.

An alternative proposal is to apply the 'consistency filter'. The *filtered content of* an imperative X will be the subset of the possible actions that it allows (and, similarly, for disallowed actions). The notion of entailment might then be restricted to the filtered content. Thus, with filtering in place, X will *entail* Y if either (a) no possible action is in compliance with X or (b)(i) every possible action in compliance with X contains an action (and hence a possible action) in compliance with Y and (ii) every possible action in compliance with Y is contained in a possible action in compliance with X.

The inference from $\neg\alpha$ and $\alpha \vee \beta$ to β will then be valid while the inference from $\neg\alpha$ and $\alpha \vee \beta$ to $\beta \vee \gamma$ will not be valid. However, this solution suffers from problems of its own. The inference from $(\alpha \wedge \neg\alpha) \vee \gamma$ to $\alpha \vee \gamma$, for example, will no longer be valid

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

(since a possible action in compliance with α will not be part of a possible action in compliance with $\alpha \wedge \neg\alpha$, given that no possible action is in compliance with $\alpha \wedge \neg\alpha$).

The inference from α to $\alpha \vee \bot$ will be valid under the operation of the consistency filter though not valid under the restriction to topsy models, since the action ■ in compliance $\bot$ may not be contained in any action in compliance with α. Thus the two solutions are incomparable in their results. A third solution, more comprehensive than the two others, replaces condition (a) in definition of entailment above with the weaker condition that the premiss X be partially unsatisfiable, i.e., that some impossible action be in compliance with X. Thus any kind of 'contamination' in the premisses will lead to an inconsistency of the worst kind, one in which everything then follows. This weaker condition might be justified on the grounds that the actions in compliance with an imperative should be permitted and no impossible action can be permitted.

Under this proposal, the inference from $\neg\alpha$ and $\alpha \vee \beta$ to *any* conclusion γ will be valid. Thus on this proposal, in contrast to all the other definitions of validity we have so far considered, an inference may be valid without being classically valid. Despite this oddity, there is, I believe, a great deal to be said in its favor, especially if one wishes to adopt a liberal conception of inference under which any conclusion whatever will follow from inconsistent premisses.

I hope it is clear from this brief discussion of the logic that it is more or less in conformity with our intuitions and that it has the flexibility to accommodate different ways in which they might be systematized.

§8. Antecedents and comparisons. I wish briefly to discuss the related work of some other authors. The work may relate either to the underlying framework, the semantics or the definition of validity and, in each case, may relate principally to indicatives or specifically to imperatives. Unfortunately, a comprehensive and detailed comparison is beyond the scope of the article.

<u>The framework</u>. I presuppose a state space in the case of indicatives and an action space in the case of imperatives. An action, as embedded in an action space, might be compared to a 'to-do' list, as found in the work of Portner (2004), Mastop (2005), and many others. But my notion of an action is, in some ways, less informative and, in other ways, more informative than the notion of a to-do list. A to-do list is a *set* of actions construed conjunctively as it were; the idea is that one do everything on the list. But in place of the set of actions $a_1, a_2, \ldots$, I have a single action $a = a_1 \sqcup a_2 \sqcup \ldots$, which is their fusion. Of course, the individual actions $a_1, a_2, \ldots$ can be recovered from the single action a, since their fusion is a. But so can many other individual actions, such a_1 and $a_2 \sqcup a_3 \sqcup \ldots$ or $a_1 \sqcup a_2$ and $a_2 \sqcup a_3 \sqcup \ldots$ or $a_1 \sqcup a_2, a_3 \sqcup a_4$, and $a_5 \sqcup a_6 \sqcup \ldots$. The correspondence between the to-do list and the single action is many-one. If there were atomic actions of which every action was a fusion, then it would not be implausible to suppose that each action a was the fusion of a unique set of atomic actions; and the correspondence between sets of atomic actions and single actions would then be one-one. But there is no reason in general to suppose that there are such atomic actions. After all, one might think that any action always takes some time and that, consequently, any action always contained a shorter action as a proper part.

It is not clear to me that there is any real point in appealing to a set of actions as opposed to their fusion. If actions are uniquely decomposable into atomic actions, then the underlying atomic actions can be recovered from their fusion. If actions are not uniquely decomposable into atomic actions, then it is not clear what else might sensibly be said about how the individual actions relate to one another which is not already implicit in the mereological structure of the fusion. It has been proposed, for example, that one might recover a preference relation on worlds from a to-do list by taking one world to be better than another if it realizes more actions on the to-list. But this would make the preference relation unduly dependent upon how the actions were mereologically related. One world might be better in regard to the actions a_1, a_2, a_3, for example, though not in regard a_1, $a_2 \sqcup a_3$. Such a criterion only has application *at best* to the case in which the actions are atomic since then one might presuppose that each of the actions should be given equal weight—though even this, to my mind, is extremely problematic.

It therefore seems to me that it is something like an action, long with its mereological relationships to other actions, that best fulfils the role that has been played by a to-do list.

Semantics. Van Fraassen (1969), as far as I am aware, was the first published statement of a truthmaker semantics for propositional logic (references to the subsequent history can be found in Fine (2017)). The truthmaker approach bears some resemblance to inquisitive semantics (as in Ciardelli *et al.*, (2013)) and 'alternative'-based semantics (as in Alonso-Ovalle (2006) and Aloni (2007)), for all three approaches take seriously the idea that there may be relevantly different ways, within a given possible world, in which a sentence may be true.

We find something like the application of truthmaker semantics to imperative or deontic sentences in Stelzner (1992), van Rooij (2000), Mastop (2005), and Aloni & Ciardelli (2013) (although in Mastop (2005), the semantics is presented as a form of update semantics in the sense of Veltman (1996)). But in all of these cases, going all the way back to van Fraassen (1969), one finds that the authors adopt what is, at best, a very particular realization of the truthmaker framework. For the 'states', or 'actions' are either identified with sets of possible worlds or the like (resulting in an 'intensional' state space) or with sets of literals or the like (resulting in a 'canonical' state space). In the first case, this means that one cannot draw a semantic distinction between 'square the circle', for example, and 'show that $2 = 3$'; in the second case, it means that there is not a proper separation between the language and what it conveys; and in both cases, it means that there is an unduly restricted conception of the different ways in which an indicative might be made be true or an imperative might be obeyed.

One also finds that these authors adopt a conception of a model, or what might be taken to be a model, in which the sentential atoms receive a special interpretation. Thus under a 'canonical' state space, the atom p might be assigned $<\{p\}, \{\neg p\}>$ as its semantic value. We ourselves will consider an interpretation of this sort in part II. However, such an interpretation is not entirely unproblematic. For even if the sentential atoms are meant to stand in for the atomic sentences of some 'host' language, there is no reason, in general to suppose that the atomic sentences of the host language will have a single truth-maker and a single falsity-maker. And, of course, such an approach is entirely unsuited to an interpretation in which the sentential atoms are allowed to stand in for any sentence whatever.

In some of this other work, there are also differences in the form of clause adopted for some of the connectives and, in particular, many of the accounts do not provide a 'dual' treatment of disjunction and conjunction under which the De Morgan Laws will hold. I might mention, in particular, the treatment of negation in Aloni & Ciardelli (2013). Suppose that $a_1, a_2, \ldots$ are the actions in compliance with the imperative X, which we may identify with subsets of the set W of possible worlds. Aloni and Ciardelli then take there to be a single action, $W - (a_1 \cup a_2 \cup \ldots)$, in compliance with ¬X. So the single action in compliance ¬¬X will be $W - (W - (a_1 \cup a_2 \cup \dots)) = (a_1 \cup a_2 \cup \dots)$. The double negation ¬¬X, in effect, flattens the meaning of X, substituting a single disjunctive action $(a_1 \cup a_2 \cup \dots)$ (the performance of one of $a_1, a_2, \dots$) in place of the individual actions $a_1, a_2, \dots$.

Quite apart from questions of empirical adequacy, this move, to my mind, destroys one of the principal rationales for adopting the truthmaker approach. For if we think of an action as an arbitrary set of possible worlds, then imperative or deontic sentences can longer serve as a guide to action. To use an example discussed at length in Fine (2014a), suppose God tells Eve: eat infinitely many of the apples from Alternative Eden (E). Then E is logically equivalent to E′: eat infinitely many apples from Eden or Alternative Eden. So is she allowed to eat the Forbidden Fruit? There is no way of telling.

This problem disappears if the actions taken to be in compliance with E are taken to be different from our actions in compliance with E′, as they are under the truthmaker approach, where the single logical equivalent is allowed to resolve itself, so to speak, into different ranges of choice.

But the problem reappears once we are allowed to flatten the meaning of the imperative. For suppose we start with E and form its double negation ¬¬E. Then just as we were unable to think of E as providing a guide to action under the standard intensional approach, we are now unable to think of ¬¬E as providing a guide to action under the present approach. As before, we have no idea what the alternative actions by which it might be satisfied should be.

Validity. As I have mentioned, the present notion of validity has application to Angell's system of analytic implication and some relevant citations are given in Fine (2016a). Related notions of validity, in application to the logic of imperatives, can be found in Stelzner (1992), van Rooij (2000), Aloni (2007), and Aloni & Ciardelli (2013). However, in all of these cases there are a number of discrepancies from the definition given here, quite apart from general issues of presentation. So, for example, Stelzner's analogues of the subsumption and subservience conditions are somewhat weaker than our own; van Rooij's conditions are relativized to a world; Aloni presents a modal form of the definition; and Aloni and Ciardelli convert the alternative actions presented by an imperative into exclusive options.

Some authors (such as Mastop (2005, 2012) and Starr (2018)) who adopt something akin to the present semantics also adopt the "dynamical" definition of validity according to which Y will follow from X if in any context c, c[X][Y] = c[X].[15] They do not necessarily face the disastrous consequence we noted above that no inference will then be valid. For they often insist that the interpretation of the atoms should be definite, that there should be only one action in compliance with or in contravention to an atomic imperative. Any atomic imperative α will then be a consequence of itself. However, they face an almost equally disastrous consequence. For say that the inference from $X_1, X_2, \dots, X_n$ to Y is *schematically valid* if each substitution-instance of the inference is valid. Then

[15] Strictly speaking, Mastop (2005) initially adopts Veltman's account (Definition 4.2 of §4.1) and then goes on to give a different account (Definition 4.9 of §4.2), under which acceptance only requires that the initial context σ be a *subset* of the updated context $\sigma[\varphi]$. Of course, his new definition of validity imports a new idea (via the subset relation) into the account of validity and is itself quite dubious as an account of acceptance, since the updated context may permit actions not already permitted in the initial context.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

no ($\top$-free) inference will be schematically valid; and, in particular, many instances of the inference from X to X (such as the inference from $\alpha \vee \beta$ to $\alpha \vee \beta$) will fail to be valid.[16]

Finally, it should be mentioned that something like the second component condition of subservience is sometimes to be found on its own as a criterion of validity. Thus the notion of p-entailment from Kamp (1973) loosely corresponds to subservience (though Kamp only has statements of permission in mind, not imperatives in general). Or again, Kenny (1966) and Geach (1966), following him, takes imperative to involve reasoning from means (as stated in the conclusion) to ends as (stated in the premisses). However, the means for Kenny are logically sufficient conditions whereas, for us, they are partially constitutive conditions.

§9. Appendix. Let me present a brief formal exposition of the proposed semantics. I give definitions only and do not state results.

An *imperative formula* is one constructed in the usual way from the imperative atoms $\alpha_1, \alpha_2, \ldots$ and the constant $\top$ by means of the connectives, $\neg$, $\wedge$ and $\vee$. We use α, β, γ and the like for arbitrary imperative atoms and X, Y, Z and the like for arbitrary imperative formulas.

An *action space* $\boldsymbol{A}$ is a structure of the form $(A, \sqsubseteq)$, where A (actions) is a set and $\sqsubseteq$ (part-whole) is a partial order on A for which each subset B of A has a least upper bound $\bigsqcup B$. We use obvious notation in connection with an action space and, in particular, use $\square = \bigsqcup \emptyset$ for the *null* action and $\blacksquare = \bigsqcup A$ for the *full* action.

A *modalized action space* **A**—or M-*space*, for short—is a structure of the form $(A, A^{\Diamond}, \sqsubseteq)$, where $(A, \sqsubseteq)$ is an action space and $A^{\Diamond}$ (possible actions) is a nonempty subset of A subject to the conditions that $b \in A^{\Diamond}$ whenever $a \in A^{\Diamond}$ and $b \sqsubseteq a$ and that $\blacksquare \notin A^{\Diamond}$. We say that an action a in an M-space is *consistent* or *possible* if $a \in A^{\Diamond}$ and *inconsistent* or *impossible* otherwise. A subset B of actions is said to be *compatible* if their fusion belongs to $A^{\Diamond}$ and to be *incompatible* otherwise.

An (*action space*) *model* $\boldsymbol{M}$ is an ordered triple $(A, \sqsubseteq, |\cdot|)$, where $(A, \sqsubseteq)$ is an action space and $|\cdot|$(valuation) is a function taking each atom α into a pair (C, C') of nonempty subsets of A. When $|\alpha| = (C, C')$, we let $|\alpha|^+$ or $|\alpha|$ be C (the set of actions in compliance with α) and let $|\alpha|^-$ be C'(the set of actions in contravention to α).

A *classical model* $\boldsymbol{M}$ is an ordered quadruple $(A, A^{\Diamond}, \sqsubseteq, |\cdot|)$, where $(A, A^{\Diamond}, \sqsubseteq)$ is an M-model and, for each atom α, (i) no member of $|\alpha|^+$ is compatible with a member of $|\alpha|^-$ and (ii) each possible action is compatible with a member of $|\alpha|^+$ or of $|\alpha|^-$.

Given a model $\boldsymbol{M} = (A, \sqsubseteq, |\cdot|)$, an action $a \in A$ and an imperative X, we may give an inductive definition of what it is for a to *comply with* X ($a \Vdash$ X) or to *contravene* X ($a \dashv\vdash$ X):

- (i)$^+$ $a \Vdash \alpha$ if $a \in |\alpha|^+$;
- (i)$^-$ $a \dashv\Vert \alpha$ if $a \in |\alpha|^-$;
- (ii)$^+$ $a \Vdash \top$ if $a = \square$;
- (ii)$^-$ $a \dashv\Vert \top$ if $a = \blacksquare$;

16 Idea behind proof: for suitably large m, replace each sentential atom in $X_1, X_2, \ldots, X_n$, and Y with $(\alpha_1 \vee \alpha_2 \vee \cdots \vee \alpha_m) \vee (\beta_1 \wedge \beta_2 \wedge \cdots \wedge \beta_m)$. If one only adopts the positive truthmaker clauses for $\wedge$ and $\vee$, the result will still hold for all positive formulas $X_1, X_2, \ldots, X_n$, and Y constructed from $\wedge$ and $\vee$.

https://doi.org/10.1017/S175502031700020X Published online by Cambridge University Press

(ii)$^+$ a ||-¬X if a -|| X;
(ii)$^-$ a -||¬X if a ||- X;
(iii)$^+$ a ||-X ∧ Y if for some b and c, b ||- X, c ||- Y and $a = b \sqcup c$;
(iii)$^-$ a -||X ∧ Y if a -|| X or a -|| Y;
(iv)$^+$ a ||-X ∨ Y if a ||- X or a ||- Y;
(iv)$^-$ a -||X ∨ Y if for some b and c, b -|| X, c -|| Y and $a = b \sqcup c$.

Relative to a model $\boldsymbol{M} = (A, \sqsubseteq, |\cdot|)$, we let $|X|^+ = \{a \in A$: a ||- X$\}$ and $|X|^- = \{a \in A$: s -|| X$\}$; and we may write $|X|^+$ as X. For an M-model $\boldsymbol{M} = (A, A^{\diamond}, \sqsubseteq)$, we define the *filtered content* $|X|^{\diamond}$ *of* X (in terms of which a modified account of validity might be given) as $\{a \in A^{\diamond}$: a ||- X$\}$.

Given an action space $\boldsymbol{A} = (A, \sqsubseteq)$ and subsets X and Y of A, we say:

(i) X *subsumes* Y if for each $a \in X$ there is a $b \in Y$ for which $a \sqsupseteq b$;
(ii) Y *subserves* X if for each $b \in Y$ there is an $a \in X$ for which $a \sqsupseteq b$;
(iii) X *analytically entails* Y if X subsumes Y and Y subserves X.

More generally, given a well-ordered sequence $X_1, X_2, \ldots$ of subsets of A, set $X_1 \sqcup X_2 \sqcup \ldots = \{a_1 \sqcup a_2 \sqcup \ldots : a_1 \in X_1, a_2 \in X_2 \ldots\}$. We then say that the imperative formulas $X_1, X_2, \ldots$ analytically *entail* Y if $X_1 \sqcup X_2 \sqcup \ldots$ entails Y.

BIBLIOGRAPHY

Aloni, M. D. (2007). Free choice, modals, and imperatives. *Natural Language Semantics*, **15**, 65–94.

Aloni, M. D. & Ciardelli, I. A. (2013). A logical account of free choice imperatives. In Aloni, M. D. and Franke, M., editors. *The Dynamic, Inquisitive, and Visionary Life of φ, ?φ, and ◇φ – A Festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman.* Available at http://festschriften.illc.uva.nl/Festschrift-JMF/.

Alonso-Ovalle, L. (2006). Disjunction in Alternative Semantics. Doctoral Dissertation, Amherst: University of Massachusetts.

Angell, R. B. (1977). Three systems of first degree entailment. *Journal of Symbolic Logic*, **47**, 147.

Chellas, B. F. (1971). Imperatives. *Theoria*, **37**, 114–129.

Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2013). Inquisitive semantics: A new notion of meaning. *Language and Linguistics Compass*, **7**, 459–476.

Fine, K. (2014a). Permission and possible worlds. *dialectica*, **68**, 317–336.

Fine, K. (2014b). Truthmaker semantics for intuitionistic logic. *Journal of Philosophical Logic*, **43.2**, 549–577. Reprinted in Philosophers' Annual for 2014.

Fine, K. (2016a). Angellic content. *Journal of Philosophical Logic*, **45**(2), 199–226. Published online https://doi.org/10.1007/s10992-015-9371-9.

Fine, K. (2016b). Constructing the impossible. In Walters, L. and Hawthorne, J., editors. *Conditionals, Probability, and Paradox: Themes from the Philosophy of Dorothy Edgington*. Oxford: Oxford University Press, forthcoming.

Fine, K. (2017). A theory of truth-maker content I & II. *Journal of Philosophical Logic*, **46**, 625–702.

Fine, K. (2018). Acts and embodiment. In Hinshelwood, A., editor. *Being and Doing*. Oxford: Oxford University Press, forthcoming.

Frege, G. (1897 [1967]). Begriffsschrift, a formula language, modeled upon that of arithmetic, for pure thought. In van Heijenoort, J., editor. *From Frege to Gödel: A Source*

Book in Mathematical Logic, 1879–1931. Cambridge, MA: Harvard University Press, pp. 5–82.

Geach, P. T. (1966). Dr. Kenny on practical inference. *Analysis*, **26**, 76–79.

Gombay, A. (1965). Imperative inference and disjunction. *Analysis*, **25**(3), 58–62.

Kamp, H. (1973). Free choice permission. *Proceedings of the Aristotelian Society, N.S.*, **74**, 57–74.

Kaufmann, M. S. (2012). *Interpreting Imperatives*. New York: Springer.

Kenny, A. J. (1966). Practical inference. *Analysis*, **26**, 65–75.

Hare, R. M. (1952). *The Language of Morals*. Oxford: Clarendon Press.

Hare, R. M. (1967). Some alleged differences between imperatives and indicatives. *Mind*, **76**, 309–326.

Mastop, R. (2005). What can you do: Imperative mood in semantic theory. Doctoral Dissertation, University of Amsterdam.

Mastop, R. (2012). Imperatives as semantic primitives. *Linguistics and Philosophy*, **34**, 305–340.

McGinn, C. (1977). Semantics for non-indicative sentences. *Philosophical Studies*, **32**, 301–311.

Parry, W. T. (1933). An axiom system for a new kind of implication (analytical implication). *Results of a Mathematical Colloquium*, **4**, 5–6.

Portner, P. (2004). The semantics of imperatives within a theory of clause types. In Watanabe, K. and Young, R. B. editors. *Proceedings of SALT 14*. New York: CLC Publications, pp. 235–252.

Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, **15**, 351–383.

Ross, A. (1941). Imperatives and logic. *Theoria*, **25**(7), 53–71.

Segerberg, K. (1989). Bringing it about. *Journal of Philosophical Logic*, **18**(4), 327–347.

Segerberg, K. (1990). Validity and satisfaction in imperative logic. *Notre Dame Journal of Formal Logic*, **31**(2), 203–221.

Starr, W. (2018). A preference semantics for imperatives. *Semantics and Pragmatics*, to appear. Available at https://philpapers.org/rec/STAAPS.

Stelzner, W. (1992). Relevant deontic logic. *Journal of Philosophical Logic*, **21**, 193–216.

van Fraassen, B. (1969). Facts and tautological entailments. *Journal of Philosophy*, **66**(15), 477–487.

van Rooij, R. (2000). Permission to change. *Journal of Semantics*, **17**(2), 119–143.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, **25**, 221–261.

Veltman, F. (2009). Imperatives at the semantics/pragmatics borderline. Unpublished manuscript, University of Amsterdam.

Williams, B. A. O. (1963). Imperative inference. *Analysis*, **23**(Suppl.), 30–36.

Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.

DEPARTMENT OF PHILSOPHY
NEW YORK UNIVERSITY
NEW YORK, NY 10003, USA
E-mail: kf14@nyu.edu