

Intention: Hyperintensional Semantics and Decision Theory

Timothy Alison Bowen (Hasen Joseph Khudairi, prior to March, 2023)

Compiled: May 26, 2015; Revised: December 8, 2023

Abstract

This paper argues that the types of intention can be modeled both as modal operators and via a multi-hyperintensional semantics. I delineate the semantic profiles of the types of intention, and provide a precise account of how the types of intention are unified in virtue of both their operations in a single, encompassing, epistemic space, and their role in practical reasoning. I endeavor to provide reasons adducing against the proposal that the types of intention are reducible to the mental states of belief and desire, where the former state is codified by subjective probability measures and the latter is codified by a utility function. I argue, instead, that each of the types of intention – i.e., intention-in-action, intention-as-explanation, and intention-for-the-future – has as its aim the value of an outcome of the agent’s action, as derived by her partial beliefs and assignments of utility, and as codified by the value of expected utility in evidential decision theory.

1 Introduction

Formal treatments of imperatival notions have been pursued both logically and semantically. In the logical setting, deontic claims have been interpreted as types of a modal operator, where a condition holding across the points of a space abbreviates the property of obligation, and its dual abbreviates the property of permissibility.¹ In the twentieth century, research in deontic logic has examined the validity of the rule of necessitation ($\vdash\phi \rightarrow \vdash\Box\phi$) (von Wright, 1981); modal axiom 4 ($\Box\phi \rightarrow \Box\Box\phi$) (cf. Barcan, 1966); and modal axiom GL [$\Box(\Box\phi \rightarrow \phi) \rightarrow \Box\phi$] (cf. Smiley, 1963). The semantic approach has been inspired by the works of Kratzer (1977, 2012), Stalnaker (1978), and Veltman (1996), arguing that there are modal operators on a set of points which are not straightforwardly truth-conditional, instead recording an update on that set which is taken to be pragmatic (cf. Yalcin, 2012). The types of obligation have proliferated, as variations on the ‘ought’-operator – e.g., what one ought to do relative to

¹Deontic logic dates from at least as early as the fourteenth century, in the writings of Ockham, Holcot, and Rosetus. See Knuuttila (1981) for further discussion.

a time and one's states of information, by contrast to what one ought to do relative to the facts – have been codified by differences in the array of intensional parameters relative to which the operator receives a semantic value (cf. Yalcin, *op. cit.*; Cariani, 2013; Dowell, 2013; et al).

This essay aims to provide a theory of the structural content of the types of intention via a similar modal analysis; to explain the role of intention in practical reasoning; and to answer thereby what I will call the unification problem: i.e., the inquiry into how the various types of intention comprise a unified mental state.² The general significance of the present contribution is that it will provide some foundational structure to the topic, where the previous lack thereof has served only to exacerbate its intransigence.³ I will argue that – similarly to the case of deontic judgment – the foregoing types of intention can be countenanced as modal operators. The defining contours of the contents of the states may thus be targeted via their intensional-semantic profile. The types of intention on which I will focus include (i) the notion of 'intention-in-action', as evinced by cases in which agents act intentionally; (ii) the notion of 'intention-with-which', where an agent's intentions figure as an explanation of their actions; and (iii) the notion of 'intention-for-the-future', as evinced by an agent's plans to pursue a course of action at a future time.

I will argue that the unification problem has at least two, consistent solutions. The first manner in which the operations of intention are unified is that they are defined on a single space, whose points are states of information or epistemic possibilities. I argue, then, that the significance of examining how the state of intention interacts with practical reasoning is that it provides a second means by which to account for the unity of intention's types. Although each type of intention has a unique formal clause codifying its structural content, the notions of 'intention-in-action', 'intention-with-which', and 'intention-for-the-future' are nevertheless unified, because each is directed toward the property of expected utility. Thus, acting intentionally, acting because of an intention, and intending to pursue a course of action at a future time, are mental states whose unification consists in that each type aims toward the satisfaction of the value of an outcome – the value of which is the product of a partial belief conditional on one's acts by the utility thereof. The dissociation between an agent's intention to pursue an action and the causal relevance of the action's outcome adduces in favor of the characterization of expected utility in the setting of evidential,

²The unification problem is first examined in Anscombe (1963), and has been pursued in contemporary research by, *inter alia*, Bratman (1984) and Setiya (2014).

³Compare the aims and methods pursued in the research projects of Fine (1981) and Williamson (2014): 'The relevance of the undertaking [...] consists mainly in the general advantages that accrue from formalizing an intuitive theory. First of all, one thereby obtains a clearer view of its primitive notions and truths. This is no small thing in a subject [...] that is so conspicuously lacking in proper foundations' (Fine, *op. cit.*); 'The aim is to gain insight into a phenomenon by studying how it works under simplified, rigorously described conditions that enable us to apply mathematical or quasi-mathematical reasoning that we cannot apply directly to the phenomenon as it occurs in the wild, with all its intractable complexity. We can then cautiously transfer our insight about the idealized model back to the phenomenon in the wild' (Williamson, *op. cit.*).

rather than causal, decision theory. The proposal that the content of intention is expected utility has, furthermore, the virtue of generalizing, in order to explain the nature of the intentions of non-human organisms. The contents of non-human organisms' intentions can here be understood as the value expected by their actions, as sensitive to both their prediction that the outcome will occur and the utility of its occurrence. Finally, because the aim of intention is expected utility, a precise account can be provided of how intention relates to the notions of belief and desire, while yet retaining its status as a unique mental state.

In Sections 2-3, I delineate the intensional-semantic profiles of the types of intention, and provide a precise account of how the types of intention are unified in virtue of both their operations in a single epistemic modal space and their role in practical reasoning, i.e., evidential decision theory. I endeavor to provide reasons adducing against the proposal that the types of intention are reducible to the mental states of belief and desire, where the former state is codified by subjective probability measures and the latter is codified by a utility function. Section 4 provides concluding remarks.

2 The Modes of Intention

The epistemic modal space of an agent can be defined via a frame, comprised of a set of points, and a relation of accessibility thereon (cf. Kripke, 1963; Blackburn et al, 2001). The points in the frame are here interpreted as an agent's states of information, while the relation of accessibility can receive various interpretations. A state of information is possible, just if there is at least one point relative to which it is true, if and only if it is not necessary for the formula to be false. One of the states of information is necessary, just if it is true everywhere, i.e. relative to all the other points in the space, if and only if it is impossible for it to be false. The distinctly epistemic interpretation of possibility comes in at least two guises, defined as the dual of epistemic necessity (' $\diamond\phi$ ' iff ' $\neg\Box\neg\phi$ '): The truth of a formula is epistemically possible, just if the formula is conceivable to the agent. The epistemic interpretation of necessity can itself come in at least two guises: The truth of a formula is epistemically invariant or necessary, just if the truth of the formula is known by an agent, or if it is inconceivable that the formula is false such that the formula is in one sense apriori.

When an agent intends to ϕ , their intention may fall into three distinct types. One type of intention concerns the intentional pursuit, by the agent, of a course of action. A second type of intention can be witnessed, when the agent cites an intention as an explanation of her pursuit of a course of action. Finally, a third type of intention can be witnessed, when the agent intends to pursue a course of action at a future time.

2.1 Intention-in-Action

If the agent acts intentionally, then her intention can be understood as an operation relative to her states of information. The agent acts intentionally, just if there is a world and a unique array of intensional parameters relative to which her intention is realized and receives a positive semantic value. The array of intensional parameters is multi-dimensional, because the value of intending to ϕ relative to one of the parameters will constrain the value of intending to ϕ relative to the subsequent parameters. Thus, we can say that an agent intends to ϕ , if and only if she acts intentionally, only if there is both a world and array of intensional parameters, relative to which her intention is realized, i.e. receives a positive value. The intensional parameters include a context comprised of a time and location, and a pair of indices on which spaces of the agents acts and of the outcomes of her actions are built. So, the agent's intention-in-action receives a positive semantic value only if there is at least one world in her epistemic modal space at which – relative to the context of a particular time and location, which constrains the admissibility of the actions as defined at a first index, and which subsequently constrains the outcome thereof as defined at a second index – the intention is realized.

- $\llbracket \text{Intention-in-Action}(\phi) \rrbracket_w = 1$ only if $\exists w' \llbracket \phi \rrbracket^{w', c(=t,l), a, o} = 1$

2.2 Intention-with-which

If the agent refers to an intention, in order to explain her pursuit of a course of action, then her intention can similarly be understood as an operation relative to her states of information. In this case, the agent intends to ϕ , just if there is a pair of formulas defined at points in her epistemic modal space, where one of the states is realized because it holds in virtue the other state being realized. Informally, the foregoing explanation can be referred to as the intention-with-which she acts. Thus, we can say that an agent intends to ϕ , if and only if her intention is an explanation for her action, only if she acts in pursuit of ψ because she intends to ϕ . In order to capture the notion of one formula holding in virtue, or because, of a distinct formula, we define grounding operators on the agent's epistemic modal space. Thus, the agent intends to ϕ because, there is an intention in virtue of which her action, ψ so as to realize ϕ , receives a positive value.

- $\llbracket \text{Intention-with-which}(\phi) \rrbracket_w = 1$ only if $\exists w' [\llbracket \psi \rrbracket^{w'} = 1 \wedge \llbracket G(\phi, \psi) \rrbracket = 1]$,

where $G(x,y)$ is a grounding operator encoding the explanatory connection between ϕ and ψ . Following Fine (2012a,b), the grounding operator can have the following properties: The grounding operator is *weak* if and only if it induces reflexive grounding. The operator is *strict* if and only if it is not weak. The operator is *full* if and only if the intention to ϕ uniquely provides the explanatory ground for the action, ψ . The operator is *part* if and only if the intention to ϕ - along with other reasons for action - provide the explanatory ground for

the action, ψ . Combinations of the foregoing explanatory operators may also obtain: $x < y$ iff ϕ is a *strict full* ground for ψ ; $x \leq y$ iff ϕ is a *weak full* ground for ψ ; $x \prec y$ iff ϕ is a *strict part* ground for ψ ; $x \preceq y$ iff ϕ is a *weak part* ground for ψ ; $x \preceq y \wedge \neg(y \preceq x)$ iff ϕ is a *strict partial* ground for ψ ; $x \prec^* y$ iff $x_1, \dots, x_n \leq y$, iff ϕ is a *partial strict* ground for ψ ; $x \prec' z$ iff $[\phi \prec^* \psi \wedge \psi \preceq \mu]$ iff ϕ is a *part strict* ground for some further action, μ .

2.3 Intention-for-the-Future

Finally, an agent can intend to ϕ , because she intends to pursue a course of action at a future time. In this case, the intensional-semantic profile which records the parameters relative to which her intention receives a positive semantic value converges with a future-directed modal operator to the effect that the agent *will* ϕ . Thus, an agent realizes an intention-for-the-future only if there is a possible world and a future time, relative to which the possibility that ϕ is realized can be defined. Thus:

- $\llbracket \text{Intention-for-the-future}(\phi) \rrbracket_w = 1$ only if $\exists w' \forall t \exists t' [t < t' \wedge \llbracket \phi \rrbracket^{w', t'} = 1]$.⁴

A multi-hyperintensional semantics for the types of intention can be provided as well. In this case, the worlds figuring in the parameters of the foregoing semantic clauses for the modal operators are replaced by topic- (i.e. subject matter) sensitive truthmakers. The truthmakers can be interpreted epistemically, and as being parts of worlds rather than whole worlds themselves. Thus epistemic topic-sensitive truthmakers figure in, and thus inform, the parameters relative to which intentions are satisfied.

Following the presentation of topic models in Berto (2018; 2019), Canavotto et al (2020), and Berto and Hawke (2021), atomic topics comprising a set of topics, T , record the hyperintensional intentional content of atomic formulas, i.e. what the atomic formulas are about at a hyperintensional level. Topic fusion is a binary operation, such that for all $x, y, z \in T$, the following properties are satisfied: idempotence ($x \oplus x = x$), commutativity ($x \oplus y = y \oplus x$), and associativity $[(x \oplus y) \oplus z = x \oplus (y \oplus z)]$ (Berto, 2018: 5). Topic parthood is a partial order, \leq , defined as $\forall x, y \in T (x \leq y \iff x \oplus y = y)$ (op. cit.: 5-6). Atomic topics are defined as follows: $\text{Atom}(x) \iff \neg \exists y < x$, with $<$ a strict order. Topic parthood is thus a partial ordering such that, for all $x, y, z \in T$, the following properties are satisfied: reflexivity ($x \leq x$), antisymmetry ($x \leq y \wedge y \leq x \rightarrow x = y$), and transitivity ($x \leq y \wedge y \leq z \rightarrow x \leq z$) (6). A topic frame can then be defined as $\{W, R, T, \oplus, t\}$, with t a function assigning atomic topics to atomic formulas. For formulas, ϕ , atomic formulas, p, q, r (p_1, p_2, \dots), and a set of atomic topics, $Ut\phi = \{p_1, \dots, p_n\}$, the topic of ϕ , $t(\phi) = \oplus Ut\phi = t(p_1) \oplus \dots \oplus t(p_n)$ (op. cit.). Topics are hyperintensional, though not as fine-grained as syntax. Thus $t(\phi) = t(\neg\neg\phi)$, $t\phi = t(\neg\phi)$, $t(\phi \wedge \psi) = t(\phi) \oplus t(\psi) = t(\phi \vee \psi)$ (op. cit.).

⁴See Rao and Georgeff (1991), for the suggestion that operators in a multi-modal logic can model the notion of goal-oriented intention. The foregoing intensional semantics is consistent with the logic that they proffer.

According to truthmaker semantics for epistemic logic, a modalized state space model is a tuple $\langle S, P, \leq, v \rangle$, where S is a non-empty set of states, i.e. parts of the elements in A , P is the subspace of possible states where states s and t comprise a fusion when $s \sqcup t \in P$, \leq is a partial order, and $v: \text{Prop} \rightarrow (2^S \times 2^S)$ assigns a bilateral proposition $\langle p^+, p^- \rangle$ to each atom $p \in \text{Prop}$ with p^+ and p^- incompatible (Hawke and Özgün, forthcoming: 10-11). Exact verification (\vdash) and exact falsification (\dashv) are recursively defined as follows (Fine, 2017: 19; Hawke and Özgün, forthcoming: 11):

- $s \vdash p$ if $s \in \llbracket p \rrbracket^+$
(s verifies p , if s is a truthmaker for p i.e. if s is in p 's extension);
- $s \dashv p$ if $s \in \llbracket p \rrbracket^-$
(s falsifies p , if s is a falsifier for p i.e. if s is in p 's anti-extension);
- $s \vdash \neg p$ if $s \dashv p$
(s verifies not p , if s falsifies p);
- $s \dashv \neg p$ if $s \vdash p$
(s falsifies not p , if s verifies p);
- $s \vdash p \wedge q$ if $\exists v, u, v \vdash p, u \vdash q$, and $s = v \sqcup u$
(s verifies p and q , if s is the fusion of states, v and u , v verifies p , and u verifies q);
- $s \dashv p \wedge q$ if $s \dashv p$ or $s \dashv q$
(s falsifies p and q , if s falsifies p or s falsifies q);
- $s \vdash p \vee q$ if $s \vdash p$ or $s \vdash q$
(s verifies p or q , if s verifies p or s verifies q);
- $s \dashv p \vee q$ if $\exists v, u, v \dashv p, u \dashv q$, and $s = v \sqcup u$
(s falsifies p or q , if s is the fusion of the states v and u , v falsifies p , and u falsifies q);
- $s \vdash \forall x \phi(x)$ if $\exists s_1, \dots, s_n$, with $s_1 \vdash \phi(a_1), \dots, s_n \vdash \phi(a_n)$, and $s = s_1 \sqcup \dots \sqcup s_n$
[s verifies $\forall x \phi(x)$ "if it is the fusion of verifiers of its instances $\phi(a_1), \dots, \phi(a_n)$ " (Fine, 2017c)];
- $s \dashv \forall x \phi(x)$ if $s \dashv \phi(a)$ for some individual a in a domain of individuals (op. cit.)
[s falsifies $\forall x \phi(x)$ "if it falsifies one of its instances" (op. cit.)];
- $s \vdash \exists x \phi(x)$ if $s \vdash \phi(a)$ for some individual a in a domain of individuals (op. cit.)
[s verifies $\exists x \phi(x)$ "if it verifies one of its instances $\phi(a_1), \dots, \phi(a_n)$ " (op. cit.)];
- $s \dashv \exists x \phi(x)$ if $\exists s_1, \dots, s_n$, with $s_1 \dashv \phi(a_1), \dots, s_n \dashv \phi(a_n)$, and $s = s_1 \sqcup \dots \sqcup s_n$ (op. cit.)
[s falsifies $\exists x \phi(x)$ "if it is the fusion of falsifiers of its instances" (op. cit.)];
- s exactly verifies p if and only if $s \vdash p$ if $s \in \llbracket p \rrbracket$;
- s inexactly verifies p if and only if $s \triangleright p$ if $\exists s' \leq s, s' \vdash p$; and
- s loosely verifies p if and only if, $\forall v, s.t. s \sqcup v \vdash p$ (35-36);

The interaction between states and topics is then that a union of topics and truthmakers can be appealed to left of the turnstile in the truthmaker clauses.

This section has endeavored to accomplish two aims. The first was to provide a precise delineation of the structural content of, and therefore the distinctions between, the types of intention. Intention was shown to be either a modal or a hyperintensional mental state, whose operations have a unique profile, and whose values are defined relative to an agent's space of states of information. The second aim was to secure one of the means by which the unity of the distinct types of intention can be witnessed. Despite that each of the types of intention has a unique structural content, the contents of those types are each defined in a single, encompassing space; i.e, relative to the agent's space of epistemic states or possibilities.

3 Intention in Decision Theory

In Section 2, I suggested that intention is a unified, modal mental state, the contents of which are defined relative to an agent's states of information. This section examines the proposal that intentions have a dual profile (cf. Bratman, *op. cit.*), because intentions figure constitutively in practical reasoning. I argue that, because expected utility theories are the only axiomatized theories of practical reasoning, an account must be provided of the role that intention plays therein. The account will illuminate a precise relationship – which I argue is not identity – between the types of intention and the mental states of belief and desire. The account will further serve to provide a second explanation for the unity of intention's types, given the uniform role that the types of intention play in decision theory.

A model of decision theory can be understood as a tuple $\langle A, O, K, V \rangle$, where A is a set of acts; O is a set of outcomes; K encodes a set of counterfactual conditionals, where an act from A figures in the antecedent of the conditional and O figures in the conditional's consequent; and V is a function assigning a real number to each outcome. The real number is a representation of the value of the outcome. The expected value of the outcome is calculated as the product of (i) the subjective probability – i.e., the agent's partial belief or credence – that the outcome will occur, as conditional on her act, and (ii) the value or utility which she assigns to the outcome's occurrence. The agent can prefer one assignment of values to the outcome's occurrence over another. (Which preference axioms ought to be adopted is a contentious issue, and will not here be examined. Cf. von Neumann and Morgenstern, 1944; Savage, 1954; Jeffrey, 1983; and Joyce, 1999.) In evidential decision theory, the expected utility of an outcome is calculated as the product of the agent's credence conditional on her action, by the utility of the outcome. In causal decision theory, the expected utility of an outcome is calculated as the product of the agent's credence, conditional on both her action and the causal efficacy thereof, by the utility of the outcome. Expected utility can further be augmented by a risk-weighting function: If the agent's expected utility diminishes with the order of the bets she might pursue – such that expected utility is sensitive to the agent's propensity to take risks relative to the total ordering of the gambles – then she might have a preference

for a sure-gain of .5 units of value, rather than prefer a bet with a 50 percent chance of winning either 0 or 1 units of value (cf. Buchak, 2014).

If intention plays a constitutive role in practical reasoning, and decision theories provide the most tractable models thereof, then what is the role of intention in decision theory? The parameters in the axiomatizations of decision theory encode variables for credences, actions, outcomes, assignments of utility, background states of information pertinent to the causal relevance of actions on outcomes, and the agent's preferences. Expected utility is derived, as noted, by the interaction between an agent's credences, actions, and utility assignments. Which, then, of these parameters do an agent's intentions concern?

There are dissociations between intention and belief and between intention and desire. An agent can have a partial belief that the sun will rise, without intending to pursue any course of action. Conversely, an agent can intend to pursue a course of action, yet appreciate that there are, unfortunately, reasons for her to disbelieve that the act will obtain. An agent can desire that the sun rises, without the intention to entrain the sun's rising as consequence. Conversely, a vegetarian can intend to consume meat, if it is the only available source of protein and they are in dire need thereof, while yet desire a distinct and orthogonal outcome.

There are dissociations between intention and preference. An agent can prefer the sun's rising to the prevalence in her life of unprovoked antagonists, without either acting intentionally, possessing an intention as an explanation for some course of action, or intending to pursue any particular course of action in the future. Conversely, whether an agent's intention to pursue an action mandates a preference for the value of the outcome of that action will depend on one's preference axioms. One such axiom might be maximin, according to which the best of the worst outcomes among a set of options should be preferred, while a distinct rule might be maximax, according to which one ought to prefer and pursue the maximally valuable outcome among a set of options. Thus, intending to ϕ is not sufficient for determining whether ϕ ought to be preferred.

There are, finally, dissociations between intention and action. One might intend to calculate the value of a formula, yet not be able so to act, because their attention might be allocated elsewhere.

Acting intentionally, intending to pursue a course of action in the future, and citing an intention as an explanation for one's course of action are each, however, in some way related to the value of a course of action. When an agent acts intentionally, she acts in such a way so as to obtain an outcome that she values. When an agent pursues a course of action, and refers to her intention so to act as the explanation for that action, the intention explains the value, for the agent, in which the action and its outcomes are supposed to consist. Finally, when an agent intends to pursue a course of action in the future, her intention is similarly guided by the value of the outcome that her action will hopefully entrain. The value of the outcome will not be her bare assessment of the utility of the outcome, because – in the setting of decision theory – utility functions codify desires, such that her intention would thereby be elided with her desire for the outcome.

Because the types of intention are all directed toward the value of an outcome of a course of action – while being irreducible to, because dissociable from, the states of belief and desire – the remaining and most suitable candidate for the role of the mental state of intention in decision theory is the aim of expected utility; i.e., the value of an outcome, as arising by the interaction between the agent’s partial belief or expectation that the outcome will occur as conditional on her act, and the utility that she associates with the outcome’s occurrence. Because of the dissociation between an agent’s intention to pursue a course of action in the future and the action’s occurrence – let alone the dissociation between the intention to act in future, and the causal efficacy of the action were it to obtain – the role of intention in practical reasoning appears to be more saliently witnessed in the setting of evidential decision theory.

That the types of intention are each directed toward expected utility evinces how an agent’s intentions can be sensitive to her beliefs and desires, without being reducible to them. Crucially, moreover, that the types of intention are each directed toward expected utility provides a second explanation of the way that the types of intention comprise a unified mental state.

Theoretical advantages accruing to the foregoing proposal include that it targets a foundational role for intention in decision theory. The proposal might be foundational, because it targets a basic role for intention in practical reasoning, which is consistent with the possible augmentation of the proposal with other approaches which assume a more cognitively demanding role for intention’s aims. Such approaches include proposals to the effect (i) that the most fundamental type of intention is intention-with-which, such that intention’s role as an explanation can be elided with its causal efficacy (cf. Anscombe, *op. cit.*; Davidson, 1963); (ii) that the content of intention is the diachronic satisfaction of self-knowledge (cf. Velleman, 1989); and (iii) that the role of intention in practical reasoning ought to be understood as an evaluative constraint on practical reasoning, as determined by the virtuous traits of an agent’s character (cf. Setiya, 2007).

4 Concluding Remarks

I have argued that the unification problem for the types of intention can be solved in two, consistent ways. The types of intention can be modeled as modal operators, where the unity of the operations consists, in the first instance, in that their values are defined relative to a single, encompassing, epistemic modal space. The second manner by which the unity of intention’s types can be witnessed is via intention’s unique role in practical reasoning. I argued that each of the types of intention – i.e., intention-in-action, intention-as-explanation, and intention-for-the-future – has as its aim the value of an outcome of the agent’s action, as derived by her partial beliefs and assignments of utility, and as codified by the value of expected utility in evidential decision theory. A precise account was thereby provided of the role of epistemic modality in the unification of the types of a unique, modal mental state, whose value figures constitutively in

decision-making and practical reason.

References

- Anscombe, G.E.M. 1957/1963. *Intention*. Harvard University Press.
- Barcan Marcus, R. 1966. Iterated Deontic Modalities. *Mind*, 75.
- Berto, F. 2018. Aboutness in Imagination. *Philosophical Studies*, 175.
- Berto, F. 2019. The Theory of Topic-Sensitive Intentional Modals. In I. Sedlar and M. Blichta (eds.), *The Logica Yearbook*, 2018. College Publications.
- Berto, F., and P. Hawke. 2021. Knowledge relative to Information. *Mind*, 130: 517.
- Blackburn, P., M. de Rijke, and Y. Venema. 2001. *Modal Logic*. Cambridge University Press.
- Bratman, M. 1984. Two Faces of Intention. *Philosophical Review*, 93:3.
- Buchak, L. 2014. *Risk and Rationality*. Oxford University Press.
- Canavotto, I., F. Berto, and A. Giordani. 2020. Voluntary Imagination: a Fine-grained Analysis. *Review of Symbolic Logic*.
- Cariani, F. 2013. 'Ought' and Resolution Semantics. *Nous*, 47:3.
- Davidson, D. 1963. Actions, Reasons, and Causes. *Journal of Philosophy*, 23.
- Dowell, J.L. 2013. Flexible Contextualism about Deontic Modals. *Inquiry*, 56.
- Fine, K. 1981. First-order Modal Theories I – Sets. *Nous*, 15:2.
- Fine, K. 2012a. The Pure Logic of Ground. *Review of Symbolic Logic*, 5:1.
- Fine, K. 2012b. Guide to Ground. In F. Correia and B. Schnieder (eds.), *Metaphysical Grounding*. Cambridge University Press.
- Fine, K. 2017a. A Theory of Truthmaker Content I: Conjunction, Disjunction, and Negation. *Journal of Philosophical Logic*, 46:6.
- Jeffrey, R. 1965/1990. *The Logic of Decision*. University of Chicago Press.
- Joyce, J. 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press.
- Knuuttila, S. 1981. The Emergence of Deontic Logic in the Fourteenth Century. In R. Hilpinen (ed.), *New Studies in Deontic Logic*. D. Reidel Publishing.
- Kratzer, A. 1977. What Must and Can Must and Can Mean. *Linguistics and Philosophy*, 1:3.
- Kratzer, A. 2012. *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press.
- Kripke, S. 1963. Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16.

- Rao, A., and M. Georgeff. 1991. Modeling Rational Agents within a BDI-Architecture. In J. Allen, R. Fikes, and E. Sandewall (eds.), *KR '91: Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann.
- Savage, L. 1954/1972. *The Foundations of Statistics*. Dover.
- Setiya, K. 2007. *Reasons without Rationalism*. Princeton University Press.
- Setiya, K. 2014. Intention. *Stanford Encyclopedia of Philosophy*.
- Smiley, T. 1963. The Logical Basis of Ethics. *Acta Philosophica Fennica*, 16.
- Stalnaker, R. 1978. Assertion. In P. Cole (ed.), *Syntax and Semantics*, vol. 9. Academic Press.
- Velleman, J.D. 1989. *Practical Reflection*. Princeton University Press.
- Veltman, F. 1996. Defaults in Update Semantics. *Journal of Philosophical Logic*, 25:3.
- von Neumann, J., and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- Williamson, T. 2014. A Note on Gettier Cases in Epistemic Logic. *Philosophical Studies*, DOI 10.1007/s11098-014-0357-1
- Yalcin, S. 2012. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, doi: 10.1111/j.1467-9264.2012.00329.x