

I OUGHT, THEREFORE I CAN OBEY*

Peter B. M. Vranas

vranas@wisc.edu

University of Wisconsin-Madison

28 December 2017

Abstract. According to typical ought-implies-can principles, if you have an obligation to vaccinate me tomorrow, then you *can* vaccinate me tomorrow. Such principles are uninformative about *conditional* obligations: what if you only have an obligation to vaccinate me tomorrow *if* you synthesize a vaccine today? Then maybe you *cannot* vaccinate me tomorrow (e.g., because you cannot synthesize a vaccine); what you *can* do instead, I propose, is make it the case that the conditional obligation is not violated (i.e., that you do not both synthesize a vaccine today and fail to vaccinate me tomorrow). More generally, I propose the *ought-implies-can-obey* principle: an agent has an obligation only at times at which the agent *can obey* the obligation (i.e., can make it the case that the obligation is not violated). I also propose another principle, which captures the idea that “ought” implies “can avoid”. I defend both principles mainly by arguing that they help explain why agents *lose* (i.e., stop having) obligations, including conditional ones.

1. Introduction

Twenty years ago, when Paul’s daughter was in kindergarten, Paul promised her that, if she ever obtained a doctorate, he would immediately buy her a three-bedroom apartment in Manhattan as a gift. Motivated by this promise, which Paul repeatedly reaffirmed in later years, Paul’s daughter has studied hard, and she has just obtained her doctorate today. Paul was very wealthy when he made the promise, but he has recently gambled away his wealth; he is now heavily in debt, and no one will give him a loan. Despite his promise, Paul *cannot* buy his daughter an apartment in Manhattan today (or indeed anytime soon). Does he nevertheless have an *obligation* to do so?

I believe that he does not. I subscribe to the following ought-implies-can principle: (OIC) if an agent at a given time has an obligation to do something, then the agent at that time can do the thing. (I formulate OIC more carefully in §2.) Using OIC and the claim that right now Paul cannot buy his daughter an apartment in Manhattan, I infer that right now he has no obligation to do so. I have extensively defended OIC in previous work (Vranas 2007), but it remains highly controversial.¹ Nevertheless, my main goal in this paper is not to defend OIC further. (I do this in Vranas 2017a.) My main goal is instead to fill two lacunas in the literature on ought-implies-can principles.

* I am grateful to Norman Dahl, Alan Hájek, Russ Shafer-Landau, Alan Sidelle, Elliott Sober, Michael Titelbaum, and especially Aviv Hoffmann and some anonymous reviewers for comments, and to John Bengson and Christopher Hitchcock for help. Thanks also to Martin Barrett, Farid Masrour, James Messina, and Larry Shapiro for interesting questions, and to my mother for typing the bulk of the paper. Material from this paper was presented at the University of Wisconsin-Madison (May 2016) and at the University of Copenhagen Facts & Norms Workshop III (August 2017).

¹ For responses to my defense of OIC, see Littlejohn 2009 and Mizrahi 2009. For recent objections to ought-implies-can principles, see: Buckwalter & Turri 2015; Chituc et al. 2016; Graham 2011a, 2011b; Henne et al. 2016; Mizrahi 2015a, 2015b; Talbot 2016; Turri 2017. For replies to some of these objections, see: Cohen 2017; Hannon 2017; Hobbs 2013: 59; Kurthy & Lawford-Smith 2015; Kurthy et al. 2017; Littlejohn 2012. For further recent discussion of ought-implies-can principles, see: Besch 2011; Bloomfield 2007; Copp 2008; Haji 2012; Heuer 2010; Hobbs 2013; Howard-Snyder 2006, 2013; King 2014, 2017; Lowry 2012; Martin 2009; Southwood 2016; Streumer 2007, 2010; Wedgwood 2013.

To explain one of the lacunas, suppose Paula has just promised her son that, if she wins a thousand dollars in the lottery this month, she will buy him a château in southern France as a gift next month. Paula is very poor, and is currently unable to buy her son a château next month. Therefore, according to OIC, Paula currently has no (unconditional) obligation to buy her son a château next month. But does she currently have a *conditional* obligation to buy her son a château next month *if* she wins a thousand dollars in the lottery this month? According to the *spirit* of OIC, she does not: a thousand dollars is not nearly enough money, even as a down payment, to enable Paula to buy a château. OIC, however, like every other ought-implies-can principle in the literature with which I am familiar, fails to deliver this result: OIC is uninformative about *conditional* obligations. In this paper, I fill this lacuna by formulating and defending an ought-implies-can principle—which I call the *ought-implies-live-option* principle (OILO)—that entails OIC but has the advantage of being also informative about *conditional* obligations.

To explain the other lacuna, suppose Pauline has won a million dollars in the lottery, and has spent it to buy her nephew a château in southern France as a gift. Her nephew, however, due to an irreversible biochemical imbalance, is unable to ever feel gratitude. Does he have an obligation to feel grateful to her? According to the *spirit* of OIC, he does not, since he cannot. OIC, however, like almost every other ought-implies-can principle in the literature with which I am familiar (see Brennan & Southwood 2007 for an exception), fails to deliver this result: OIC is formulated only in terms of obligations to *do*, not in terms of obligations to *feel*. In this paper, I fill this lacuna by formulating my preferred ought-implies-can principle—namely OILO—in terms of *all* obligations that agents have, not just obligations to *do*. I argue that OILO has the advantages of (1) explaining why people normally have no obligations to *feel* and (2) being also informative about obligations to *believe*.

To sum up, my main project in this paper is to generalize OIC so as to reach an ought-implies-can principle that is informative about all kinds of obligations that agents have. I also have a secondary project, namely to defend a principle which captures the idea that “ought” implies “can avoid”. These projects may be most appealing to those who are sympathetic to ought-implies-can principles to start with. Nevertheless, even opponents of such principles should find parts of the paper interesting; for example, I propose novel counterexamples to some ultimately unsatisfactory ought-implies-can principles before I reach my preferred principle, namely OILO. I reach OILO in three steps. (1) In §2, I generalize OIC to reach the *ought-implies-can-satisfy* principle (OICS), which is informative about *unconditional* obligations. (2) In §3, I generalize OICS to reach the *ought-implies-can-obey* principle (OICO), which is also informative about *conditional* obligations. In parallel, I also defend the *ought-implies-possible-violation* principle (OIPV), which captures the idea that “ought” implies “can avoid”. (3) Finally, in §4, I conjoin OICO with OIPV to get OILO.

2. Unconditional obligations

2.1. The starting point: OIC

Since my project in this paper is to generalize the ought-implies-can principle that I have defended in previous work, namely OIC, I start by formulating OIC precisely (Vranas 2007: 169-71):

(OIC) By virtue of conceptual necessity, if an agent at a given time has an (objective, pro tanto) obligation to do something, then the agent at that time can (i.e., has both the ability and the opportunity to) do the thing.

Several remarks are in order. (1) I formulate OIC in terms of *conceptual necessity*, such as the necessity of the claim that, if I am taller than you, then you are not taller than me. Other authors, by contrast, formulate ought-implies-can principles in terms of logical or metaphysical necessity, or in terms of presupposition or conversational implicature. (2) I formulate OIC in terms of *objective, pro tanto* (i.e., *prima facie*) obligations. Other authors, by contrast, formulate ought-implies-can principles in terms of subjective or all-things-considered obligations, or in terms of ought-claims (such as the claim that you ought to grant someone's request for a favor, a claim which arguably does not entail that you have any obligation). (3) I formulate OIC in terms of *ability plus opportunity*: I understand the claim that you can do something as the claim that you have both the ability to do it (i.e., you have the requisite skills, physical and mental capacities, and knowledge) and the opportunity to do it (i.e., you are in a situation that allows you to exercise your ability; e.g., you have entered a competition, and thus you have the opportunity to win).² (For simplicity, I omit mention of opportunity in the rest of this paper.) Other authors, by contrast, use the term "can" in different ways.³ (4) I made the above choices in order to formulate a *defensible* ought-implies-can principle, namely OIC. I believe that many alternative ways of using the terms "implies", "ought", and "can" yield *indefensible* ought-implies-can principles. Nevertheless, if you are unsympathetic to OIC, you can substitute your preferred ought-implies-can principle as an alternative starting point in the quest for a general principle. Mutatis mutandis, you may still be able to benefit from much of the discussion that follows.

Obligations and abilities are relative to *times*. In the morning you do not have an obligation to run in tomorrow's marathon, but in the afternoon you do: at noon you promise to run. In the afternoon you can run in tomorrow's marathon, but in the evening you cannot: at 5pm you have an

² I understand the claim that you can do something as entailing that you have the relevant know-how. For example, if you do not know how to open your safe (because you cannot remember the combination and you have no way to find out what it is), then you *cannot* open the safe without forcing it, even if you *can* dial the particular ten-digit sequence (say, 2-5-7-6-9-0-1-2-8-8) that happens to be the combination (because you can dial *any* ten-digit sequence; cf.: Carlson 1999: 91-3, 2000: 280-1; Haji 2012: 23; Andrić 2016: 69-71), and even if (by pure luck) you *in fact* dial that sequence and thus open the safe without forcing it (cf. Southwood & Wiens 2016; contrast Miller 2003: 54). According to OIC, in this case (1) you have no *objective* obligation to open the safe without forcing it (even if you have promised to do so). On the other hand, OIC is compatible with the claim that (2) you have an objective obligation to dial 2-5-7-6-9-0-1-2-8-8. It is an interesting question how to formulate a theory of objective obligations which accommodates both (1) and (2); cf. Bergström 1996: 81-4.

³ Here are some references on alternative formulations of ought-implies-can principles. (1) Conceptual necessity: Brennan & Southwood 2007: 5; Hobbs 2013: 82-4; Manning 1981: 118; cf. Jacobs 1985. Logical necessity: Gensler 1996: 49-52. Metaphysical necessity: Howard-Snyder 2006: 236, 2013: 3749. Presupposition: Atkinson 1958: 58-9; Besch 2011: 352; Cooper 1966: 46; Hare 1963: 53-4; King 2017: 640-9, 656-9; Shaw 1965: 197; cf. Hampshire 1951: 163-4. Conversational implicature: Driver 2011: 191-6; King 2017: 649-56; Sinnott-Armstrong 1984, 1988: 121-6; cf.: Oppenheim 1987: 376-7; Pigden 1990: 16-21; Vogelstein 2012: 255. (2) Objective obligations: Frankena 1950/1963: 158-9; Graham 2011b: 340-1. Pro tanto obligations: Brennan & Southwood 2007: 7; Velleman 1998: 99. Subjective obligations: O'Connor 1971: 26. All-things-considered obligations: Copp 2008: 68; Dahl 1974: 487; Graham 2011b: 339; Hobbs 2013: 44; Howard-Snyder 2006: 234; Littlejohn 2009: 364; van Someren Greve 2014: 913. Ought-claims: Wedgwood 2013: 72. (3) Ability plus opportunity: Haji 1999: 180, 2002: 22, 65-6; Hobbs 2013: 82. Other uses of "can": Besch 2011: 353; Chuard & Southwood 2009: 614-5; Dahl 1974: 487-9; Lichtenberg 2010: 126-8; Smith 1961; cf. Kenny 1975: 131, 1976: 217.

accident that leaves you paralyzed.⁴ OIC is a *synchronic* principle: it relates obligations and abilities at the *same* time. For example, according to OIC, if *at noon* you have an obligation to call me by midnight, then *at noon* you can call me by midnight. In support of OIC, note that it explains why at any time after midnight you have no obligation to call me by midnight: because, setting the possibility of backwards causation aside (as I do throughout this paper), at any time after midnight you cannot call me by midnight (regardless of whether you called me by midnight). I believe that this is properly called an *explanation*;⁵ if you disagree, you can replace in this paper “explain(s) why” with “entail(s) that” whenever appropriate.

2.2. A new and more useful formulation of OIC

One might argue that the explanatory power (or the range of consequences) of OIC is rather limited. If at noon you have an obligation to call me by midnight and you first call me at 4pm, then at any time after 4pm you no longer have that obligation. One might argue that OIC does not explain why after 4pm you no longer have the obligation to call me by midnight: after you call me at 4pm, can’t you still call me by midnight? In reply, I will argue that OIC does explain this. Besides providing further support for OIC, my argument will lead to a new formulation of OIC, a formulation that will prove useful to my project of generalizing OIC. Before I give my argument, I introduce some important concepts.

Your obligation to (sooner or later) visit me is satisfied exactly if you visit me and is violated exactly if you do not visit me; call the proposition that you visit me the “satisfaction proposition” of your obligation, and call the proposition that you do not visit me the “violation proposition” of your obligation. More generally, for every obligation there are two corresponding (logically incompatible) propositions: the *satisfaction proposition* of the obligation, which specifies the conditions under which the obligation is satisfied, and the *violation proposition* of the obligation, which specifies the conditions under which the obligation is violated.⁶ Say that an agent at a

⁴ One might claim that agents only have abilities to do things in the present, not in the future: when we say that now you can run in tomorrow’s marathon, we mean that, as far as we now know, tomorrow you will be able to run. I reply that, although we can indeed mean that, in many cases agents have present abilities to do things in the future, not in the present. For example, suppose that your hands are not touching your nose, and consider your present ability to scratch your nose “now”. Strictly speaking, this is not an ability to scratch your nose *at the present time instant*; it is instead an ability to scratch your nose *in the very near future*. This is because it will take at least a few nanoseconds before one of your hands touches your nose. But if you now have the ability to scratch your nose in the very near future, why can’t you now also have the ability to scratch your nose in the more distant future? (You may not *know* whether you have the latter ability, but that is a different matter.) In response, one might grant that agents have present abilities to do things in the distant future, but might claim that, for example, you now have the ability to scratch your nose at noon tomorrow *only if* tomorrow you also have this ability. In reply, suppose that you do not scratch your nose at noon tomorrow: agents do not always do what they can do. Then, one nanosecond before noon tomorrow, you no longer have the ability to scratch your nose at noon tomorrow: it is too late for your hands to move so as to scratch your nose at noon. But if you now have the ability to scratch your nose at noon tomorrow although you no longer have it *shortly* before noon tomorrow, why can’t you now also have it even if you no longer have it *long* before noon tomorrow?

⁵ Cf. Andrić 2017: 75. It is a *non-causal* explanation, but some explanations are definitely non-causal, at least in mathematics (Lange 2014), and maybe elsewhere (Lipton 2009: 621-2; contrast Skow 2014). I am not claiming that *every* entailment provides an explanation.

⁶ Although every obligation *has* both a satisfaction proposition and a violation proposition, *identifying* the satisfaction proposition and the violation proposition of a particular obligation is not always straightforward. For example, one might argue that your obligation (arising from your promise) to visit me is satisfied *not* exactly if you visit me, but rather exactly if you visit me *with the motive of keeping your promise* (cf.: King 2014; Williams 1981: 117). In reply, distinguish (1) an unconditional obligation O_1 whose satisfaction proposition is the proposition that you visit

given time *can satisfy* an obligation exactly if the agent at that time can make it the case that the obligation is satisfied; in other words, exactly if the agent at that time can *actualize* the satisfaction proposition of the obligation (i.e., can cause that proposition to be true).⁷ Say also that a proposition is *settled*—in other words, is *historically necessary*—at a given time exactly if it is logically entailed by the *history* of the world up to and including that time (understood as the conjunction of all true propositions that are not about any later time).⁸ For example, the proposition that the sun rose yesterday is settled today. Since I am setting the possibility of backwards causation aside, I am assuming that an agent at a given time cannot actualize a proposition that is already settled (or whose negation is already settled) at that time. For example, today you cannot actualize the proposition—i.e., make it the case—that you skipped breakfast yesterday (or the proposition that you did not skip breakfast yesterday). Consequently, *if the satisfaction (or the violation) proposition of an obligation is settled at a given time, then an agent at that time cannot satisfy the obligation.*

Go back now to the example in which at noon you have an obligation to call me by midnight and you first call me at 4pm. As I said, one might argue that OIC does not explain why after 4pm you no longer have the obligation to call me by midnight: after you call me at 4pm, can't you still call me by midnight? Yes, after 4pm you can still call me by midnight, but you can no longer satisfy the obligation: you can no longer actualize its satisfaction proposition (i.e., the proposition that you call me by midnight), since this proposition is already settled (given that you called me at 4pm). So if one understands (as I do), in the formulation of OIC in §2.1, “can ... do the thing” as “can satisfy the obligation”, then OIC does explain why after 4pm you no longer have the obligation to call me by midnight: because after 4 pm you can no longer satisfy this obligation (although you can still call me by midnight). Here is then my new formulation of OIC (omitting for simplicity—both here and in the rest of this paper—the phrase “by virtue of conceptual necessity” and the parenthetical qualifications of the formulation in §2.1):

(OIC) If an agent at a given time has an obligation to do something, then the agent at that time can satisfy the obligation (i.e., can actualize its satisfaction proposition).

I take the new formulation of OIC to be an improvement over the formulation in §2.1 because the new formulation enables one to generalize OIC more easily (see §2.3 below), and also because the new formulation enables one to see that OIC has greater explanatory power than one

me from (2) an unconditional obligation O_2 whose satisfaction proposition is the proposition that you visit me with the motive of keeping your promise. For simplicity, I understand your obligation to visit me as O_1 ; I need not deny that, if you promise to visit me, sometimes you acquire O_2 instead of O_1 (but see note 18).

⁷ The claim that an agent causes a proposition to be true (i.e., *actualizes* the proposition) does not entail that the causal relata are the agent and the proposition. I take instead causes and effects to be facts (Mellor 2004); for example, the fact that the agent does something and the fact that the proposition is true. (However, nothing substantive in this paper hangs on this issue.) One might argue that, because the locution “an agent causes a proposition to be true” is not an intuitive or ordinary bit of English, the concept of actualizing requires explication. I reply that, to a first approximation, I understand actualizing as *bringing about*. Nevertheless, I introduce the term “actualize” (instead of using the familiar term “bring about”) because I want to avoid taking a stand on certain debates in the literature concerning bringing about (see note 41).

⁸ This conjunction includes propositions that are not about any time, like the proposition that some prime number is even, but excludes propositions that correspond to “soft facts”, like the proposition that before t you knew that it was going to rain after t (cf.: Hasker 1989: chap. 5; Hoffman & Rosenkrantz 1984). Note that the concept of being settled does not become trivial in a deterministic world: in such a world, every true proposition is logically entailed by the history of the world (up to and including any particular time) *in conjunction with the laws of nature*, and thus need not be settled at *every* time.

might initially think:⁹ one might initially think that OIC does not explain why after 4pm you no longer have the obligation to call me by midnight, but the new formulation enables one to see that OIC does explain this.¹⁰

2.3. The “ought-implies-can-satisfy” principle (OICS)

I formulated OIC in terms of obligations to *do* (something), which I understand as obligations to *make something the case* (i.e., to actualize a proposition). For example, your obligation to lock the door is an obligation to do: it is an obligation to make it the case that you lock the door.¹¹ Obligations to do are *unconditional* obligations (unlike obligations to do something *if* some non-tautologous condition obtains, which are *conditional* obligations; I distinguish unconditional from conditional obligations more precisely in §3.1).¹² Arguably, however, some unconditional obligations are not obligations to do. For example, suppose you have an (unconditional) obligation to *attend* my wedding, understood as satisfied exactly if you are present at my wedding. This is not an obligation to *make it the case* that you attend my wedding: if someone kidnaps you and brings you to my wedding by force, then your obligation to attend my wedding is satisfied (since you attend) but an obligation to make it the case that you attend is not satisfied (since you do not make it the case that you attend).¹³ Since your obligation to attend my wedding is not an

⁹ One might object as follows. Suppose that you have offered to treat some friends to dinner for your birthday. When you ask for the check, you realize to your embarrassment that you have forgotten your wallet at home, so you cannot pay. One of your friends offers to pay in your stead. One might argue that the “old” formulation of OIC (in §2.1) enables one to see that OIC explains why you have no obligation to pay the check (since you *cannot* pay) but the new formulation does not enable one to see this (since you *can* actualize the satisfaction proposition of such an obligation, by allowing your friend to pay). I reply that you *cannot* (actualize the proposition that you) pay the check *yourself*, so both the old and the new formulation enable one to see that OIC explains why you have no obligation to *pay the check yourself*. (Moreover, assuming that by—asking and then—allowing your friend to pay you *can* actualize the proposition that the check *is paid*, OICS—see §2.3—is compatible with the claim that you have an obligation to *the effect that the check be paid*.)

¹⁰ One might disagree with a component of the above explanation, namely with my claim that after 4pm you can no longer satisfy the obligation to call me by midnight: one might insist that after 4pm you can still satisfy this obligation (since you can still call me by midnight), so one might reject my definition of the expression “can satisfy”. In reply, I can just replace my new formulation of OIC in the text with the following formulation: “If an agent at a given time has an obligation to do something, then the agent at that time can actualize the satisfaction proposition of the obligation”. This formulation of OIC does not contain the expression “can satisfy”, and still enables one to see that OIC explains why after 4pm you no longer have the obligation to call me by midnight: after 4pm you can no longer actualize the satisfaction proposition of this obligation (since this proposition is already settled). For convenience, I will continue to use the formulation of OIC in the text.

¹¹ I do not understand an obligation to make it the case that you lock the door as satisfied only if you *get yourself* (through self-manipulation) to lock the door (contrast Baier 1970: 653); I understand it instead as satisfied exactly if you make it the case that you lock the door (which happens even if you lock the door directly, without self-manipulation), and thus exactly if you lock the door.

¹² One might claim that in practice our obligations are almost always conditional. For example, if you promise to lock the door, typically you do not acquire an unconditional obligation to lock the door; you acquire instead a *conditional* obligation to lock the door *if*, for example, no bomb destroys the door before you have a chance to lock it. I reply that I need not deny this claim: its truth would render even more urgent and important my project of reaching an ought-implies-can principle that is informative about conditional obligations.

¹³ Cf. McNamara 2004: 121. One might similarly argue that, contrary to what I claimed, your obligation to lock the door is not an obligation to make it the case that you lock the door: if someone gives you a drug that makes you lock the door, then (1) your obligation to lock the door is satisfied (since you lock the door) but (2) an obligation to make it the case that you lock the door is not satisfied (since you do not make it the case that you lock the door). In reply, distinguish two ways in which the drug might work. First, suppose the drug makes you *decide* to lock the door (and then you lock it as a result of this decision). Then you do make it the case that you lock the door: the drug causes

obligation to make it the case that you attend, and is clearly not an obligation to make anything else the case either, it is not an obligation to do.¹⁴ But are there any circumstances in which you would have an obligation to attend my wedding, as opposed to an obligation to make it the case that you attend? There are: if you have promised to attend my wedding and you have explicitly distinguished your promise from a promise to make it the case that you attend, then arguably you have an obligation to attend—assuming that you can attend.¹⁵ But what if you cannot attend? Then, according to the *spirit* of OIC, you have no obligation to attend. OIC, however, fails to deliver this result, since OIC is formulated only in terms of obligations to do. These considerations lead me to propose a generalization of OIC to all unconditional obligations, namely the *ought-implies-can-satisfy* principle:

(OICS) If an agent at a given time has an unconditional obligation, then the agent at that time can satisfy the obligation.

I explain in §3.1 why I formulate OICS only in terms of *unconditional* obligations. OICS entails OIC: since all obligations to do are unconditional obligations, if something holds (as per OICS) for all unconditional obligations, then it also holds (as per OIC) for all obligations to do. If one accepts OIC, to accept OICS one need not accept that some unconditional obligations are not obligations to do: it is enough to accept instead that, *if* some unconditional obligations are not obligations to do, then they are related to abilities in the same way in which, according to OIC, obligations to do are related to abilities.¹⁶

you to make it the case (and thus *you* make it the case) that you lock the door. But then (2) is false. Second, suppose the drug makes your hand move so as to lock the door against your will. Then, as in a case in which someone forcibly moves your hand, it is false that *you* lock the door. But then (1) is false. So (1) and (2) are not both true, and the above argument fails.

¹⁴ For another example of an unconditional obligation which is not an obligation to do, suppose you have an (unconditional) obligation O_1 *not* to kill me, understood as satisfied exactly if it is not the case that you kill me. This is not an obligation O_2 to *make it the case* that you do not kill me: if I kill you and thus you do not kill me, then O_1 is satisfied (since it is not the case that you kill me) but O_2 is not satisfied (since you do not make it the case that you do not kill me; by contrast, if you *refrain* from killing me, then you do make it the case that you do not kill me). For further examples of obligations *not* to do, see notes 21 and 31.

¹⁵ My distinction between obligations to do and unconditional obligations that are not obligations to do is the distinction between agential and non-agential unconditional obligations (cf. McNamara 2004: 120-3): an unconditional obligation is *agential* exactly if its satisfaction proposition is a proposition to the effect that an agent makes something the case, and is *non-agential* otherwise. This distinction differs from the distinctions between “ought to do” and “ought to be” (*Tunsollen* vs *Seinsollen*; cf.: Castañeda 1970: 449-52; García 1986; Geach 1982/1991: 35-8; Hansson 2013: 197; Harman 1986: 131-2; Horty 2001: chap. 3; Krogh & Herrestad 1996; Montefiore 1958: 30-1; von Wright 1963: 112, 1981a: 9, 1981b: 409-11) and between “oughts” that are *prescriptive* vs *evaluative*, *praxiological* vs *axiological* (Guendling 1974: 122-3), *agent-implicating* vs *situational* (Humberstone 1971: 8), *deliberative* vs *judgmental* (Thomason 1981: 179-80), and *action-guiding* (Hansson 1999) vs *ideal* (Copp 1997: 446; Haji 2002: 15; Mason 2003: 319; Robinson 1971: 193-6; cf. Haines 1972), *ideal-expressing* (McConnell 1989: 438; Stocker 1971: 304; cf.: Moore 1922: 315-23; Sidgwick 1907/1981: 33), or *Utopian* (Brown 2000: 103-4).

¹⁶ When I say that they are related to abilities in the same way, I do not just mean that, for example, an obligation to attend my wedding and an obligation to make it the case that you attend are related in the same way to *corresponding* abilities; I also mean that, according to a plausible argument, they are related in the same way to the *very same* ability. Here is the argument: (1) you can satisfy an obligation to attend my wedding exactly if (2) you can make it the case that you attend, and thus (see below) exactly if (3) you can make it the case that you make it the case that you attend, and thus exactly if (4) you can satisfy an obligation to make it the case that you attend. ((2) is equivalent to (3) on the plausible assumption that any proposition p to the effect that you do something—i.e., that you make something the case—is equivalent to the proposition that you make p the case (cf. the “stit paraphrase thesis” in

Here are three observations in support of OICS. (1) Like OIC, OICS explains why at any time after midnight you have no obligation to call me by midnight: because at any time after midnight you cannot make it the case that you call me by midnight (§2.1). (2) Unlike OIC, OICS explains why today you have no obligation to the effect that the sun rises tomorrow (understood not as an obligation to do, but instead as an unconditional obligation satisfied exactly if the sun rises tomorrow): because today you cannot make it the case that the sun rises tomorrow. (OIC, by contrast, explains why today you have no obligation to *make* the sun rise—i.e., to make it the case that the sun rises—tomorrow; OICS also explains this.) (3) Unlike OIC, OICS explains why normally you have no obligation to feel grateful to me right away (understood not as an obligation to do, but instead as an unconditional obligation satisfied exactly if you feel grateful to me right away): because normally you cannot make it the case that you feel grateful to me right away.¹⁷ (OIC, by contrast, explains why normally you have no obligation to *make* yourself feel grateful to me right away; OICS also explains this.) For further discussion and references, see Vranas 2007: 174-5, 199-200 n. 7.¹⁸

Belnap et al. 2001: 7-8); in particular, the proposition that you make it the case that you attend is equivalent to the proposition that you make it the case that you make it the case that you attend.) Given the result that (1) is equivalent to (4), one might ask: is it plausible to propose (as OICS does) the same condition (namely the condition of having the ability to attend my wedding) as necessary both for having an obligation to attend my wedding and for having an obligation to make it the case that you attend? I reply that it *is* plausible: if you cannot attend my wedding, then you have neither an obligation to attend nor an obligation to make it the case that you attend. To propose the same condition as necessary for having either obligation is not to confuse the two obligations.

¹⁷ Admittedly, in some cases (for example, if you know that I have just saved your daughter's life) it is natural to say that you *ought* to feel grateful to me right away. Arguably, however, to say this is not to ascribe an obligation; it is instead to say that it would be fitting or good if you were to feel grateful to me right away. One might respond that, if you fail to feel grateful to me right away, it makes sense for me to blame (and resent) you, and this indicates that you violated an obligation (cf. McHugh 2012: 87). In reply, I suggest that blaming you makes sense not because you violated an obligation to feel grateful to me, but rather because you violated in the past an obligation to cultivate a disposition to feel gratitude whenever appropriate. This suggestion (in conjunction with OICS) also explains why it makes no sense to similarly blame a person who, due to an irreversible biochemical imbalance, is unable to ever feel gratitude (§1): such a person had no obligation to cultivate a disposition to feel gratitude. Adams (1985) disagrees with my suggestion: "what we chiefly blame in the present immoral state of mind is not the imprudence of the previous voluntary omissions" (1985: 14). Adams suggests instead that blaming you can make sense even if you violated no obligation at all: "If someone says to me that I am incapable of feeling gratitude, ... this claim about my feelings ... is already an ethical indictment There is no need to search for guilty actions or omissions of which I may be accused" (1985: 13). I reply first that Adams's suggestion does not explain why it makes no sense to blame a person who is unable to ever feel gratitude. Moreover, Adams's suggestion would support my main point that (contrary to what McHugh 2012 presupposes) it can make sense to blame you (if you fail to feel grateful to me right away) even if you had no obligation to feel grateful to me.

¹⁸ Similar points to those I made in note 17 can be made about cases in which it seems that an agent has an obligation which (is not an obligation to feel, but instead) requires for its satisfaction that the agent have a certain motive or feeling. For example, one might argue that one's obligation to be a good parent is satisfied only if one acts with certain motives (cf. Card 1988: 117; see also note 6). For another example, King (2014) argues that one's obligation to apologize is satisfied only if one says that one is sorry *while feeling sorrow or regret*: if a child said she was sorry "but sneakily thought to herself that she was really fooling Grandma, she was merely pretending to apologize without actually doing so" (2014: 318). As King in effect recognizes, proponents of OICS can say that "the sulky child, if she's sufficiently and irreparably sulky, is under no obligation to apologize. She has, at best, only an obligation to utter the words 'I'm sorry' [and] may have an additional obligation to try to improve her attitude" (2014: 326). King responds: "This is already a costly counterintuitive conclusion. It seems like her being *irreparably* sulky shouldn't exempt her from obligations that the *less* sulky incur" (2014: 326). In reply, I ask: in what sense is the child "irreparably" sulky? If it is in the sense of having an irreversible biochemical imbalance that prevents her from ever feeling sorrow or regret, why shouldn't this exempt her from obligations that the less sulky incur? But if it is instead in the

2.4. OICS and the loss of obligations

Just as one can *acquire* (i.e., start having) obligations, for example by making promises, one can *lose* (i.e., stop having) obligations, for example by being released from promises. OICS provides a sufficient condition for losing an unconditional obligation: becoming unable to satisfy it. More precisely, say that an agent *loses* an obligation at time t exactly if (1) *right before* t (i.e., at every time before t in some open time interval that includes t) the agent *has* the obligation but (2) *right after* t (i.e., at every time after t in some open time interval that includes t) the agent *does not have* the obligation. Similarly, say that at time t an agent *becomes unable* to satisfy an obligation exactly if (1') *right before* t the agent *can* satisfy the obligation but (2') *right after* t the agent *cannot* satisfy the obligation.¹⁹ OICS entails that, if at t an agent becomes unable to satisfy an unconditional obligation that the agent has right before t , then the agent loses the obligation at t .²⁰

In addition to the *timeless* concepts of *being* satisfied or violated that I have been using so far, it is useful to have the following *time-indexed* concepts of *becoming* satisfied or violated at a given time: say that a proposition *becomes settled* at a given time exactly if it is not settled at any earlier time but it is settled at every later time, and say that an obligation *becomes satisfied (or violated)* at a given time exactly if its satisfaction (or violation) proposition becomes settled at that time. For example, if you first visit me at 2pm, then the satisfaction proposition of your obligation to visit me (i.e., the proposition that you visit me) becomes settled at 2pm, so your obligation becomes satisfied at 2pm.²¹ For another example, if you burn my diary at 2pm, then the violation proposition of your obligation *not* to burn my diary (i.e., the proposition that you burn my diary) becomes settled at 2pm, so your obligation becomes violated at 2pm. It is a consequence of OICS that in both of these examples you no longer have the obligation after 2pm (since, as I explained in §2.2, if the satisfaction or the violation proposition of an obligation is settled at a

sense of being *extremely* sulky, then strictly speaking she *can* (with great effort, which may take some time) bring herself to recognize that she behaved inappropriately and as a result feel sorrow or regret, so the alleged counterexample to OICS fails (unless one—implausibly—insists that the child's obligation to apologize as understood above, namely as distinct from an obligation to just say something like "I'm sorry", is an obligation to *quickly* apologize).

¹⁹ To avoid artificial precision, these definitions leave it open whether the agent has or can satisfy the obligation *exactly* at t . Nevertheless, it might still be considered artificially precise to say that an agent loses or becomes unable to satisfy an obligation at a specific time instant. To alleviate this worry, one can similarly define what it is to lose an obligation over (for example) a closed time interval T (rather than at a specific time instant t). For that purpose, define "right before (after) T " as "at every time earlier (later) than all members of T in some open time interval that has T as a subset". I ignore this complication in the text.

²⁰ If the qualification "that the agent has right before t " is omitted, one gets the *false* claim that, if at t an agent becomes unable to satisfy an unconditional obligation, then the agent loses the obligation at t . To see that this claim is false, suppose that at 1pm you promise to meet me tomorrow, at 2pm you are released from your promise, and at 3pm you become unable to meet me tomorrow. Then you lose your obligation (to meet me tomorrow) at 2pm, *before* you become unable (at 3pm) to satisfy it. (Right before 3pm you can still satisfy the obligation, although you no longer have it: you can still actualize its satisfaction proposition, namely the proposition that you meet me tomorrow.)

²¹ Some obligations are such that it is impossible for them to ever *become* satisfied (even if they *are* satisfied). An example is your obligation to never torture. Indeed, it is impossible for the proposition that you never torture to ever become settled: necessarily, for any time t , the history of the world up to and including t does not logically entail that you never torture (because it does not logically entail that you never torture *after* t ; even if you are dead at t , it is logically possible that you rise from the dead and you torture after t). Nevertheless, you *can* satisfy this obligation: you can actualize the proposition that you never torture, assuming that you can refrain from ever torturing.

given time, then the agent at that time cannot satisfy the obligation). One might object to OICS by claiming that, although indeed you no longer have an obligation after it becomes satisfied, you keep having your obligations after they become violated: after 2pm, you still have the obligation not to burn my diary (the very diary that you have burned). In reply, I ask: if you still have the obligation not to burn my diary after you burn it, when do you lose this obligation? It would be implausible to claim that you still have this obligation on your deathbed.²² Moreover, if you keep having your obligations after they become violated, then you have obligations to do things *in the past*: for example, in 2015 you still have the obligation to meet me for lunch in 1985, an obligation that became violated when you failed to show up. But this is implausible if (as I assume in this paper) backwards causation is impossible. I conclude that the objection to OICS fails.

In the examples in the previous paragraph, your obligations (to visit me and not to burn my diary) are lost approximately when they become satisfied or violated. (See the last paragraph of §2.5 below for an explanation of “approximately”.) In other examples, however, an obligation is lost *long before* it becomes satisfied or violated. To see this, suppose that right before noon you are in Detroit and you have an obligation—because you have promised—to attend my wedding (which is scheduled to start in Boston at 4pm), but at noon you miss your flight from Detroit and thus you become unable to attend my wedding (because the Detroit airport is shutting down due to a major snowstorm, and no further flights are scheduled to depart from it—or from any other airport that you can reach on time—for the rest of the day). Although your obligation to attend my wedding does not become violated (i.e., its violation proposition does not become settled) until 4pm, it is a consequence of OICS that you lose the obligation at noon. This consequence of OICS is controversial (especially if you miss your flight on purpose, to avoid attending my wedding), but I have in effect provided an extensive defense of it in previous work (Vranas 2007: 175-82). I summarize that defense in a note;²³ here I bolster that defense by addressing an objection to my claim that at noon you become unable to attend my wedding.

²² Moreover, the view under consideration does not explain why you no longer have this obligation *after* you die. OICS, by contrast, explains why dead people have no obligations: because dead people cannot actualize the satisfaction proposition of any obligation.

²³ My defense has two components: (1) I give two arguments against the competing claim that you lose the obligation at 4pm, and (2) I reply to objections against my claim that you lose the obligation at noon. (1a) It is implausible to claim that right before 4pm you still have the obligation to be in Boston at 4pm: one nanosecond before 4pm, it is logically *possible* for you to move faster than light so as to be in Boston at 4pm, but such *exotic* possibilities are irrelevant to whether you have the obligation. (1b) The above competing claim presumably relies on the view that, putting cases of revocation (e.g., being released from a promise) aside (see note 20), you never lose an obligation before it becomes satisfied or violated. But this view is false, as one can see by considering examples in which you acquire a compensatory obligation before the original obligation becomes satisfied or violated: if at 1pm someone destroys a rare book that you have promised to return to me by 4pm and at 2pm you become able to buy the only other existing copy of the book, then shortly after 2pm you have the obligation to buy the other copy and give it to me by 4pm, so you do not *also* have the obligation to return my copy of the book by 4pm. (2a) It seems implausible to claim that you can get rid of unwanted obligations by making yourself unable to satisfy them; e.g., by missing your flight on purpose. I reply that, if you kill your teacher, then you get rid of the obligation to turn in a paper to your teacher; so it is definitely possible to get rid of unwanted obligations. (When you do so, however, usually you acquire compensatory obligations, so my view does not let people off the hook. For example, it is plausible that at noon you acquire an obligation to *try* to become again able to attend my wedding.) (2b) If at 4pm you no longer have the obligation to attend my wedding, then failing to attend my wedding is not impermissible for you at 4pm, so why is it natural to say that (assuming you missed your flight on purpose) you are blameworthy *for failing to attend* my wedding? Because, I reply, failing to attend my wedding is a bad consequence of missing your flight (which is

Here is the objection: it is *false* that at noon you become unable to attend my wedding, because it is *possible* that, at 1pm (after you miss your flight at noon), you are unexpectedly offered a ride to Boston in a military plane specially equipped to fly in a snowstorm. I reply that mere possibility is not enough for ability: although it is possible that tomorrow you will win ten million dollars in the lottery, the fact remains that *today* you cannot donate ten million dollars to charity next week. In response, one might ask: what if at 1pm you are *in fact* offered—and you accept—a ride to Boston in a military plane, so that you in fact attend my wedding? How can you be unable to attend if you in fact attend? In reply, I can grant (though see Southwood & Wiens 2016) that if you do something at a given time then at *some* earlier time you can do it. It does not follow, however, that at *every* earlier time you can do it: in some cases you *become able* to do things. (For example, today you are not in a position to—i.e., you cannot—donate ten million dollars to charity next week, but tomorrow you become able to do so because you win the lottery.) If you attend my wedding at 4pm, then at *some* earlier time (shortly after 1pm) you can attend; but this is compatible with my claim that at some even earlier time (shortly after noon) you cannot attend.

If at 1pm you are in fact offered a ride to Boston in a military plane and as a result you become again able to attend my wedding, then it is plausible that at 1pm you acquire again an obligation to attend, since the impediment to your having such an obligation (namely your inability to attend) is removed. (I do not need to take a stand in this paper on whether at 1pm you reacquire the *same* obligation to attend that you lost at noon or you acquire instead a *different* obligation to attend; in Vranas 2017b, I defend the former view.) But why exactly do you acquire again at 1pm an obligation to attend? Because your promise to attend my wedding is still in effect. Compare: if you are a soldier and at 9am your commanding officer orders you to run in tomorrow’s marathon but you cannot do so because you are paralyzed, then at 9am you acquire no obligation to run; but if at 1pm you suddenly recover from your paralysis and thus you become able to run in tomorrow’s marathon, then at 1pm you acquire an obligation to run because the officer’s order is still in effect.²⁴

2.5. The “ought-implies-possible-violation” principle (OIPV)

It turns out that OICS does not *always* explain why agents lose unconditional obligations. To see this, suppose that I am too squeamish to euthanize my terminally ill dog, so I ask you to do it. At 9am, you promise me that the dog will be dead by midnight, and as a result you acquire a corresponding obligation.²⁵ At noon, you administer to the dog a quick-acting and invariably fatal

blameworthy because right before you miss it you have an obligation not to miss it), and it is natural to hold people blameworthy for the bad consequences of their blameworthy actions.

²⁴ One might argue that this explanation of why at 1pm you acquire an obligation to run (or to attend my wedding) is incomplete: a complete explanation (one might continue) must also appeal to the fact that at 9am you acquire a *conditional* obligation to run *if* at some time before the marathon starts you become able to run. I reply that it is more plausible to claim instead that at 9am you acquire a conditional obligation to run if at some time before the marathon starts you become able to run *and you keep being able to run* until the marathon starts. But the fact—if it is a fact—that you acquire such a conditional obligation (and you still have it at 1pm) does not help explain why *at 1pm* you acquire an obligation to run. See also note 59.

²⁵ I assume that, right before you make the promise, it is all-things-considered permissible for you to make and keep the promise, so the promise is not “wicked” (Altham 1985; cf. Smith 1997). I also assume that the promise is not obtained by coercion or deception (cf. Owens 2007). I make similar assumptions concerning promises throughout this paper.

poison to which there is no antidote, with the result that the dog will inevitably be dead within a few minutes (and thus by midnight). Then you lose the obligation at noon: you are done with it (i.e., there is nothing more you need to do—or to refrain from doing—to ensure that it becomes satisfied). Indeed, if you do not lose the obligation at noon, when do you lose it? One might argue that you lose it when the dog dies, presumably because right before the dog dies it is still possible—but right after the dog dies it is no longer possible—that the dog miraculously recovers from the poisoning. I reply that such *exotic* possibilities are irrelevant to whether you have the obligation (see Vranas 2007: 177-8).²⁶ Now given that you lose the obligation at noon, OICS does not explain why you lose it, since you do not become unable to satisfy it: after you administer the poison but before the poison acts, you can kill the dog by shooting it, so you can still satisfy the obligation (i.e., you can still make it the case that the dog is dead by midnight).²⁷

One might propose that you lose the obligation in the above example because, when you administer the poison to the dog, you become *unable to violate* the obligation: you become unable to make it the case that the obligation is violated, namely that the dog is *not* dead by midnight. So one might propose the *ought-implies-can-violate* principle:

(OICV) If an agent at a given time has an unconditional obligation, then the agent at that time can violate the obligation.

I reply that OICV is false. To see this, suppose your teenage daughter is considering whether to participate in a psychological study at school, and you can make it the case that she does not participate because you can withhold your consent: participation requires both your consent and her own (and this requirement is strictly enforced). Then, if you promise your mother that your daughter will *not* participate in the study, right after you promise you have a corresponding obligation (as evidenced by the fact that right after you promise you have a reason to withhold your consent). You can satisfy this obligation by withholding your consent, but you cannot violate the obligation: you cannot make it the case that your daughter participates in the study, assuming

²⁶ See also note 23. In response, one might ask: if such possibilities are “exotic”, then why do you acquire at noon an obligation to wait and make sure that the dog dies by midnight? I reply that, if (as I may assume, since I am talking about *objective* obligations) you *know* that the poison is quick-acting and invariably fatal and that there is no antidote to it, then you acquire no obligation to wait and make sure that the dog dies by midnight.

²⁷ Here is also a second example. Suppose that the safe in your home can be opened only by *both* turning a key and dialing a combination. At 11:30am, you are at home, and your spouse calls you and tells you that a thief has obtained the combination; your spouse’s key is secure, but the thief is now driving to your home to get your key and open the safe. Your spouse has called the police, but they will only arrive shortly before noon. You promise your spouse that *either the thief will not get your key until noon or you will change the combination before noon (or both)*, and as a result you acquire a corresponding obligation. As soon as you hang up, you destroy your key by dissolving it in a very potent acid. When you destroy your key, you lose the above obligation: you are done with it (i.e., there is nothing more you need to do—or to refrain from doing—to ensure that it becomes satisfied). But OICS does not explain why you lose the obligation, since you do not become unable to satisfy it: after you destroy your key but before noon (and before the thief arrives), you can still satisfy the obligation by changing the combination. (Suppose that, whenever the combination is changed, it cannot be changed again for the next 24 hours. If you destroy your key and then change the combination before noon, your obligation becomes satisfied when you change the combination, not when you destroy your key: when you change the combination, it becomes settled that you change the combination before noon, but when you destroy your key, it does not become settled that the thief will not get your key until noon, since it is logically *possible* that the constituent parts of your key will spontaneously reassemble.)

that you have no influence over whether she decides to give her own consent. So OICV is false.²⁸

I propose instead that you lose the obligation in the dog example because, when you administer the poison, it becomes *inevitable* for you that the dog will be dead by midnight. More precisely, say that a proposition is inevitable for an agent at time t exactly if it is *ability-necessary* for the agent at t : it *would* be true no matter what the agent were to do, among the things that at t the agent can do.²⁹ I propose, then, that you lose the obligation because its satisfaction proposition—namely the proposition that the dog is dead by midnight—becomes ability-necessary for you; equivalently, its violation proposition becomes ability-impossible for you. So I propose the *ought-implies-possible-violation* principle:

(OIPV) If an agent at a given time has an obligation, then the violation proposition of the obligation is *ability-possible* for the agent at that time (i.e., it *might* be true if the agent were to do something that at the given time the agent can do).³⁰

²⁸ To my knowledge, this is a novel counterexample to OICV. On variants of OICV, see: al-Hibri 1978: 14-6; Anderson 1956/1967: 181-3, 196-9; Åqvist 1984: 616-7, 1987: 19-20; Bailhache 1991: 17-9, 23-4; Brown 1977: 211-2; Bykvist 2007: 113-4; Copp 1997: 455 n. 8; Dahl 1974: 488-9, 499, 510 n. 18; Feis 2014; Feldman 1988a: 241-2; Føllesdal & Hilpinen 1971: 13; Haji 1994: 120-1, 2002: 28-31, 2012: 53-68; Henderson 1966: 111; Kekes 1984: 459; Loewer & Belzer 1983: 306; Narveson 1978: 44; Nelkin 2011: 102-3; Oppenheim 1987: 376; Prior 1958: 138, 1962: 221-2; Stenius 1963: 253; Stocker 1971: 304-5; Taylor 1967: 89; Tranøy 1972: 121; van Rijen 1993: 265-6; von Wright 1951: 10-1, 1981a: 8, 1981b: 402, 1988: 60; Zimmerman 1993: 52-3, 1996: 81-2, 89-90.

²⁹ (1) In the definition of ability-necessity, the clause “among the things that at t the agent can do” is not redundant. To see this, suppose that your parents would die in the next minute no matter what you were to do, among the things that now you can do: they are about to be crushed by a boulder which is too heavy for you to deflect. Then the proposition that your parents will die in the next minute—call this proposition P —is ability-necessary for you now, although it is false that, no matter what you were to do, P would be true: P might be false if you were to do something that now you *cannot* do, namely deflect the boulder. (2) Ability-necessity (in contrast to what I call “practical necessity”: see Vranas 2016: 1743) *does not entail truth*. To see this, suppose that the only thing you do is something that now you *cannot* do, and as a result P is false (e.g., the only thing you do is deflect the boulder: you suddenly acquire superhuman strength, you deflect the boulder, and then you die). This supposition (and thus the claim that P is false) is compatible with the claim that, if you were to do anything that now you *can* do, P would be true (and thus is compatible with the claim that P is ability-necessary for you now). (3) Ability-necessity is relative to *agents*. To see this, suppose that now I can deflect the boulder, so that P is not ability-necessary *for me* now. Still, P may be ability-necessary *for you* now, because I would not deflect the boulder no matter what you were to do (for example, because deflecting the boulder would kill my children), among the things that now you can do. (4) Ability-necessity is also relative to *times*. To see this, suppose that a second from now you will suddenly acquire superhuman strength and will become able to deflect the boulder. Then P , which is ability-necessary for you *now*, will not be ability-necessary for you *a second from now*. If something like this happens in the dog example (e.g., if someone miraculously synthesizes an antidote and gives it to the dog), I think that you reacquire the obligation, as in the last paragraph of §2.4.

³⁰ If there is nothing that at time t an agent can do (for example, because at t the agent is in a persistent vegetative state, or because at t the agent is not yet born or is already dead), then (according to my definitions in the text) every proposition is ability-necessary but no proposition is ability-possible for the agent at t ; this is counterintuitive, and also has the consequence that ability-necessity does not entail ability-possibility. To avoid these problems, I propose the following official definitions. (1) A proposition P is *ability-necessary* for an agent at time t exactly if both (a) there is something that at t the agent can do and (b) P would be true no matter what the agent were to do, among the things that at t the agent can do. (2) A proposition P is *ability-possible* for an agent at time t exactly if either (a) there is nothing that at t the agent can do or (b) P might be true if the agent were to do something that at t the agent can do. According to these definitions, ability-necessity entails ability-possibility, and if there is nothing that at t an agent can do, then every proposition is ability-possible but no proposition is ability-necessary for the agent at t . In the text, I stick with the simpler definitions, since I only examine cases in which there is something that at a given time an agent can do. If the negation of a “would” counterfactual is a “might” counterfactual with the same anteced-

I explain in §3.1 why I formulate OIPV (in contrast to OICS) in terms of *all* obligations that agents have, including conditional obligations. My objection to OICV leaves OIPV unscathed: although (now) you cannot make it the case that your daughter participates in the study, the proposition that your daughter will participate is ability-possible for you (since it might be true if you were to do something that you can do, namely give your consent), so OIPV (in contrast to OICV) does not conflict with the claim that, if you have promised that she will not participate, then you have a corresponding obligation.³¹ Note that OIPV does not render OICS redundant: OICS explains why you have no obligation to win the lottery (since you cannot make it the case that you win, assuming that you have no influence over the outcome of the lottery), but OIPV does not explain this (since the proposition that you do *not* win the lottery is ability-possible for you).³²

Here is an objection both to OIPV and to OICS. Suppose that at 9am you promise me to attend my wedding at noon, and as a result you acquire a corresponding obligation (which is satisfied exactly if you attend—i.e., you are present at—my wedding at noon, even if you do not make it the case that you attend). Suppose further that at 11am it becomes inevitable for you both that you attend my wedding at noon and that you do not make it the case that you attend: you are kidnapped, knocked unconscious, brought to my wedding, and woken up at noon. Since you do attend my wedding at noon, your obligation to attend becomes satisfied at noon, so according to the objection you have this obligation until noon. But you lose the obligation at 11am according to OIPV: its satisfaction proposition (namely the proposition that you attend my wedding at noon) becomes inevitable for you. And you also lose the obligation at 11am according to OICS: you become unable to satisfy it (i.e., to make it the case that you attend my wedding at noon). I reply that indeed you lose the obligation at 11am: it is fallacious to infer from the premise that the obligation becomes satisfied at noon the conclusion that you have the obligation until noon. To see the fallacy, suppose that at 10am I *release* you from your promise to attend my wedding at noon, and as a result you lose the corresponding obligation. Then you do *not* have the obliga-

ent and negated consequent (Lewis 1973/1986: 8, 1973: 21), then ability-necessity and ability-possibility are inter-definable in the usual way (this also holds for the simpler definitions in the text).

³¹ One might object to OIPV as follows. Suppose that at 9am you are in a room that has no windows, and you promise that you will not leave the room until noon, acquiring as a result a corresponding obligation. At 11:55 am, you are still in the room, but someone welds the only door of the room to the frame of the door. Then, one might argue, at 11:56am it is inevitable for you that you will not leave the room until noon, but you still have the obligation not to leave the room until noon, so OIPV is false. In reply, I claim that you lose the obligation at 11:55am: there is nothing more you need to do—or to refrain from doing—to ensure that it becomes satisfied (e.g., you no longer need to restrain any urge to leave the room). In support of my claim, note the implausibility of the competing claim that you lose the obligation at noon: this competing claim presumably relies on the observation that before noon it is logically *possible* that (for example) you suddenly acquire superhuman strength and force the door open, but such *exotic* possibilities are irrelevant to whether you have the obligation (cf. note 23).

³² Here is another objection to OIPV. Suppose you promise me that (*P*) *either you will exercise today or you will not run faster than light tomorrow*. OIPV entails that you do not acquire a corresponding obligation, since the proposition that you will not run faster than light tomorrow—and thus also *P*—is inevitable for you. One might object that, since you *can* actualize *P* (because you can exercise today) and you promise to do so, you do acquire an obligation to actualize it—an obligation that you can satisfy by exercising today. I reply that, if your promise is to *actualize P* (an *agential* promise), OIPV does not entail that as a result of that promise you do not acquire a corresponding obligation (since it is not inevitable for you that you actualize *P*: it is not inevitable for you that you exercise today). But if your promise is instead (as supposed above) that *P is true* (a *non-agential* promise), then it is correct that (as OIPV entails) you do not acquire a corresponding obligation, since your promise will become satisfied no matter what you do.

tion until noon (assuming that you do not reacquire it before noon), even if you attend my wedding at noon and thus the obligation becomes satisfied at noon (i.e., its satisfaction proposition becomes settled at noon).³³

One might modify the example in the previous paragraph so as to propose a “Frankfurt-type” objection to OIPV (similar to an objection to OIC I examined in Vranas 2007: 195-6). Suppose again that at 9am you promise me to attend my wedding at noon, and as a result you acquire a corresponding obligation. Suppose further that at 11am it becomes inevitable for you that you attend my wedding at noon: a “counterfactual intervener” starts monitoring your activities, and would kidnap you and bring you to my wedding at noon were you not to come on your own initiative. According to OIPV, you lose your promissory obligation at 11am. But this is implausible if (as one may also suppose) you in fact attend my wedding at noon on your own initiative, without any interference from the counterfactual intervener—or so the objection goes. I reply first that the information that the intervener does not actually intervene is a red herring: as the possibility of being *released* from your promise shows (see my explanation in the previous paragraph), from the claim that in fact you attend my wedding at noon on your own initiative it does not follow that you have the promissory obligation until noon. Moreover, I suspect that the apparent implausibility of the claim that you lose the promissory obligation at 11am is largely due to the implicit assumption that you are unaware of the presence of the intervener. But the question is whether you lose an *objective* obligation (§2.1), so I may assume instead that you *know* that the intervener arrives on the scene at 11am. Then it should be clear that indeed you lose the promissory obligation at 11am: you are done with it (i.e., there is nothing more you need to do—or to refrain from doing—to ensure that it becomes satisfied), since it is inevitable for you that it will become satisfied (even if you do not *make it the case* that you attend).

If OICS and OIPV are both true, then obligations are lost at the latest shortly *before* they become satisfied or violated. To see this, suppose that at 9am you have an obligation to demolish building A by noon, and you also have an obligation *not* to demolish building B until noon. (1) Suppose further that you demolish both buildings by pushing at 11am a button that detonates explosives placed inside the buildings. After you push the button but before the buildings are demolished (and thus before your obligations become satisfied or violated), it is inevitable for you that you demolish both buildings by noon (assuming that you cannot stop the signal that is traveling from the button to the explosives). So, according to OIPV, you no longer have the obligation to demolish building A by noon (since the satisfaction proposition of this obligation is inevitable for you). Moreover, according to OICS, you no longer have the obligation *not* to demolish building B until noon (since you can no longer satisfy this obligation). (2) Suppose alternatively that you demolish neither building until noon. A few nanoseconds before noon (and thus before your obligations become satisfied or violated), it is inevitable for you that you demolish neither building until noon (assuming that pushing the button now would only demolish the buildings shortly

³³ (1) One might claim that, if you do not have the obligation until noon, it is misleading to say that the obligation becomes satisfied at noon. I reply that, if by saying that the obligation becomes satisfied at noon one presupposes that you have the obligation until noon, then it is question-begging to propose the above objection to OIPV and to OICS: proponents of the objection infer from the premise that you attend my wedding at noon the conclusion that the obligation becomes satisfied at noon, and thus *presuppose* that you have the obligation until noon. (2) A variant of the above objection to OIPV and to OICS supposes that at 9am, when you promise me to attend my wedding at noon, it is *already* inevitable for you both that you attend my wedding at noon and that you do not make it the case that you attend. But then, I reply, by promising to attend you do not acquire an obligation to attend: it is like promising that the sun will rise tomorrow.

after noon). So, according to OIPV, you no longer have the obligation *not* to demolish building B until noon (since the satisfaction proposition of this obligation is inevitable for you). Moreover, according to OICS, you no longer have the obligation to demolish building A by noon (since you can no longer satisfy this obligation). I think that the reasoning in this example generalizes: as far as I can see, if OICS and OIPV are both true, there is no case in which an agent loses an obligation at the *same* time at which the obligation becomes satisfied or violated.

2.6. Comparing OICS with OIPV

The two principles I have defended, namely OICS and OIPV, provide two necessary conditions for an agent to have an unconditional obligation at time *t*: (1) at *t* the agent must be able to actualize the satisfaction proposition of the obligation (OICS), and (2) the satisfaction proposition of the obligation must not be ability-necessary (equivalently, the violation proposition of the obligation must be ability-possible) for the agent at *t* (OIPV). Neither condition entails the other. First, (1) does not entail (2): even if at *t* an agent can actualize a proposition, the proposition may be ability-necessary for the agent at *t* (e.g., you can shoot the dog and thus cause its death by midnight, but you have already poisoned the dog and thus its death by midnight is inevitable for you). Second, (2) does not entail (1): even if a proposition is not ability-necessary for an agent at *t*, at *t* the agent may be unable to actualize the proposition (e.g., the proposition that you win the lottery is not inevitable for you, but you cannot actualize it). I propose that both conditions (and both principles) are needed to explain why (or to correctly predict when) agents lose or fail to have unconditional obligations: OIPV explains but OICS does not explain why you lose the obligation in the dog example, and OICS explains but OIPV does not explain why you have no obligation to win the lottery (§2.5).

In some cases, OICS and OIPV provide different explanations of why an agent loses an unconditional obligation. Suppose that you promise your daughter's high school principal that your daughter will not enter the school premises tomorrow (because she has been suspended), and as a result you acquire a corresponding obligation. If your daughter dies today, then you immediately lose that obligation. Why do you lose it? Not because the obligation becomes satisfied today, since it does not: it is logically *possible* that your daughter will rise from the dead and will enter the school premises tomorrow, so the satisfaction proposition of the obligation—namely the proposition that your daughter will *not* enter the school premises tomorrow—is not yet *settled*. OICS provides one explanation of why you lose the obligation: because you become unable to actualize its satisfaction proposition. OIPV provides a different explanation of why you lose the obligation: because its satisfaction proposition becomes ability-necessary for you. I propose that both explanations are correct: this is a case of explanatory overdetermination. If so, then this kind of case does not reveal any conflict between OICS and OIPV.

One might be worried by the *asymmetry* between satisfaction and violation exhibited by the OICS/OIPV pair: OICS combines satisfaction with *ability*, but OIPV combines violation with *ability-possibility*. It will not do to remove the asymmetry by rejecting OIPV and adopting instead (in addition to OICS) OICV: as I argued, OICV is false. But one might propose to remove the asymmetry by rejecting OICS and adopting instead (in addition to OIPV) the *ought-implies-possible-satisfaction* principle:

(OIPS) If an agent at a given time has an unconditional obligation, then the satisfaction proposition of the obligation is ability-possible for the agent at that time.

In reply, note first that OICS entails OIPS. This is because, if an agent at a given time can satisfy an obligation (as per OICS), then the satisfaction proposition S of the obligation is ability-possible for the agent at that time (as per OIPS): S might (indeed, *would*) be true if the agent were to do something that at the given time the agent can do, namely actualize S . So OIPS is true if (as I believe) OICS is true. Nevertheless, it will not do to reject OICS and adopt instead OIPS, because in some cases OICS explains but OIPS does not explain (and neither does OIPV) why an agent fails to have an obligation. For example, OICS explains (as I said) why you have no obligation to win the lottery (since you cannot make it the case that you win the lottery), but OIPS does not explain this (since the proposition that you win the lottery is ability-possible for you: it might be true if you were to do something that you can do, namely buy a lottery ticket).³⁴ To conclude, I see no way to remove the asymmetry between satisfaction and violation exhibited by the OICS/OIPV pair. Later on, however, I argue that the asymmetry is unproblematic (§3.3).

To complete my discussion of unconditional obligations in this section, I address two objections to OICS: one objection in §2.7, and another one in §2.8.

2.7. OICS and obligations partly about the past

One might object to OICS (and to OIC) as follows. Suppose that at 9am a law is enacted that indefinitely forbids gambling starting at 11am, and as a result at 9am you acquire an obligation to never gamble starting at 11am. Suppose further that at 1pm you gamble (for the first time); then your obligation becomes violated (i.e., its violation proposition becomes settled) at 1pm, and as a result at 3pm you can no longer satisfy it. Contrary to OICS, however, at 3pm you still have the obligation, since the law remains in effect—or so the objection goes. In reply, I deny that (1) at 3pm you still have the obligation to never gamble *starting at 11am*. I maintain instead that (2) at 3pm you have an obligation to never gamble *starting at 1pm*; you acquire this obligation shortly before (see §2.5) 1pm, when you become unable to satisfy (and thus you lose) the obligation to never gamble starting at 11am.³⁵ To defend (2), suppose that at 5pm you gamble again (for the second time). Then the law that forbids gambling is broken *twice*: once at 1pm, when you first gamble, and once at 5pm, when you gamble again. Appealing to (2), I propose the following explanation of why the law is broken twice: a law is broken whenever an obligation that arises from the law becomes violated, so the law that forbids gambling is broken both at 1pm, when

³⁴ In response, one might note that the proposition (P) that you *make it the case* that you win the lottery (as opposed to the proposition that you *win* the lottery) is *not* ability-possible for you: there is nothing you can do such that, if you were to do it, P might be true (because, even if you win, it is by chance, so you do not make it the case that you win). So an explanation of why you have no obligation to win the lottery is provided by the following principle: (OIPS*) if an agent at a given time has an unconditional obligation, then it is ability-possible for the agent at that time to make it the case that the obligation is satisfied. In reply, note first that OICS entails OIPS* (and OIPS* entails OIPS; I omit the proofs); so OIPS* is true if (as I believe) OICS is true. Nevertheless, in some cases OICS explains but OIPS* does not explain (and neither does OIPV) why an agent fails to have an obligation. For example, suppose that today you cannot donate ten million dollars to charity next week. Then today you have no obligation to do so, and OICS explains why. OIPS*, by contrast, provides no explanation, because the proposition that you make it the case that you donate ten million dollars to charity next week is ability-possible for you today: it might be true if you were to do something that today you can do, namely buy a lottery ticket (because you might win ten million dollars tomorrow and donate them to charity next week).

³⁵ An alternative view is that you acquire this obligation at 9am, when the law is enacted: for any time t not earlier than 11am, at 9am you acquire an obligation to never gamble starting at t . I believe that this view results in an undesirable proliferation of obligations, but for my purposes in this paper I do not need to insist on this: this view is compatible with my denial of (1) in the text, and thus also provides a response to the objection to OICS.

your obligation to never gamble *starting at 11am* becomes violated, and at 5pm, when your obligation to never gamble *starting at 1pm* becomes violated. Proponents of (1), by contrast, have no explanation of why the law is broken twice, since your obligation to never gamble *starting at 11am* only becomes violated once (namely at 1pm, when its violation proposition becomes settled).³⁶ My explanation works only if it is possible for multiple obligations to arise from a single law, but this is indeed possible: when the law that forbids gambling is enacted, both you and I acquire an obligation to never gamble starting at 11am, and our obligations are distinct because they have distinct satisfaction propositions (your obligation is satisfied exactly if *you* never gamble starting at 11am, and my obligation is satisfied exactly if *I* never gamble starting at 11am).

To further defend my view, consider a different example. Suppose that in December you take out a loan repayable in ten equal monthly installments, due at 5pm on the fifth day of each month starting in January, and as a result you acquire an obligation to pay all ten installments on time. Suppose further that you pay the first but not the second installment on time; then your obligation becomes violated on February 5 (at 5pm, when the second installment was due), and as a result on February 6 you can no longer satisfy it. By analogy with the gambling example, opponents of OICS might argue that (1') on February 6 you still have the obligation to pay all ten installments on time (since the terms of your loan have not changed). I maintain instead that (2') on February 6 you have an obligation to pay the second installment together with a penalty *as soon as possible* (assuming that this is what the terms of your loan specify), as well as an obligation to pay *the remaining eight installments* on time; you acquire these obligations on February 5, when you become unable to satisfy (and thus you lose) the obligation to pay all ten installments on time.³⁷ To argue against (1'), I ask: if on February 5 you do not lose the obligation to pay all ten installments on time, when do you lose it? It will not do to answer that you lose it whenever you *fully* repay the loan (together with any penalties): if you fully repay the loan on March 4, your obligation to pay all ten installments *on time* does not become satisfied on March 4 (since you did not pay the second installment on time), so why would you lose it on March 4?³⁸ It seems then that proponents of (1') have no satisfactory answer to the question of when you

³⁶ More generally, on my definition of becoming violated at a given time (§2.4), it is impossible for an obligation to become violated more than once: necessarily, once the violation proposition of an obligation—or indeed *any* proposition—becomes settled, it remains settled forever, and thus it never becomes settled again. In response, proponents of (1) might reject my definition, and might propose instead that an obligation becomes violated at time *t* exactly if (a) what the obligation requires about *t* is false, but (b) what the obligation requires about times right before *t* is true. On this alternative definition, your obligation to never gamble starting at 11am does become violated again at 5pm: what it requires about 5pm (namely that you fail to gamble at 5pm) is false, but what it requires about times right before 5pm (namely that you fail to gamble at those times) is true. I reply that this alternative definition does not work for obligations that do not require anything about particular times, for example a student's obligation to turn in a paper to a professor *by 5pm*. By contrast, my definition of becoming violated (or satisfied) at a given time is fully general.

³⁷ Alternatively, on February 5 you may acquire a *single* obligation to both (a) pay the second installment together with a penalty as soon as possible and (b) pay the remaining eight installments on time. For my purposes in this paper, it does not matter whether this alternative or what I say in the text is the case.

³⁸ One might answer: because on March 4, since you fully repay the loan, you lose all obligations that arise from the loan agreement (and that you still have right before you fully repay the loan). I reply that this is false. Suppose the loan agreement specifies that, if you fail to pay an installment by the due date, then within two months of the due date you must send your name to a company that publishes lists of “bad payers”. Then on March 4, although you fully repay the loan, you do not lose the obligation to send your name to that company by April 5 (assuming you have not already done so).

lose the obligation to pay all ten installments on time. OICS, by contrast, provides a satisfactory answer: you lose the obligation when you become unable to satisfy it, namely on February 5.

One might object that the above answer is not satisfactory because, if (on February 5 you lose and thus) on February 4 you still have the obligation to pay *all ten* installments on time, then on February 4 you have an obligation to pay the *January* installment on time—an obligation about the past! In reply, go back for a moment to the gambling example. The view that (1) at 3pm you still have the obligation to never gamble *starting at 11am* and the view that (2) at 3pm you have instead an obligation to never gamble *starting at 1pm* agree on this much: at 3pm you have an obligation which is *partly* about the past (more precisely, an obligation whose satisfaction proposition is partly about the past of 3pm). To see that it is indeed possible to have obligations partly about the past, suppose that at 8am you promise your spouse that you will never smoke again, *starting immediately*. It is then plausible that at 8am you acquire an obligation to never smoke starting at 8am. But if it is impossible to have obligations partly about the past, then at any time t after 8am—e.g., one nanosecond after 8am—you have no obligation to never smoke starting at 8am, since such an obligation would be partly about the past of t . But then, although in effect you promise to never smoke starting at 8am, you have *for at most a single time instant* (namely at most at 8am) an obligation to never smoke starting at 8am. I take this to be a *reductio* of the view that (3) it is impossible to have obligations partly about the past. Return now to the loan example and to my claim that (4) on February 4 you still have the obligation to pay all ten installments on time. The fact that (4) contradicts (3) is no problem for (4), since (3) is false. But what about the point that (4) entails the false claim that (5) on February 4 you have an obligation to pay the *January* installment on time? I agree that (5) is false: OICS entails that it is impossible to have unconditional obligations *wholly* about the past (setting the possibility of backwards causation aside, as I do throughout this paper). But I deny that (4) entails (5); more generally, I deny that, (6) necessarily, if at time t you have an unconditional obligation whose satisfaction proposition is P and P entails P' , then at t you have an unconditional obligation whose satisfaction proposition is P' . Indeed, (6) is false: possibly, you currently have an obligation to help me but no obligation to help *or kill* me.³⁹

One might argue that there is another problem with my claim that (4) on February 4 you still have the obligation to pay all ten installments on time: if OICS is true, (4) entails that (7) on February 4 you can actualize the proposition that you pay *all ten* installments on time, and (7) in turn entails that (8) on February 4 you can actualize the proposition that you pay the *January* installment on time—which is clearly false. In reply, I agree that (8) is false, but I deny that (7) entails (8); more generally, I deny that, (9) necessarily, if at time t you can actualize a proposition P and P entails P' , then at t you can actualize P' . Indeed, (9) is false: possibly, today you can actualize the proposition that you help me, but today you cannot actualize the proposition that *either you help me or the sun rose yesterday* (since this disjunction is already settled). The falsity of (9) de-

³⁹ On the other hand, if you currently have an obligation to help me, then it is currently *obligatory* for you to help *or kill* me (and this may account for the intuitive appeal of (6) in the text): in my view, (unconditional pro tanto) obligatoriness corresponds to *necessary* conditions (subject to ability) for the satisfaction of (unconditional pro tanto) obligations (although there is also another kind of obligatoriness—which I call *safe obligatoriness*—that corresponds to conditions which are *both necessary and sufficient* for the satisfaction of obligations). Defending my view about obligatoriness requires an excursion into deontic logic, and thus lies beyond the scope of this paper. For references to “Ross’s paradox” (concerning the move from “help me” to “help or kill me”), see Vranas 2011: 412 n. 52.

fuses the above objection to (7) (and to (4)),⁴⁰ but I can also argue for (7) by appealing to the following principle (cf. Ginet 1990: 102-3): (10) necessarily, if a proposition Q is settled at time t and at t you can actualize a proposition R , then at t you can actualize the conjunction of Q with R . By (10), since the proposition Q that you pay the *January* installment on time is settled on February 4 and on February 4 (I assume) you can actualize the proposition R that you pay the *remaining nine* installments on time, on February 4 you can actualize the conjunction of Q with R , namely the proposition that you pay *all ten* installments on time (i.e., (7) is true). One might object that (10) has the counterintuitive consequence that today you can actualize the proposition that *the sun rose yesterday and you help me*. I reply that this consequence is not counterintuitive because (as I understand it) the claim that you can actualize that conjunction does not entail that you can be the *only* cause of the fact that the conjunction is true; the claim is instead true because you can (by helping me) *bring to completion a causal process* resulting in the conjunction becoming settled. Compare: if you are the last runner of your team in a $4 \times 100\text{m}$ relay race and you can run fast enough, then you can make it the case that your team wins (although you cannot be the *only* cause of the fact that your team wins) because you can bring to completion a causal process resulting in your team's victory.⁴¹

2.8. OICS and epistemic obligations

Here is the final objection to OICS that I will address: agents sometimes have *epistemic* obligations to form certain beliefs, but normally they cannot satisfy these obligations because normally the formation of one's beliefs is not under one's direct voluntary control (e.g., one cannot simply decide to form a belief in order to collect a reward⁴²). In reply, note first that, if (1) you have an

⁴⁰ In response, one might grant that (9) is false, but might argue that (7) entails (8) by appealing instead to the following principle: (9') necessarily, if at time t you can actualize the conjunction of two *logically independent* propositions (i.e., propositions Q and R such that Q entails neither R nor the negation of R , and neither of these entails Q), then at t you can actualize both propositions. This principle avoids my counterexample to (9) because the propositions that you help me and that either you help me or the sun rose yesterday are not logically independent. By (9'), (7) entails (8) because the proposition that you pay *all ten* installments on time is the conjunction of the logically independent propositions that you pay the *January* installment on time and that you pay the *remaining nine* installments on time. I reply that (9') is false. To see this, suppose that today you can help me. Consider the logically independent propositions that (i) *either you help me or the sun rose yesterday* and that (ii) *either you help me or the sun neither rose nor set yesterday*; their conjunction is logically equivalent to the proposition that you help me, so today you can actualize this conjunction. But today you cannot actualize (i), since (i) is already settled. (Similar remarks apply if one grants that (6) is false but argues that (4) entails (5) by appealing instead to the following principle: (6') necessarily, if at t you have an unconditional obligation whose satisfaction proposition is the conjunction of two logically independent propositions, then at t you have two unconditional obligations whose satisfaction propositions are those two propositions.) My counterexample to (9') works because (i) is already settled, so to avoid the counterexample one might propose instead the following principle: (9'') necessarily, if at t you can actualize the conjunction of two logically independent propositions Q and R , and Q is not settled at t , then at t you can actualize Q . In reply, I agree that (9'') is true, but (9'') cannot be used to argue that (7) entails (8): the proposition Q that you pay the *January* installment on time is settled on February 4.

⁴¹ The above considerations suggest that my use of "can actualize" (that vindicates (10)) is reasonable, but I am not saying it is the *only* reasonable use (see the end of note 3). I am not committed to claims analogous to (10) concerning *making true* (cf. Rodriguez-Pereyra 2006, 2009) or *bringing about the truth of* (cf. Guigon 2009: 454) a conjunction: I can grant that you cannot make true (although you can actualize), and you cannot bring about the truth of, the proposition that the sun rose yesterday and you help me.

⁴² See: Alston 1988: 263, 2005: 62-3; Bennett 1990: 89-90; Bondy 2015: 758-9, 764-7; Chrisman 2008: 346; Chuard & Southwood 2009: 605-6; Feldman 2008: 341; Heller 2000: 131; McHugh 2012: 68-9; Plantinga 1993: 24; Reynolds 2011: 23; Ryan 2003: 63; Weatherson 2008: 544, 547.

epistemic obligation to start believing a proposition P ,⁴³ then (2) you are in a situation in which you have *evidence* for P .⁴⁴ But then OICS and (1) jointly entail that (3) *in that situation* you can start believing P ; they do not entail that (4) you can start believing P even in a situation in which you have no evidence for P , and thus they do not entail that (5) you can start believing P *at will* (so whether or not you have evidence for P). So OICS (in conjunction with (1)) does not conflict with the claim that (6) the formation of your beliefs is not under your direct voluntary control.⁴⁵ In response, one might grant that OICS and (1) jointly entail (3) instead of (5), but might claim that this is still a problem for OICS because (3) is false: in some situations, you have evidence for P but you cannot start believing P . I examine next an (alleged) example of such a situation.

Suppose that you are presented with overwhelming evidence for the proposition G that your spouse is guilty of murder (Mizrahi 2012), but you are “psychologically unable” to start believing G : believing G “would just be too painful and damaging for you” (Ryan 2003: 59; cf. Feldman 2008: 350). Then, one might argue, you have an epistemic obligation to start believing G , but you cannot start believing G . I reply that I find this objection to OICS no more convincing than the following common objection to OIC: a kleptomaniac has a moral obligation to refrain from stealing, but is psychologically unable to refrain from stealing and thus cannot refrain from stealing. Addressing this objection to OIC, in previous work (Vranas 2007: 182-6) I noted that I

⁴³ Strictly speaking, I am talking about an epistemic obligation to start believing P (within a reasonable amount of time) *in response to appropriate evidence for P* (and to *keep* believing P until the evidence charges): if you start believing P not because you are convinced by your evidence for P but rather because you take a pill that makes you start believing P , then your epistemic obligation is not satisfied. So the satisfaction proposition of the epistemic obligation that I am talking about is *not* the proposition that *you believe P over a time interval* (that starts within a reasonable amount of time and ends when your evidence appropriately charges). It is also definitely not the proposition that *you believe P at a particular time* (or *at some time or other*): I doubt that there is any epistemic obligation which is satisfied if an agent has a belief at just a *single* time instant (cf. note 49).

⁴⁴ One might object as follows: if you have negligently failed to acquire strong evidence for P that you could have easily acquired, then you may have an epistemic obligation to start believing P even if you have no evidence for P (cf. Kornblith 1983: 35-6). In reply, I ask: do you *know* that there is strong evidence for P that you have failed to acquire? If you do, then you do have evidence for P ; but if you do not (and you have no evidence for P), then you do not have an (objective) epistemic obligation to start believing P . (Cf.: Feldman 2000: 687-8; Feldman & Conee 1985: 21-2. On the dependence of objective epistemic—unlike objective moral (though see note 2)—obligations on the agent’s “perspective”, see: Feldman 1988b: 408-11; Russell 2001: 40-1.)

⁴⁵ Here is the reasoning more explicitly: since the objection to OICS relies on (6), I charitably understand (6) weakly, namely so that its negation is strong enough to entail (5); so if one supposes for *reductio* that OICS & (1) conflicts with—i.e., entails the negation of—(6), then OICS & (1) entails (5), contradicting the claim that—because (5) entails (4) but OICS & (1) does not entail (4)—OICS & (1) does not entail (5). To my knowledge, this is a novel (although to some extent anticipated in Reynolds 2011: 26-7) response to the alleged conflict between OICS, (1), and (6)—or to the alleged conflict between (a) epistemic ought-implies-can principles, (b) *epistemic deontology* (defined in Chrisman 2008: 347 as the claim that beliefs are proper subjects of epistemic oughts), and (c) *doxastic involuntarism* (defined in Chrisman 2008: 346 as the claim that the formation of one’s beliefs is not within one’s direct voluntary control). Like me, Chuard and Southwood (2009: 614-9), as well as McHugh (2012: 86-8), respond to the alleged conflict by arguing in effect that OICS and (1) do not jointly entail the negation of (6); however, these authors do not provide my explanation of why the entailment fails. (Ryan (2003: 63-6) makes a distinction similar to my distinction between (3) and (4), but does not use her distinction to argue that the above entailment fails.) Other responses to the alleged conflict include: (i) rejecting epistemic ought-implies-can principles (Feldman 1988a: 240-3, 2000: 676, 2001: 87-8; Kornblith 2001: 237-8; Ryan 2003: 59-60; cf. Qu 2017); (ii) rejecting epistemic deontology (Alston 1988, 2005: chap. 4); (iii) rejecting doxastic involuntarism (Bondy 2015; Heller 2000; Russell 2001: 42-3; Ryan 2003: 62-75; Shaffer 2013; Shah 2002; Steup 2000; Weatherson 2008; cf.: Chisholm 1968: 224-5; Ginet 2001; Goldstick 2010; Steup 2008, 2012); and (iv) arguing that epistemic ought-claims are not obligation claims (Chrisman 2008: 358-66; cf. Plantinga 1993: 24). See also: Côté-Bouchard 2017; Engel 2009; Peels 2017.

understand the claim that you can do something as compatible with the claim that you are psychologically “unable” to do it: even if you are psychologically unable to stick your arm into a cesspool to retrieve your wallet (Stocker 1971: 311), strictly speaking you *can* stick your arm in (cf. Bykvist & Hattiangadi 2007: 279). Similarly, even if it is extremely hard for kleptomaniacs to resist their urges to steal, strictly speaking they *can* resist them, as evidenced by the fact that they “will generally avoid stealing when immediate arrest is probable (e.g., in full view of a police officer)” (American Psychiatric Association 2013: 478; cf. Pickard 2015). But if kleptomaniacs (who have a recognized mental disorder) are able to resist their urges to steal, why aren’t you (who, I assume, have no mental disorder) able to accept the implications of the evidence for *G*? Mizrahi (2012: 835, 839) argues in effect that you cannot start believing *G* because *you have reasons to doubt G*: you have evidence that your spouse has a good moral character. I reply that having reasons to doubt a proposition need not render one unable to start believing the proposition: some people who have reasons to doubt creationism (because they know that the vast majority of biologists believe in evolution) nevertheless become creationists.

One might argue that, since I understand the claim that you can start believing a proposition *P* as the claim that you can *make it the case* that you start believing *P*, there is a general reason why, even if you have evidence for *P*, you *cannot* start believing *P*: if you start believing *P* in response to the evidence for *P*, what makes it the case that you start believing *P* is the *evidence*, not *you*. As Alston puts it, “the belief follows automatically, without intervention from the will, from the way things seem at the moment to the subject” (1988: 266; cf. 2005: 65). This objection relies on something like what Williams calls “the picture offered by Hume of belief as a passive phenomenon, something that happens to us” (1970/1973: 148; cf. Pojman 1993: 528-32). I reply that this picture is inaccurate: in many cases, you do make it the case that you start believing *P*. To see this, suppose (to adapt an example from Shah & Velleman 2005: 512-3) that you add a long column of numbers, and you arrive at the total 987. You suspect that the sum is 987, but you suspend belief until you check the addition. You add the numbers again, and you arrive at the same total. Although you could decide to go through one more iteration, you decide instead to stop checking, and as a result you start believing that the sum is 987. In this example, by means of your decision to stop checking, you make it the case that you start believing that the sum is 987. This is not to say that you are the *only* cause of the fact that you start believing that the sum is 987: your evidence also played a significant causal role.⁴⁶ Still, the ultimate causal responsibility was yours: you could have decided instead to keep suspending belief. The example generalizes: in many cases in which you have evidence for a proposition *P*, you decide to suspend belief in *P* and start deliberating, but after some time you decide to stop deliberating, and as a result you start believing *P*. And in many cases in which you *do not* go through this process, you *can* go through it, so you *can* start believing *P*. It may be worth repeating that I am not saying you can start believing *P at will*, since I am not saying you can (decide to stop deliberating and) start believing *P* even if you have no (convincing) evidence for *P*.

One might respond by considering cases in which there is no room for deliberation. Suppose that you glance at my shirt and it looks red to you, and as a result you start believing the proposition *R* that my shirt is red. One might argue that your belief is “evidentially compelled”: there is no

⁴⁶ Recall from note 7 that I take causes and effects to be facts, not agents or propositions (or pieces of evidence). Strictly speaking, then, I claim that in the above example *the fact that you decide to stop checking* causes (but is not the *only* cause of) the fact that you start believing that the sum is 987.

way you can prevent its formation⁴⁷ and start deliberating instead (just as there is no way you can prevent your pupils from contracting when they are suddenly exposed to bright light), so my argument in the previous paragraph does not apply. Moreover (one might continue), if you cannot make it the case that you do *not* start believing *R*, then you cannot make it the case that you *start* believing *R* either: the claim that you can (make it the case that you) start believing *R* presupposes that you have a *choice* on whether you start believing *R*.⁴⁸ I reply that, regardless of whether you can start believing *R*, this objection to OICS is unconvincing because it is not clear that there is any time at which you have an epistemic obligation to start believing *R*. Here is why. (1) *Before* you glance at my shirt, you have (I assume) no evidence for *R*, so you have no epistemic obligation to start believing *R*. (2) *After* (or *when*) you start believing *R*, the proposition that you start believing *R* is already settled, so you have no obligation (epistemic or not) to start believing *R*.⁴⁹ So, (3) if at a given time you have an epistemic obligation to start believing *R*, then at that time you have already glanced at my shirt but you have not yet started believing *R*. But (4) it is not clear that there is any such time: arguably, you start believing *R as soon as* you glance at my shirt. (Compare: if a first billiard ball strikes a second and as a result the second ball starts to move, is there any time at which the first ball has already struck the second but the second ball has not yet started to move?⁵⁰) In response, one might try to come up with a modified example in which it is clear that you start believing *R* shortly *after* you glance at my shirt; e.g.,

⁴⁷ See: Alston 1985: 64-5, 1988: 264, 270, 2005: 63-4, 68; Feldman 1988a: 238, 242; Peels 2017: 55; Plantinga 1993: 24; Pojman 1993: 529; Reynolds 2011: 20; Ryan 2003: 62-3, 73; Scott-Kakures 1994: 79; Steup 2000: 28, 45-6, 2012: 149; Weatherson 2008: 556-61.

⁴⁸ Actually, there are counterexamples to the claim that, if you cannot make it the case that the negation of a proposition is true, then you cannot make it the case that the proposition is true either. A first counterexample is provided by my objection to OICV (§2.5): you cannot make it the case that your daughter participates in the study, but you can make it the case that she does not participate. For a second counterexample, consider a “Frankfurt-type” example (cf. §2.5). Suppose that, without any outside interference, you decide to stand up, and as a result you stand up. Unbeknownst to you, however, a “counterfactual intervener” was monitoring your brain and would have made you decide to stand up (and then stand up) if your brain waves had not shown that you were going to decide on your own to stand up. In this example, you cannot make it the case that you do not stand up, but you can (and you do) make it the case that you stand up. Cf. Chuard & Southwood 2009: 604-5, 627 n. 17.

⁴⁹ Here I am appealing to the principle that, if an agent at a given time has an unconditional obligation, then neither the satisfaction nor the violation proposition of the obligation is settled at that time. This principle follows from OICS, but appealing to the principle in the course of addressing an objection to OICS is not circular because the principle is independently plausible: you do not now have an obligation to call me yesterday, since either it is now settled that you called me yesterday or it is now settled that you did not call me yesterday. This principle entails that an agent at a given time has no epistemic obligation to *believe P at that time* (as opposed to an epistemic obligation to *start believing P* at a *later* time; on this distinction, see: Alston 1988: 261-2, 2005: 61; McHugh 2012: 78-82; Weatherson 2008: 544-5). For example, you do not now have an epistemic obligation to believe that you are reading this paper right now (contrast Chrisman 2008: 347)—although you may be now *justified* in believing that you are reading this paper right now. On the other hand, the above principle (like OICS) is compatible with the claim that an agent at a given time has an epistemic obligation to believe *P* at a *later* time (or to *keep believing P* in the near future).

⁵⁰ One might argue that there is exactly one such time, namely the time at which the first ball strikes the second, because at any later time the second ball is moving. Similarly, one might argue that there is exactly one time at which you have an epistemic obligation to start believing *R*, namely the time at which you glance at my shirt. I reply that there is no unique such time. Contrary to the way that (for simplicity) I talk in the text, glancing at my shirt (and perceiving it as red) occurs during a short time interval. At any time in the interior of that interval, you have not yet fully processed the perceptual information, so you have no evidence for *R* and thus no epistemic obligation to start believing *R*. The objection to OICS is tenuous if it relies on the assumptions that the interval has a last time and that at that time you do not start believing *R*.

suppose that you hesitate for a few seconds to believe that my shirt is red because you entertain the possibility that there is a source of red light in the room. I reply that, in this modified example, it is not clear that your belief is evidentially *compelled*: arguably, during those few seconds of hesitation, you *can* (deliberate and) prevent its formation.⁵¹ But then the objection to OICS does not get off the ground.

3. Conditional obligations

3.1. Why a new principle is needed

Suppose you promise me that, if you drink, you will let me drive, and as a result you acquire a conditional obligation to let me drive if you drink. This obligation is violated exactly if you drink but you do not let me drive. Under what conditions is the obligation satisfied? One might propose that it is satisfied exactly if it is not violated; equivalently, exactly if either you do not drink or you let me drive. I see two problems with this proposal (cf. Vranas 2008: 534-5, 555-7). First, on this proposal, the obligation is satisfied if you do not drink and yet you let me drive, and I find this counterintuitive: intuitively, the obligation is “inoperative” (rather than satisfied) if you do not drink. Second (and more important), on this proposal, your conditional obligation to let me drive if you drink is satisfied (or violated) exactly if so is an unconditional obligation to either not drink or let me drive; similarly, on a generalization of this proposal, every conditional obligation is indistinguishable in terms of satisfaction and violation from an unconditional obligation, so the distinction between unconditional and conditional obligations collapses. (I take it for granted that there is such a distinction; defending this lies beyond the scope of this paper.) To avoid these problems, I propose instead that your conditional obligation to let me drive if you drink is satisfied exactly if you *both* drink and let me drive. On my proposal, the satisfaction proposition of your conditional obligation is not the negation of the violation proposition: your obligation is *neither satisfied nor violated*—say, equivalently, that it is *avoided*—exactly if you do not drink.⁵² More generally, an obligation is *unconditional* exactly if its satisfaction proposi-

⁵¹ In response, one might try to come up with an example in which it is clear that at time *t* you acquire evidence that *compels* you to (sooner or later) start believing a proposition *P*, but it is also clear that you only start believing *P* shortly *after t*. I reply that, if at *t* your belief is evidentially *compelled*, then at *t* it is *inevitable* (i.e., ability-necessary) for you that you (sooner or later) start believing *P*. But then—as OIPV entails—at *t* you have no obligation to start believing *P*, so no such example is a counterexample to OICS (regardless of whether at *t* you *can* start believing *P*). One might object that OIPV fails for epistemic obligations (Feldman 1988a: 242). I reply that considerations similar to those I adduced in §2.5 support OIPV for epistemic obligations. Suppose that, before you glance at my shirt, you have an epistemic obligation to start believing that it is red (because I have told you that it is red, and you know that I never lie), but you are deliberating on whether to start believing that it is red (because you know that I have never worn a red shirt before). When you glance at my shirt, even if you do not immediately start believing that it is red, it is plausible that you lose the obligation: you are done with it, since the compelling perceptual evidence settles the matter for you. OIPV explains why you lose the obligation.

⁵² One might object that, if you do not drink *in order* to keep your promise (for example, because you do not want to let me drive), then you keep your promise and thus your promissory obligation (to let me drive if you drink) is *satisfied* rather than avoided. In reply, note first that no similar move is available if you promise to let me drive *if it snows*: there is no plausible case in which it does not snow because you do not want to let me drive. Moreover, there are two kinds of cases in which you keep (i.e., you do not break) your promise to let me drive if you drink: cases in which you do not drink, and cases in which you both drink and let me drive. If one wants to say (as the objection in effect proposes) that in both kinds of cases your promissory obligation is satisfied, then one can distinguish two kinds of satisfaction, satisfaction₁ and satisfaction₂, corresponding to the two kinds of cases. If one does this, then the distinction between unconditional and conditional obligations does not collapse: although every conditional obligation is indistinguishable *in terms of satisfaction and violation* from an unconditional obligation, the two obliga-

tion is (equivalent to) the negation of its violation proposition, and an obligation is *conditional* exactly if it is not unconditional.

Given this understanding of conditional obligations, I can now explain why in §2.3 I formulated OICS only in terms of unconditional obligations. The reason is that the following analog of OICS for conditional obligations fails: (OICS*) if an agent at a given time has a conditional obligation, then the agent at that time can satisfy the obligation. To see that OICS* fails, suppose that you promise your mother to take shelter in the basement if a tornado hits the town (within the next few hours). Then you have a conditional obligation to take shelter in the basement if a tornado hits the town. But you cannot satisfy this obligation: you cannot actualize its satisfaction proposition, namely the proposition that both a tornado hits the town and you take shelter in the basement, since you have no influence over whether a tornado hits the town.⁵³

Although the example in the previous paragraph falsifies OICS*, it does not falsify OIPV (which, to repeat, is the principle that, if an agent at a given time has an obligation, then the violation proposition of the obligation is ability-possible for the agent at that time). Indeed, the violation proposition of your conditional obligation in the above example—namely the proposition that a tornado hits the town but you do not take shelter in the basement—is ability-possible for you: it might be true if you were to do something that you can do, namely refrain from taking shelter in the basement. I believe that no example of a conditional (or unconditional) obligation falsifies OIPV; this is why in §2.5 I formulated OIPV (in contrast to OICS) in terms of *all* obligations that agents have, including conditional obligations. Nevertheless, OIPV is not the whole story concerning conditional obligations: in some cases, OIPV does not explain why an agent loses or fails to have a conditional obligation. To see this, consider two examples. (1) Suppose that in the morning you have a conditional obligation (because you have promised) to run in tomorrow’s marathon if it does not snow tonight. If you have an accident at noon that leaves you paralyzed and thus unable to run in tomorrow’s marathon, then you lose this conditional obligation at noon. But OIPV does not explain why you lose it: right after the accident, the violation proposition of the obligation—namely the proposition that it does not snow tonight but you do not run in tomorrow’s marathon—is still ability-possible for you.⁵⁴ (2) If you (now) have no influence over the outcome of the lottery, then you have no conditional obligation to win the lottery if I do not win. But OIPV does not explain why not: the violation proposition of such an obligation—namely the proposition that neither of us wins the lottery—is ability-possible for you.

tions are distinct because satisfaction₁ is possible for the conditional but not for the unconditional obligation. But then, I note, my proposal—which is to talk of avoidance in place of satisfaction₁ and of satisfaction in place of satisfaction₂—is a terminological variant of what the objection in effect proposes.

⁵³ This inference relies on the following consequence of principle (9'') in note 40: if you can actualize the conjunction of the logically independent propositions that (*Q*) a tornado hits the town and that (*R*) you take shelter in the basement, and *Q* is not yet settled, then you can actualize *Q*. Note that, if *Q* is already settled (i.e., a tornado has already hit the town) and you can actualize *R* (i.e., you can take shelter in the basement), then you *can* actualize the conjunction of *Q* with *R* (see principle (10) in the last paragraph of §2.7). (I take the satisfaction proposition of your conditional obligation—to take shelter in the basement if a tornado hits the town—to be the conjunction of *Q* with *R*: for the purpose of keeping your promise, it does not matter whether you take shelter in the basement *before* or *after* a tornado hits the town.)

⁵⁴ One might argue that you do not lose the obligation at noon: the violation proposition of the obligation is still impermissible (because your promise is still in force) after the accident, so you still have the obligation after the accident. I reply that, even if the violation proposition of the obligation is *impersonally* impermissible after the accident, it is not impermissible *for you* after the accident; so, even if the obligation is still *in force* after the accident, you no longer *have* it after the accident (cf. Vranas 2017b).

(The observation that you *cannot satisfy* such an obligation provides no explanation either: as I argued, OICS* is false.) To explain why in these examples you lose or you fail to have a conditional obligation, a new principle is needed. I propose next such a new principle, which also generalizes (and thus supersedes) OICS.

3.2. The “ought-implies-can-obey” principle (OICO)

Before I formulate the new principle, I introduce some concepts. Say that an obligation is *obeyed* exactly if it is not violated (i.e., it is either satisfied or avoided). Let the *obedience proposition* of an obligation be the negation of its violation proposition (Vranas 2011: 374). For example, the obedience proposition of your conditional obligation to confess if you repent is the negation of the proposition that you both repent and do not confess; equivalently, it is the proposition that either you do not repent or you confess (i.e., the proposition that, if you repent, then you confess). By contrast, the obedience proposition of your unconditional obligation to confess is the negation of the proposition that you do not confess (i.e., it is the proposition that you confess), and thus is the same as the satisfaction proposition of your unconditional obligation. More generally, for *unconditional* obligations, avoidance (i.e., neither-satisfaction-nor-violation) is impossible, so obedience (i.e., satisfaction-or-avoidance) amounts to satisfaction. Given these concepts, here is my proposed new principle, namely the *ought-implies-can-obey* principle:

(OICO) If an agent at a given time has an obligation, then the agent at that time *can obey* the obligation (i.e., can actualize its obedience proposition).

OICO entails OICS: if OICO is true, then any agent who (at a given time) has an *unconditional* obligation can *obey* the obligation, and then OICS is also true because the agent can *satisfy* the obligation (since, for unconditional obligations, obedience amounts to satisfaction). Since OICO entails OICS, OICO yields correct results concerning the loss of—and the failure to have—unconditional obligations (if my arguments in §2 succeed). Concerning conditional obligations, note first that my counterexample to OICS* (§3.1) leaves OICO unscathed: if you have a conditional obligation to take shelter in the basement if a tornado hits the town, then you *can* actualize the obedience proposition of this obligation (namely the proposition that either no tornado hits the town or you take shelter in the basement), since you can take shelter in the basement. Consider also again the two examples I examined in the last paragraph of §3.1. (1) Suppose that in the morning you have a conditional obligation to run in tomorrow’s marathon if it does not snow tonight, but you have an accident at noon that leaves you paralyzed. In contrast to OIPV, OICO does explain why you lose this conditional obligation at noon: because at noon you become unable to actualize its obedience proposition, namely the proposition that either it snows tonight or you run in tomorrow’s marathon. (2) Suppose that you (now) have no influence over the outcome of the lottery. In contrast to OIPV, OICO does explain why you have no conditional obligation to win the lottery if I do not win: because you cannot actualize the obedience proposition of such an obligation, namely the proposition that either I win the lottery or you win.

To make some further remarks on OICO, I introduce a definition. Let an *unconditional counterpart* of a conditional obligation be an unconditional obligation with the same violation proposition (and thus the same obedience proposition) as the conditional obligation. For example, an unconditional obligation for you to either drive or fly is an unconditional counterpart of a conditional obligation for you to fly if you do not drive. An unconditional counterpart of a conditional obligation is *satisfied* exactly if the conditional obligation is *obeyed*. Therefore, according to OICO, if an agent at a given time has a conditional obligation, then the agent at that time can sat-

isfy any unconditional counterpart of the conditional obligation. One might then ask: is it plausible to propose (as OICO does) the same condition as necessary both for having a conditional obligation and for having an unconditional counterpart of the conditional obligation? I reply that it *is* plausible: if you cannot make it the case that either you drive or you fly, then you have neither a conditional obligation to fly if you do not drive nor an unconditional obligation to either drive or fly. To propose the same condition as necessary for having either obligation is not to confuse the two obligations (cf. note 16).

Here is an objection to OICO. Suppose you have contracted a virus that has made you unable to walk for the foreseeable future, and you have no influence over whether or when you will be cured. Nevertheless, if on Monday you promise me that you will walk on Friday if you are cured by Wednesday, then on Monday you acquire a corresponding conditional obligation. One might argue, however, that according to OICO you do not acquire such an obligation because on Monday you cannot actualize the obedience proposition of such an obligation, namely the proposition that either (1) you are not cured by Wednesday or (2) you will walk on Friday: on Monday you cannot actualize (1) because you have no influence over whether or when you will be cured, and on Monday you cannot actualize (2) because the virus has made you unable to walk for the foreseeable future. I reply first that this reasoning is invalid: in some cases, one can actualize a disjunction although one cannot actualize any of the disjuncts. For example, although you can *neither* actualize the proposition that you will buy an umbrella and it will rain *nor* actualize the proposition that you will buy an umbrella and it will not rain (since you have no influence over whether it will rain), you *can* actualize the disjunction of these two propositions, which is logically equivalent to the proposition that you will buy an umbrella.⁵⁵ Similarly, I argue next that on Monday you *can* actualize the disjunction of (1) with (2) (equivalently, the conditional “if you are cured by Wednesday, you will walk on Friday”).

My argument appeals to the principle that you are currently able to actualize a proposition P if there is a future time t such that, if P were not settled by t , then by t you would be able to actualize P .⁵⁶ In support of this principle, suppose that you are currently competing against me in a five-kilometer race, and I have taken a drug such that, if it did not kill me within a minute (if it did, you would win the race by default), it would significantly weaken me within a minute (and in that case you would become able to run faster than me and thus to win the race). In this example, there is a future time t (namely one minute from now) such that, if the proposition P that you win the race were not settled by t , then the drug would not kill me but would significantly weaken me by t , and then by t you would be able (to run faster than me and thus) to actualize P ; so the above principle yields the intuitively correct consequence that you are currently able to (actual-

⁵⁵ For another example, although you can neither cause the coin to come up heads nor cause the coin to come up tails, you *can* cause the coin to come up heads or tails: you can toss the coin.

⁵⁶ The principle in the text does not entail the claim that, if there is a future time at which you will be able to actualize P , then you are currently able to actualize P . That claim is false: even if tomorrow you will win the lottery and thus you will be able to donate ten million dollars to charity tomorrow, today you are unable to donate ten million dollars to charity tomorrow (§2.4). By contrast, if tomorrow (is your payday, so) you will receive your paycheck and thus you will be able to pay your bills tomorrow, today you are already able to pay your bills tomorrow (even if you are unable to pay them today). To explain the contrast between the two cases, note that today there is no reasonable expectation (even if it is true) that you will win the lottery tomorrow, but today there is a reasonable expectation that you will receive your paycheck tomorrow. Similarly, the principle in the text is informative about cases in which by some future time you *would* be able—and thus there is a reasonable expectation that you *will* be able—to actualize P if P is not settled by that time.

ize the proposition P that you) win the race. Go back now to the disjunction of (1) with (2) in the previous paragraph. There is a future time t (namely the end of the day Wednesday) such that, if the disjunction were not settled by t , then you would be cured by t , and then by t you would be able (to walk on Friday and thus) to actualize the disjunction; so the above principle yields the consequence that on Monday you are already able to actualize the disjunction. One might respond that this consequence cannot be correct because, if you are *not* cured by Wednesday, then you do not actualize the disjunction (since it becomes settled without your involvement), and subsequently you are unable to actualize it (since it is already settled). I reply that the claim that you are currently able to actualize a proposition is compatible both with the claim that you will not in fact actualize it and with the claim that in the future you will no longer be able to actualize it. I conclude that the objection to OICO in the previous paragraph fails.

3.3. Why OIPV is still needed

Despite its successes, OICO does not *always* explain why agents lose conditional obligations. To see this, consider a modification of the example I used in §2.5 to argue that OICS does not always explain why agents lose unconditional obligations. Suppose I ask you to euthanize my terminally ill dog. At 9am, you promise me that the dog will be dead by midnight *if* it is alive at 1pm, and as a result you acquire a corresponding conditional obligation. At noon, you administer to the dog a quick-acting and invariably fatal poison to which there is no antidote, with the result that the dog will inevitably be dead within a few minutes (and thus by midnight). Then you lose the conditional obligation at noon: you are done with it. But OICO does not explain why you lose it, since you do not become unable to obey it: after you administer the poison but before the poison acts, you can kill the dog by shooting it, so you can still actualize the obedience proposition of the obligation, namely the proposition that the dog is (not alive at 1pm or) dead by midnight. By contrast, when you administer the poison, the obedience proposition of the obligation becomes ability-necessary for you; equivalently, the violation proposition of the obligation becomes ability-impossible for you, so OIPV does explain why you lose the obligation.⁵⁷

Moreover, in some cases OIPV explains but OICO does not explain why agents fail to have a conditional obligation. For example, just as OIPV explains why you (now) have no unconditional obligation to *not* run faster than light tomorrow (since running faster than light tomorrow is ability-impossible for you), OIPV explains why you have no conditional obligation to *not* run faster than light tomorrow *if* you fail to exercise today: the violation proposition of such an obligation, namely the proposition that both you fail to exercise today and you run faster than light tomorrow, is ability-impossible for you. But OICO does not explain why you have no such conditional obligation: you *can* actualize the obedience proposition of such an obligation, namely the proposition that either you exercise today or you do not run faster than light tomorrow, since

⁵⁷ I can now explain why in §2.5 I formulated OIPV in terms of violation instead of satisfaction. Consider the “ought-implies-not-necessary-satisfaction” principle: (OINNS) if an agent at a given time has an obligation, then the satisfaction proposition of the obligation is not ability-necessary for the agent at that time. OIPV entails OINNS, and the restrictions of the two principles to unconditional obligations are equivalent. Nevertheless, OINNS does not explain why you lose the obligation in the above example. This is because, when you administer the poison, the satisfaction proposition of your conditional obligation—namely the proposition that the dog is alive at 1pm and dead by midnight—does not become ability-necessary for you: it might be false if you were to do something that you can do, namely shoot the dog before the poison acts.

you can exercise today.⁵⁸ I propose then (taking also into account some results from §3.2) that both OICO and OIPV are needed to explain why (or to correctly predict when) agents lose or fail to have conditional obligations.⁵⁹

One might be worried by the *asymmetry* between obedience and violation exhibited by the OICO/OIPV pair: OICO combines obedience with *ability*, but OIPV combines violation with *ability-possibility*. In reply, I ask: what would count as a *symmetric* pair of principles? In a first sense, two principles are symmetric exactly if they propose necessary conditions (for having an obligation) that ascribe the *same property* (e.g., being ability-possible) to *different propositions* (e.g., to the satisfaction and the violation proposition of an obligation). OIPS (§2.6) and OIPV are symmetric in this sense, but clearly OICO and OIPV are not. In a second sense, two principles are symmetric exactly if they propose necessary conditions that ascribe *different properties* (e.g., being ability-possible and not being ability-necessary) to the *same proposition* (e.g., to the obedience proposition of an obligation). It turns out that OICO and OIPV *are* symmetric in this sense, as one can see by noting that OIPV can be equivalently reformulated in terms of obedience as follows:

(OIPV) If an agent at a given time has an obligation, then the obedience proposition of the obligation is *not inevitable* (i.e., is not ability-necessary) for the agent at that time.

⁵⁸ One can similarly see that OIPV explains but OICO does not explain why you (now) have no conditional obligation to exercise today if you run faster than light tomorrow. More generally, it is a consequence of OIPV (but not of OICO) that, if a proposition *P* (like the proposition that you run faster than light tomorrow) is *ability-impossible* for an agent at time *t*, then, for any proposition *Q*, the agent at *t* has no conditional obligation to the effect that *Q* is true if *P* is true (or to the effect that *P* is false if *Q* is false): the violation proposition of such an obligation, namely the conjunction of *P* with the negation of *Q*, is also ability-impossible for the agent at *t*. On the other hand, it is a consequence of OICO (but not of OIPV) that, if (1) a proposition *P* (like the proposition that you do *not* run faster than light tomorrow) is *ability-necessary* for an agent at time *t*, then, for any proposition *Q*, if (2) the agent at *t* has a conditional obligation to the effect that *Q* is true if *P* is true (or to the effect that *P* is false if *Q* is false), then (3) the agent at *t* can actualize *Q*. *Proof.* By OICO, (2) entails that (4) the agent at *t* can actualize the disjunction *Q'* of *Q* with the negation of *P*. By (1), however, (5) *Q* and *Q'* are *ability-equivalent* for the agent at *t*, in the sense that it is ability-necessary for the agent at *t* that *Q* and *Q'* are materially equivalent (i.e., that *Q* is true exactly if *Q'* is true). Finally, (4) and (5) jointly entail (3) by the *ability-equivalence-implies-can-equivalence* principle: (AEICE) if propositions *Q* and *Q'* are ability-equivalent for an agent at time *t*, then the agent at *t* can actualize *Q* exactly if the agent at *t* can actualize *Q'*. This principle is of more general interest: it has as a corollary principle (10) in the last paragraph of §2.7.

⁵⁹ One might object to OIPV as follows. Suppose you promise at noon that, if you ever become able to levitate, you will do so (as soon as you become able). According to OIPV, you do not acquire a corresponding conditional obligation at noon: the violation proposition of such an obligation, namely the proposition that at some time you become able to levitate but you do not do so, is not ability-possible for you right after you promise, assuming that right after you promise it is ability-necessary for you that you will never become able to levitate. Now recall that ability-necessity does not entail truth (see (2) in note 29), so suppose further that (1) at 3pm you suddenly become able to levitate. Then (2) at 3pm you become obligated to levitate, and one might argue that this is because, contrary to OIPV, (3) until 3pm you do have a conditional obligation to levitate if you ever become able to do so. In reply, note first that the argument from (1) and (3) to (2) is invalid: the conjunction of its premises is compatible with the possibility that at 3pm you lose the conditional obligation (for example, because you are released from your promise) and so at 3pm you do not become obligated to levitate. This suggests that to explain (2) one needs to appeal instead to the premise that (3') right *after* 3pm you have a conditional obligation to levitate if you ever become able to do so. But (3') does not conflict with OIPV: (3') is compatible with the possibility that (because your promise is still in effect: see the end of §2.4) you acquire such a conditional obligation *at 3pm*.

Since OICO and OIPV are symmetric in the second sense, I see no reason to be worried by the fact that they are not symmetric in the first sense: it is impossible for two principles to be symmetric in *both* senses.⁶⁰

3.4. Are OICO and OIPV the whole story?

In some cases in which an obligation is *revoked*, neither OICO nor OIPV explains why an agent loses the obligation. For example, suppose that in the morning you promise me (and you never promise anyone else) to call me in the evening, and as a result you acquire a corresponding obligation, but at noon I release you from your promise. Then at noon you lose the obligation to call me in the evening (cf. Ross 1939: 110), but neither OICO nor OIPV explains why you lose it: you can still actualize the obedience proposition of the obligation, namely the proposition that you call me in the evening, and that proposition is still not inevitable for you. But *apart* from cases of revocation, are there any cases in which neither OICO nor OIPV explains why an agent loses or fails to have an obligation? To conclude my discussion of conditional obligations, I will examine three alleged—but, as I will argue, ultimately unsuccessful—examples of such cases.

First example. One might argue that neither OICO nor OIPV explains why you (now) have no conditional obligation to run faster than light tomorrow *if* you fail to exercise today: you *can* actualize the obedience proposition of such an obligation, namely the proposition that either you exercise today or you run faster than light tomorrow (since you can exercise today), and that proposition is *not* inevitable for you (since it would be false if you were to do something that you can do, namely refrain from exercising today). I reply that it is possible for you to have such a conditional obligation. To see this, suppose that you are a soldier and your commanding officer gives you the following order at noon: “if you fail to exercise today, run faster than light tomorrow”. It will not do for you to protest that you cannot run faster than light tomorrow, since the officer could respond: “I did not order you to run faster than light tomorrow, something that indeed you cannot do. I ordered you instead to run faster than light tomorrow *if* you fail to exercise today. You *can* obey this order: you can exercise today.” And if you ask the officer why she did not just order you instead to exercise today, the officer could respond: “I wanted you to use your brain to figure out what you need to do to obey my order”.⁶¹ But since what you need to do to obey the order is exercise today, one might argue that the officer just ordered you in a roundabout way to exercise today; as a result, at noon you acquired only an *unconditional* obligation to exercise today, not a *conditional* obligation to run faster than light tomorrow if you fail to exercise today. Before I reply to this argument, I examine a second example.

Second example. Suppose that in the morning you have a conditional obligation (because you have promised) to run in tomorrow’s marathon if you do not withdraw from the marathon today, but you have an accident at noon that leaves you paralyzed and thus unable to run in tomorrow’s

⁶⁰ Similar remarks apply to the special case of unconditional obligations (§2.6): although OICS and OIPV are not symmetric in the second sense (because OIPV cannot be equivalently reformulated in terms of *satisfaction*), OICS and the *restriction* of OIPV to unconditional obligations (namely the principle that, if an agent at a given time has an *unconditional* obligation, then the obedience proposition of the obligation is not ability-necessary for the agent at that time) *are* symmetric in the second sense (because the restriction of OIPV *can* be equivalently reformulated in terms of satisfaction: for unconditional obligations, obedience amounts to satisfaction).

⁶¹ Could one similarly argue that (contrary to what I claimed in §3.3) it is possible for you to have a conditional obligation to *not* run faster than light tomorrow if you fail to exercise today? No: if your commanding officer gives you a corresponding order, then (by OIPV) you do not acquire any obligation, since there is nothing you need to do to obey the order.

marathon (but still able to withdraw). Then, one might argue, you lose the conditional obligation at noon, but neither OICO nor OIPV explains why you lose it: you can still actualize the obedience proposition of the obligation, namely the proposition that either you withdraw from the marathon today or you run in the marathon tomorrow (since you can still withdraw from the marathon today), and that proposition is still not inevitable for you (since it would be false if you were to do something that you still can do, namely refrain from withdrawing from the marathon today). I reply that you do *not* lose the conditional obligation at noon. To see this, note first that at noon you become obligated to withdraw from the marathon today. Now suppose that at 3pm, not having yet withdrawn from the marathon, you suddenly recover from your paralysis and thus you become again able to run in the marathon tomorrow. Then at 3pm you stop being obligated to withdraw from the marathon today. But your ability to withdraw remains unimpaired, so why do you stop being obligated to withdraw? The only plausible explanation that I see is this: until 3pm, you are obligated to withdraw only because *both* (1) you have the conditional obligation to run if you do not withdraw and (2) you are unable to run, so when you stop (at 3pm) being unable to run you also stop being obligated to withdraw. But this explanation entails that you have the conditional obligation until 3pm,⁶² and thus that (as I claimed) you do not lose it at noon.⁶³ Similarly, going back to the example in the previous paragraph, suppose that at 3pm, not having yet exercised today, you suddenly become able to run faster than light tomorrow (this is logically *possible*). Then at 3pm you stop being obligated to exercise today, and the only plausible explanation that I see is this: until 3pm, you are obligated to exercise today only because *both* (1') you have the conditional obligation to run faster than light tomorrow if you fail to exercise today and (2') you are unable to run faster than light tomorrow, so when you stop being unable to run faster than light tomorrow you also stop being obligated to exercise today.⁶⁴ But this explanation entails that you have the conditional obligation until 3pm.⁶⁵

⁶² One might propose the following alternative explanation: until 3pm, you are obligated to withdraw only because *both* (1*) you have an *unconditional* obligation to either withdraw or run—this is an *unconditional counterpart* (see §3.2) of the conditional obligation in (1) in the text—and (2) you are unable to run, so when you stop being unable to run you also stop being obligated to withdraw. In reply, I ask: why do you have such an unconditional obligation? If one answers by claiming that whenever you have a conditional obligation to do B if you do not do A you also have an unconditional obligation to do either A or B, I reply that this answer in effect grants that you have the conditional obligation until 3pm. But if one alternatively answers by claiming that whenever for example you promise to do B if you do not do A you acquire an unconditional obligation to do either A or B but you do not acquire a conditional obligation to do B if you do not do A, I reply that it is implausible to deny (as this answer in effect does) that people ever acquire conditional obligations by making conditional promises.

⁶³ Could one similarly argue that (contrary to what I claimed in §3.1 and §3.2), if you become paralyzed at noon, you do not lose the conditional obligation to run in tomorrow's marathon if it does not snow tonight? No. My argument in the text starts with the premise that, if you become paralyzed at noon, it becomes obligatory for you (i.e., you become obligated) to withdraw from the marathon today. But a similar argument concerning the example in §3.1 and §3.2 does not get off the ground: if you become paralyzed at noon, it does *not* become obligatory for you that it snows tonight, since you cannot make it the case that it snows tonight.

⁶⁴ I say that until 3pm you are (unconditionally) *obligated* to exercise today (or, in the second example, to withdraw from the marathon today) because I do not wish to take a stand on whether until 3pm you also have an *unconditional obligation* to exercise today (see note 39). I do not need to take a stand on this because, if until 3pm you also have such an unconditional obligation, I can still claim that until 3pm you are obligated to exercise today only because both (1') and (2') hold (since I can claim that until 3pm you have such an unconditional obligation only because both (1') and (2') hold).

⁶⁵ This argument relies on the assumption that running faster than light is logically *possible*, so the argument does not show that it is possible for you to have a conditional obligation with a logically *impossible* satisfaction proposition; for example, a conditional obligation to fail to torture if you torture. Nevertheless, my argument about the pos-

Third example. One might argue that neither OICO nor OIPV explains why you (now) have no conditional obligation to run in tomorrow's marathon *if* you are struck with total paralysis tonight: you *can* actualize (or so one might argue) the obedience proposition of such an obligation, namely the proposition that either (1) you are not struck with total paralysis tonight or (2) you run in tomorrow's marathon (assuming that now you can run in tomorrow's marathon), and that proposition is *not* inevitable for you (since it *might* be false if you were to do something that you can do, namely pray to be struck with total paralysis tonight). I reply that OICO does explain why you have no such conditional obligation: you *cannot* actualize the obedience proposition of such an obligation, namely the disjunction of (1) with (2) above, even if you can actualize (2). To see this, note first that it is fallacious to infer the conclusion that you can actualize a disjunction from the premise that you can actualize one of the disjuncts. For example, even if you can run in tomorrow's marathon, you cannot actualize the proposition that either the sun sets tonight or you run in tomorrow's marathon, because it is inevitable for you that you will not actualize it: it is inevitable for you that it will become settled tonight (when the sun sets) without your involvement. Similarly, you cannot actualize the disjunction of (1) with (2), because it is inevitable for you that you will not actualize it:⁶⁶ (a) if you were *not* struck with total paralysis tonight, then the disjunction would become settled tonight without your involvement (so you would not actualize it),⁶⁷ but (b) if you *were* struck with total paralysis tonight, then you would not run in tomorrow's marathon, and then the disjunction would be false (so again you would not actualize it).⁶⁸

sibility of being given a corresponding order by your commanding officer provides a reason for thinking that it is possible for you to have such a conditional obligation. Note that OICS (and thus also OICO) entails that no agent ever has an *unconditional* obligation with a logically impossible satisfaction proposition.

⁶⁶ I am not relying on the fallacious inference from the premise that you *will not* actualize a proposition to the conclusion that you *cannot* actualize it; I am relying instead on the valid inference from the premise that it is *inevitable* (i.e., ability-necessary) for you that you will not actualize a proposition to the conclusion that you *cannot* actualize it. To see that the latter inference is valid, note that its contrapositive inference is valid: the premise that you *can* actualize a proposition *P* entails the conclusion that the proposition *Q* that you actualize *P* is *ability-possible* for you (since *Q* might—indeed, *would*—be true if you were to do something that you can do, namely actualize *P*).

⁶⁷ I assume that you *cannot* make it the case that you are *not* struck with total paralysis tonight; if you *can*, then it is possible for you to have a conditional obligation to run in tomorrow's marathon if you are struck with total paralysis tonight (see my argument in the text concerning the first example, on running faster than light tomorrow if you fail to exercise today).

⁶⁸ One might propose a variant of the third example: one might argue that neither OICO nor OIPV explains why you (now) have no conditional obligation to beat me at chess tomorrow *if* you have a mild stroke tonight that leaves you able to play chess at an elementary level but unable to beat an experienced player like me. I grant that my reply in the text to the third example does not work here: it is *not* inevitable for you that you will not actualize the obedience proposition of such an obligation—namely the proposition that either (1') you do not have such a stroke tonight or (2') you beat me at chess tomorrow—because, if you had such a stroke tonight, you *might* still beat me tomorrow (if I were to—have a similar stroke and thus—play unusually poorly). Nevertheless, it seems clear that you cannot actualize the disjunction of (1') with (2'): if you did *not* have a stroke tonight, then the disjunction would become settled tonight without your involvement (so you would not actualize it), but if you *did* have a stroke tonight, then you would become unable to beat me at chess tomorrow (so you would be unable to actualize the disjunction). My claim that you cannot actualize the above disjunction is compatible with my claim that you *can* actualize the disjunction of (1) with (2) in the penultimate paragraph of §3.2.

4. Conclusion

Say that a proposition is *feasible*—in other words, is an *option*—for an agent at a given time exactly if the agent at that time can actualize the proposition. Say also that a proposition is a *live option* for an agent at a given time exactly if the proposition is feasible but not inevitable for the agent at that time.⁶⁹ The main conclusion of this paper, namely the conjunction of OICO with OIPV, can be equivalently reformulated as the *ought-implies-live-option* principle:

(OILO) If an agent at a given time has an obligation, then the obedience proposition of the obligation is a live option (i.e., is feasible but not inevitable) for the agent at that time.

This is my preferred ought-implies-can principle. This principle is formulated in terms of *all* obligations that agents have—including *conditional* obligations, which have been neglected in the literature on ought-implies-can principles. I defended OILO (via defending OICO and OIPV) against several objections, and I argued that in many (non-revocation) cases OILO explains why an agent loses or fails to have an obligation.

Perhaps more important than the particular principle that I defended (namely OILO) is the *methodology* that I introduced in this paper for formulating and assessing ought-implies-can principles. In the literature, ought-implies-can principles are typically formulated in terms of what an agent *ought* or is *obligated* to do (at a given time). By contrast, I formulated ought-implies-can principles in terms of what *obligations* an agent has (at a given time). One might argue—and I can grant (though see note 39)—that there is no substantive difference between these two kinds of formulations: an agent is obligated to ϕ exactly if the agent has an obligation to ϕ . Nevertheless, I reply, because obligations—as opposed to being obligated—can be *satisfied*, *violated*, or *obeyed*, formulating ought-implies-can principles in terms of obligations has a major advantage: it enables one to view such principles as *providing conditions on the satisfaction, violation, or obedience proposition of an obligation* that are necessary for an agent to have the obligation at a given time.⁷⁰ And, as I hope this paper demonstrates, this way of viewing ought-implies-can principles is fruitful: it enables one to ask questions that the literature does not address. (For example: for an agent to have an obligation at a given time, must the *satisfaction* or the *obedience* proposition of the obligation be feasible for the agent at that time?) Concerning finally the *assessment* of ought-implies-can principles, the main methodological innovation of this paper con-

⁶⁹ Cf. Greenspan 1975: 264, 265 n. 7. A live option can be contrasted with a *dead option*, namely a feasible but inevitable proposition (for example, the proposition that either you raise your arm today or the sun rises tomorrow: you can actualize this proposition, since you can raise your arm today, but if you do not, the proposition will become settled anyway). A live option need not be a *controllable* proposition, namely a feasible proposition whose negation is also feasible. For example, the proposition that Smith does not win the election is a live option for you (it is not inevitable for you, and you can actualize it because you can kill Smith, who would otherwise win), but is not controllable by you (you cannot actualize its negation, namely the proposition that Smith wins the election). See also note 48.

⁷⁰ One might argue that formulating ought-implies-can principles in terms of obligations also has a major disadvantage: because (1) being all-things-considered obligated to ϕ does not require having an obligation to ϕ , (2) ought-implies-can principles formulated in terms of obligations fail to explain why being all-things-considered obligated to ϕ requires being able to ϕ . One might support (1) as follows. Suppose that you are all-things-considered obligated to both exercise and diet just because you have promised your mother to exercise and you have promised your father to diet. Then, one might argue, you have no obligation to both exercise and diet: no such obligation is generated by your promise to your mother (which generates instead an obligation for you to exercise), and no such obligation is generated by your promise to your father (which generates instead an obligation for you to diet). (Cf. Brandt 1964: 378.) I reply that the fact that you have made *both* promises does generate an obligation for you to both exercise and diet. So I claim that (1) is false, and thus that so is (2).

sists in assessing such principles in terms of their *explanatory power*; more specifically, their power to explain why agents lose or fail to have obligations. To my knowledge, the question of whether an ought-implies-can principle explains why agents lose (i.e., stop having) obligations has not been raised in the literature. I hope that this paper demonstrates the fruitfulness of raising it.

REFERENCES

- Adams, Robert Merrihew (1985). Involuntary sins. *The Philosophical Review*, 94, 3-31.
- al-Hibri, Azizah (1978). *Deontic logic: A comprehensive appraisal and a new proposal*. Washington, DC: University Press of America.
- Alston, William P. (1985). Concepts of epistemic justification. *The Monist*, 68, 57-89.
- Alston, William P. (1988). The deontological conception of epistemic justification. In J. E. Tomberlin (Ed.), *Philosophical perspectives: Vol. 2. Epistemology* (pp. 257-299). Atascadero, CA: Ridgeview.
- Alston, William P. (2005). *Beyond "justification": Dimensions of epistemic evaluation*. Ithaca, NY: Cornell University Press.
- Altham, J. E. J. (1985). Wicked promises. In I. Hacking (Ed.), *Exercises in analysis: Essays by students of Casimir Lewy* (pp. 1-21). New York: Cambridge University Press.
- American Psychiatric Association (2013). *Diagnostic and statistical manual of mental disorders: DSM-5* (5th ed.). Arlington, VA: American Psychiatric Association.
- Anderson, Alan Ross (1967). The formal analysis of normative systems. In N. Rescher (Ed.), *The logic of decision and action* (pp. 147-213). Pittsburgh: University of Pittsburgh Press. (Originally published 1956.)
- Andrić, Vuko (2016). Is objective consequentialism compatible with the principle that "ought" implies "can"? *Philosophia: Philosophical Quarterly of Israel*, 44, 63-77.
- Andrić, Vuko (2017). Objective consequentialism and the rationales of ' "ought" implies "can" '. *Ratio (new series)*, 30, 72-87.
- Åqvist, Lennart (1984). Deontic logic. In D. Gabbay & F. Guenther (Eds.), *Handbook of philosophical logic* (Vol. 2, pp. 605-714). Dordrecht: Reidel.
- Åqvist, Lennart (1987). *Introduction to deontic logic and the theory of normative systems*. Napoli: Bibliopolis.
- Atkinson, R. F. (1958). The autonomy of morals. *Analysis*, 18, 57-62.
- Baier, Annette C. (1970). Act and intent. *The Journal of Philosophy*, 67, 648-658.
- Bailhache, Patrice (1991). *Essai de logique déontique*. Paris: Librairie Philosophique J. Vrin.
- Belnap, Nuel, Perloff, Michael, & Xu, Ming (2001). *Facing the future: Agents and choices in our indeterminist world*. New York: Oxford University Press.
- Bennett, Jonathan (1990). Why is belief involuntary? *Analysis*, 50, 87-107.
- Bergström, Lars (1996). Reflections on consequentialism. *Theoria: A Swedish Journal of Philosophy*, 62, 74-94.
- Besch, Thomas M. (2011). Factualism, normativism and the bounds of normativity. *Dialogue: Canadian Philosophical Review*, 50, 347-365.
- Bloomfield, Paul (2007). Two dogmas of metaethics. *Philosophical Studies*, 132, 439-466.
- Bondy, Patrick (2015). Epistemic deontologism and strong doxastic voluntarism: A defense. *Dialogue: Canadian Philosophical Review*, 54, 747-768.
- Brandt, Richard B. (1964). The concepts of obligation and duty. *Mind*, 73, 374-393.
- Brennan, Geoffrey, & Southwood, Nicholas (2007). Feasibility in action and attitude. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, & D. Egonsson (Eds.), *Hommage à Wlodek: Philosophical papers dedicated to Wlodek Rabinowicz*. <http://www.fil.lu.se/hommageawlodek>.
- Brown, James (1977). Moral theory and the ought—can principle. *Mind*, 86, 206-223.
- Brown, Mark A. (2000). Conditional obligation and positive permission for agents in time. *Nordic Journal of Philosophical Logic*, 5, 83-111.
- Buckwalter, Wesley, & Turri, John (2015). Inability and obligation in moral judgment. *PLoS ONE*, 10(8), e0136589.
- Bykvist, Krister (2007). Violations of normative invariance: Some thoughts on shiftiness. *Theoria: A Swedish Journal of Philosophy*, 73, 98-120.
- Bykvist, Krister, & Hattiangadi, Anandi (2007). Does thought imply ought? *Analysis*, 67, 277-285.
- Card, Claudia (1988). Gratitude and obligation. *American Philosophical Quarterly*, 25, 115-127.
- Carlson, Erik (1999). The oughts and cans of objective consequentialism. *Utilitas*, 11, 91-96.
- Carlson, Erik (2000). Incompatibilism and the transfer of power necessity. *Noûs*, 34, 277-290.
- Castañeda, Hector-Neri (1970). On the semantics of the ought-to-do. *Synthese*, 21, 449-468.
- Chisholm, Roderick M. (1968). Lewis' ethics of belief. In P. A. Schilpp (Ed.), *The philosophy of C. I. Lewis* (pp. 223-242). LaSalle, IL: Open Court.
- Chituc, Vladimir, Henne, Paul, Sinnott-Armstrong, Walter, & De Brigard, Felipe (2016). Blame, not ability, impacts moral "ought" judgments for impossible actions: Toward an empirical refutation of "ought" implies "can". *Cognition*, 150, 20-25.
- Chrisman, Matthew (2008). Ought to believe. *The Journal of Philosophy*, 105, 346-370.
- Chuard, Philippe, & Southwood, Nicholas (2009). Epistemic norms without voluntary control. *Noûs*, 43, 599-632.

- Cohen, Yishai (2017). An analysis of recent empirical data on 'ought' implies 'can'. *Philosophia*, doi 10.1007/s11406-017-9892-2.
- Cooper, Neil (1966). Some presuppositions of moral judgments. *Mind*, 75, 45-57.
- Copp, David (1997). Defending the principle of alternate possibilities: Blameworthiness and moral responsibility. *Noûs*, 31, 441-456.
- Copp, David (2008). 'Ought' implies 'can' and the derivation of the Principle of Alternate Possibilities. *Analysis*, 68, 67-75.
- Côté-Bouchard, Charles (2017). 'Ought' implies 'can' against epistemic deontologism: Beyond doxastic involuntarism. *Synthese*, doi 10.1007/s11229-017-1531-8.
- Dahl, Norman O. (1974). "'Ought' implies 'can'" and deontic logic. *Philosophia: Philosophical Quarterly of Israel*, 4, 485-511.
- Driver, Julia (2011). Promising too much. In H. Sheinman (Ed.), *Promises and agreements: Philosophical essays* (pp. 183-197). New York: Oxford University Press.
- Engel, Pascal (2009). Epistemic responsibility without epistemic agency. *Philosophical Explorations*, 12, 205-219.
- Feis, Guglielmo (2014). The OIC/PAP dispute: Two ways of interpreting 'ought' implies 'can'. In S. Bonicalzi, L. Caffo, & M. Sorgon (Eds.), *Naturalism and constructivism in metaethics* (pp. 172-181). Newcastle upon Tyne, UK: Cambridge Scholars Publishing.
- Feldman, Richard (1988a). Epistemic obligations. In J. E. Tomberlin (Ed.), *Philosophical perspectives: Vol. 2. Epistemology* (pp. 235-256). Atascadero, CA: Ridgeview.
- Feldman, Richard (1988b). Subjective and objective justification in ethics and epistemology. *The Monist*, 71, 405-419.
- Feldman, Richard (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60, 667-695.
- Feldman, Richard (2001). Voluntary belief and epistemic evaluation. In M. Steup (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 77-92). New York: Oxford University Press.
- Feldman, Richard (2008). Modest deontologism in epistemology. *Synthese*, 161, 339-355.
- Feldman, Richard, & Conee, Earl (1985). Evidentialism. *Philosophical Studies*, 48, 15-34.
- Føllesdal, Dagfinn, & Hilpinen, Risto (1971). Deontic logic: An introduction. In R. Hilpinen (Ed.), *Deontic logic: Introductory and systematic readings* (pp. 1-35). Dordrecht: Reidel.
- Frankena, William K. (1963). Obligation and ability. In M. Black (Ed.), *Philosophical analysis: A collection of essays* (pp. 148-165). Englewood Cliffs, NJ: Prentice-Hall. (Originally published 1950.)
- García, J. L. A. (1986). The *Tunsollen*, the *Seinsollen*, and the *Soseinsollen*. *American Philosophical Quarterly*, 23, 267-276.
- Geach, Peter T. (1991). Whatever happened to deontic logic? In P. T. Geach (Ed.), *Logic and ethics* (pp. 33-48). Dordrecht: Kluwer. (Originally published 1982.)
- Gensler, Harry J. (1996). *Formal ethics*. New York: Routledge.
- Ginet, Carl (1990). *On action*. New York: Cambridge University Press.
- Ginet, Carl (2001). Deciding to believe. In M. Steup (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 63-76). New York: Oxford University Press.
- Goldstick, D. (2010). Does epistemic "ought" imply "can"? *Dialogue: Canadian Philosophical Review*, 49, 155-158.
- Graham, Peter A. (2011a). Fischer on blameworthiness and "ought" implies "can". *Social Theory and Practice*, 37, 63-80.
- Graham, Peter A. (2011b). 'Ought' and ability. *The Philosophical Review*, 120, 337-382.
- Greenspan, Patricia S. (1975). Conditional oughts and hypothetical imperatives. *The Journal of Philosophy*, 72, 259-276.
- Guendling, John E. (1974). Modal verbs and the grading of obligations. *The Modern Schoolman*, 51, 117-138.
- Guigon, Ghislain (2009). Bringing about and conjunction: A reply to Bigelow on omnificence. *Analysis*, 69, 452-458.
- Haines, Nicolas (1972). Ought and can. *Philosophy*, 47, 263.
- Haji, Ishtiyaque (1994). Doing the best one can and the principle of alternative possibilities. *Southwest Philosophy Review*, 10, 113-127.
- Haji, Ishtiyaque (1999). Moral anchors and control. *Canadian Journal of Philosophy*, 29, 175-203.
- Haji, Ishtiyaque (2002). *Deontic morality and control*. New York: Cambridge University Press.
- Haji, Ishtiyaque (2012). *Reason's debt to freedom: Normative appraisals, reasons, and free will*. New York: Oxford University Press.
- Hampshire, Stuart (1951). Freedom of the will. *The Aristotelian Society: Supplementary Volume*, 25, 161-178.
- Hannon, Michael (2017). Intuitions, reflective judgements, and experimental philosophy. *Synthese*, doi 10.1007/s11229-017-1412-1.
- Hansson, Sven Ove (1999). But what should I do? *Philosophia: Philosophical Quarterly of Israel*, 27, 433-440.
- Hansson, Sven Ove (2013). The varieties of permission. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (pp. 195-240). London: College Publications.
- Hare, Richard M. (1963). *Freedom and reason*. Oxford: Clarendon Press.
- Harman, Gilbert (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Hasker, William (1989). *God, time, and knowledge*. Ithaca, NY: Cornell University Press.
- Heller, Mark (2000). Hobartian voluntarism: Grounding a deontological conception of epistemic justification. *Pacific Philosophical Quarterly*, 81, 130-141.
- Henderson, G. P. (1966). "'Ought' implies 'can' ". *Philosophy*, 41, 101-112.

- Henne, Paul, Chituc, Vladimir, De Brigard, Felipe, & Sinnott-Armstrong, Walter (2016). An empirical refutation of ‘ought’ implies ‘can’. *Analysis*, 76, 283-290.
- Heuer, Ulrike (2010). Reasons and impossibility. *Philosophical Studies*, 147, 235-246.
- Hobbs, James Austin (2013). *“Ought” claims and blame in a deterministic world*. Doctoral dissertation, Cornell University.
- Hoffman, Joshua, & Rosenkrantz, Gary (1984). Hard and soft facts. *The Philosophical Review*, 93, 419-434.
- Horty, John F. (2001). *Agency and deontic logic*. New York: Oxford University Press.
- Howard-Snyder, Frances (2006). “Cannot” implies “not ought”. *Philosophical Studies*, 130, 233-246.
- Howard-Snyder, Frances (2013). Ought implies can. In H. LaFollette (Ed.), *International encyclopedia of ethics* (Vol. 6, pp. 3748-3757). Oxford: Blackwell.
- Humberstone, I. L. (1971). Two sorts of ‘ought’s. *Analysis*, 31, 8-11.
- Jacobs, Russell A. (1985). Is “ought implies can” a moral principle? *Southwest Philosophy Review*, 2, 43-54.
- Kekes, John (1984). ‘Ought implies can’ and two kinds of morality. *The Philosophical Quarterly*, 34, 459-467.
- Kenny, Anthony (1975). *Will, freedom and power*. Oxford: Blackwell.
- Kenny, Anthony (1976). Human abilities and dynamic modalities. In J. Manninen & R. Tuomela (Eds.), *Essays on explanation and understanding: Studies in the foundations of humanities and social sciences* (pp. 209-232). Dordrecht: Reidel.
- King, Alex (2014). Actions that we ought, but can’t. *Ratio (new series)*, 27, 316-327.
- King, Alex (2017). ‘Ought implies can’: Not so pragmatic after all. *Philosophy and Phenomenological Research*, 95, 637-661.
- Kornblith, Hilary (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, 92, 33-48.
- Kornblith, Hilary (2001). Epistemic obligation and the possibility of internalism. In A. Fairweather & L. Zagzebski (Eds.), *Virtue epistemology: Essays in epistemic virtue and responsibility* (pp. 231-248). New York: Oxford University Press.
- Krogh, Christen, & Herrestad, Henning (1996). Getting personal: Some notes on the relationship between personal and impersonal obligation. In M. A. Brown & J. Carmo (Eds.), *Deontic logic, agency and normative systems. AEOON ’96: Third international workshop on deontic logic in computer science, Sesimbra, Portugal, 11-13 January 1996* (pp. 134-153). Berlin: Springer.
- Kurthy, Miklos, & Lawford-Smith, Holly (2015). A brief note on the ambiguity of ‘ought’. Reply to Moti Mizrahi’s ‘Ought, can and presupposition: An experimental study’. *Method: Analytic Perspectives*, 4(6), 244-249.
- Kurthy, Miklos, Lawford-Smith, Holly, & Sousa, Paulo (2017). Does ought imply can? *PLoS ONE*, 12(4), e0175206.
- Lange, Marc (2014). Aspects of mathematical explanation: Symmetry, unity, and salience. *The Philosophical Review*, 123, 485-531.
- Lewis, David K. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, David K. (1986). Counterfactuals and comparative possibility. In D. K. Lewis, *Philosophical papers* (Vol. 2, pp. 3-31). New York: Oxford University Press. (Originally published 1973.)
- Lichtenberg, Judith (2010). Oughts and cans: Badness, wrongness, and the limits of ethical theory. *Philosophical Topics*, 38(1), 123-142.
- Lipton, Peter (2009). Causation and explanation. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook of causation* (pp. 619-631). New York: Oxford University Press.
- Littlejohn, Clayton (2009). “Ought,” “can,” and practical reasons. *American Philosophical Quarterly*, 46, 363-372.
- Littlejohn, Clayton (2012). Does ‘ought’ still imply ‘can’? *Philosophia: Philosophical Quarterly of Israel*, 40, 821-828.
- Loewer, Barry, & Belzer, Marvin (1983). Dyadic deontic detachment. *Synthese*, 54, 295-318.
- Lowry, Rosemary (2012). Reasons for action and psychological capacities. *Ethical Theory and Moral Practice*, 15, 521-531.
- Manning, Rita C. (1981). “Ought implies can” and the price of duty. *The Southern Journal of Philosophy*, 19, 117-121.
- Martin, Wayne (2009). Ought but cannot. *Proceedings of the Aristotelian Society*, 109, 103-128.
- Mason, Elinor (2003). Consequentialism and the “ought implies can” principle. *American Philosophical Quarterly*, 40, 319-331.
- McConnell, Terrance (1989). “‘Ought’ implies ‘can’” and the scope of moral requirements. *Philosophia: Philosophical Quarterly of Israel*, 19, 437-454.
- McHugh, Conor (2012). Epistemic deontology and voluntariness. *Erkenntnis*, 77, 65-94.
- McNamara, Paul (2004). Agential obligation as non-agential personal obligation plus agency. *Journal of Applied Logic*, 2, 117-152.
- Mellor, D. H. (2004). For facts as causes and effects. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 309-323). Cambridge, MA: MIT Press.
- Miller, Dale E. (2003). Actual-consequence act utilitarianism and the best possible humans. *Ratio (new series)*, 16, 49-62.
- Mizrahi, Moti (2009). ‘Ought’ does not imply ‘can’. *Philosophical Frontiers*, 4, 19-35.
- Mizrahi, Moti (2012). Does ‘ought’ imply ‘can’ from an epistemic point of view? *Philosophia: Philosophical Quarterly of Israel*, 40, 829-840.
- Mizrahi, Moti (2015a). Ought, can, and presupposition: An experimental study. *Method: Analytic Perspectives*, 4(6), 232-243.
- Mizrahi, Moti (2015b). Ought, can, and presupposition: A reply to Kurthy and Lawford-Smith. *Method: Analytic Perspectives*, 4(6), 250-256.
- Montefiore, Alan (1958). ‘Ought’ and ‘can’. *The Philosophical Quarterly*, 8, 24-40.
- Moore, G. E. (1922). *Philosophical Studies*. New York: Harcourt, Brace & Co.
- Narveson, Jan (1978). Future people and us. In R. I. Sikora & B. Barry (Eds.), *Obligations to future generations* (pp. 38-60). Philadelphia: Temple University Press.

- Nelkin, Dana K. (2011). *Making sense of freedom and responsibility*. New York: Oxford University Press.
- O'Connor, D. J. (1971). *Free will*. New York: Macmillan.
- Oppenheim, Felix E. (1987). National interest, rationality, and morality. *Political Theory*, 15, 369-389.
- Owens, David (2007). Duress, deception, and the validity of a promise. *Mind*, 116, 293-315.
- Peels, Rik (2017). *Responsible belief: A theory in ethics and epistemology*. New York: Oxford University Press.
- Pickard, Hanna (2015). Psychopathology and the ability to do otherwise. *Philosophy and Phenomenological Research*, 90, 135-163.
- Pigden, Charles R. (1990). Ought-implies-can: Erasmus Luther and R. M. Hare. *Sophia*, 29, 2-30.
- Plantinga, Alvin (1993). *Warrant: The current debate*. New York: Oxford University Press.
- Pojman, Louis P. (1993). Believing, willing, and the ethics of belief. In L. P. Pojman (Ed.), *The theory of knowledge: Classical and contemporary readings* (pp. 525-544). Belmont, CA: Wadsworth.
- Prior, Arthur N. (1958). Escapism: The logical basis of ethics. In A. I. Melden (Ed.), *Essays in moral philosophy* (pp. 135-146). Seattle: University of Washington Press.
- Prior, Arthur N. (1962). *Formal logic* (2nd ed.). Oxford: Clarendon Press.
- Qu, Hsueh (2017). Hume's doxastic involuntarism. *Mind*, 126, 53-92.
- Reynolds, Steven L. (2011). Doxastic voluntarism and the function of epistemic evaluations. *Erkenntnis*, 75, 19-35.
- Robinson, Richard (1971). Ought and ought not. *Philosophy*, 46, 193-202.
- Rodriguez-Pereyra, Gonzalo (2006). Truthmaking, entailment, and the conjunction thesis. *Mind*, 115, 957-982.
- Rodriguez-Pereyra, Gonzalo (2009). The disjunction and conjunction theses. *Mind*, 118, 427-443.
- Ross, W. David (1939). *Foundations of ethics: The Gifford lectures delivered in the University of Aberdeen, 1935-6*. Oxford: Clarendon Press.
- Russell, Bruce (2001). Epistemic and moral duty. In M. Steup (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 34-48). New York: Oxford University Press.
- Ryan, Sharon (2003). Doxastic compatibilism and the ethics of belief. *Philosophical Studies*, 114, 47-79.
- Scott-Kakures, Dion (1994). On belief and the captivity of the will. *Philosophy and Phenomenological Research*, 54, 77-103.
- Shaffer, Michael J. (2013). Doxastic voluntarism, epistemic deontology, and belief-contravening commitments. *American Philosophical Quarterly*, 50, 73-81.
- Shah, Nishi (2002). Clearing space for doxastic voluntarism. *The Monist*, 85, 436-445.
- Shah, Nishi, & Velleman, J. David (2005). Doxastic deliberation. *The Philosophical Review*, 114, 497-534.
- Shaw, P. D. (1965). Ought and can. *Analysis*, 25, 196-197.
- Sidgwick, Henry (1981). *The methods of ethics* (7th ed.). Indianapolis: Hackett. (7th ed. originally published 1907.)
- Sinnott-Armstrong, Walter (1984). 'Ought' conversationally implies 'can'. *The Philosophical Review*, 93, 249-261.
- Sinnott-Armstrong, Walter (1988). *Moral dilemmas*. Oxford: Blackwell.
- Skow, Bradford (2014). Are there non-causal explanations (of particular events)? *The British Journal for the Philosophy of Science*, 65, 445-467.
- Smith, Holly M. (1997). A paradox of promising. *The Philosophical Review*, 106, 153-196.
- Smith, James W. (1961). Impossibility and morals. *Mind*, 70, 362-375.
- Southwood, Nicholas (2016). "The thing to do" implies "can". *Noûs*, 50, 61-72.
- Southwood, Nicholas, & Wiens, David (2016). "Actual" does not imply "feasible". *Philosophical Studies*, 173, 3037-3060.
- Stenius, Erik (1963). The principles of a logic of normative systems. *Acta Philosophica Fennica*, 16, 247-260.
- Steup, Matthias (2000). Doxastic voluntarism and epistemic deontology. *Acta Analytica*, 15, 25-56.
- Steup, Matthias (2008). Doxastic freedom. *Synthese*, 161, 375-392.
- Steup, Matthias (2012). Belief control and intentionality. *Synthese*, 188, 145-163.
- Stocker, Michael (1971). 'Ought' and 'can'. *Australasian Journal of Philosophy*, 49, 303-316.
- Streumer, Bart (2007). Reasons and impossibility. *Philosophical Studies*, 136, 351-384.
- Streumer, Bart (2010). Reasons, impossibility and efficient steps: Reply to Heuer. *Philosophical Studies*, 151, 79-86.
- Talbot, Brian (2016). The best argument for "ought implies can" is a better argument against "ought implies can". *Ergo: An Open Access Journal of Philosophy*, 3, 377-402.
- Taylor, Paul W. (Ed.). (1967). *Problems of moral philosophy: An introduction to ethics*. Belmont, CA: Dickenson.
- Thomason, Richmond H. (1981). Deontic logic and the role of freedom in moral deliberation. In R. Hilpinen (Ed.), *New studies in deontic logic: Norms, actions, and the foundations of ethics* (pp. 177-186). Dordrecht: Reidel.
- Tranøy, Knut Erik (1972). 'Ought' implies 'can': A bridge from fact to norm? Part I. *Ratio*, 14, 116-130.
- Turri, John (2017). How "ought" exceeds but implies "can": Description and encouragement in moral judgement. *Cognition*, 168, 267-275.
- van Rijen, Jeroen (1993). Review of Fred Feldman's *Doing the best we can: An essay in informal deontic logic* (Dordrecht: Reidel, 1986). *Noûs*, 27, 264-267.
- van Someren Greve, Rob (2014). 'Ought', 'can', and fairness. *Ethical Theory and Moral Practice*, 17, 913-922.
- Velleman, J. David (1998). Is motivation internal to value? In C. Fehige & U. Wessels (Eds.), *Preferences* (pp. 88-102). Berlin: de Gruyter.
- Vogelstein, Eric (2012). Subjective reasons. *Ethical Theory and Moral Practice*, 15, 239-257.
- von Wright, Georg Henrik (1951). Deontic logic. *Mind*, 60, 1-15.

- von Wright, Georg Henrik (1963). *Norm and action: A logical enquiry*. New York: Humanities Press.
- von Wright, Georg Henrik (1981a). On the logic of norms and actions. In R. Hilpinen (Ed.), *New studies in deontic logic: Norms, actions, and the foundations of ethics* (pp. 3-35). Dordrecht: Reidel.
- von Wright, Georg Henrik (1981b). Problems and prospects of deontic logic: A survey. In E. Agazzi (Ed.), *Modern logic—A survey: Historical, philosophical, and mathematical aspects of modern logic and its applications* (pp. 399-423). Dordrecht: Reidel.
- von Wright, Georg Henrik (1988). Action logic as a basis for deontic logic. In G. di Bernardo (Ed.), *Normative structures of the social world* (pp. 39-63). Amsterdam: Rodopi.
- Vranas, Peter B. M. (2007). I ought, therefore I can. *Philosophical Studies*, 136, 167-216.
- Vranas, Peter B. M. (2008). New foundations for imperative logic I: Logical connectives, consistency, and quantifiers. *Noûs*, 42, 529-572.
- Vranas, Peter B. M. (2011). New foundations for imperative logic: Pure imperative inference. *Mind*, 120, 369-446.
- Vranas, Peter B. M. (2016). New foundations for imperative logic III: A general definition of argument validity. *Synthese*, 193, 1703-1753.
- Vranas, Peter B. M. (2017a). “Ought” implies “can” but does not imply “must”: An asymmetry between becoming infeasible and becoming overridden. *The Philosophical Review* (forthcoming).
- Vranas, Peter B. M. (2017b). Imperative foundations for the metaphysics of obligations. Unpublished.
- Weatherson, Brian (2008). Deontology and Descartes’s demon. *The Journal of Philosophy*, 105, 540-569.
- Wedgwood, Ralph (2013). Rational ‘ought’ implies ‘can’. In E. Sosa & E. Villanueva (Eds.), *Philosophical Issues: Vol. 23. Epistemic agency* (pp. 70-92). Oxford: Wiley Blackwell.
- Williams, Bernard (1973). Deciding to believe. In B. Williams, *Problems of the self: Philosophical papers 1956-1972* (pp. 136-151). New York: Cambridge University Press. (Originally published 1970.)
- Williams, Bernard (1981). *Ought* and moral obligation. In B. Williams, *Moral luck: Philosophical papers 1973-1980* (pp. 114-123). New York: Cambridge University Press.
- Zimmerman, Michael J. (1993). Obligation, responsibility and alternate possibilities. *Analysis*, 53, 51-53.
- Zimmerman, Michael J. (1996). *The concept of moral obligation*. New York: Cambridge University Press.