

EXPRESSIVISM DETRIVIALIZED

JOHN CANTWELL

ABSTRACT

It is argued that David Lewis' two triviality results (the probability of the conditional cannot be the conditional probability; desire cannot be belief) both present a potential problem for expressivism, are related, and can both be resolved in the same way: by allowing for gappy propositions (propositions that can lack truth value). In particular, a semantics for '*A* is good' is provided that allows one to embrace the major premises leading up to Lewis' triviality result while avoiding its conclusion.

1. Introduction

David Lewis (1976; 1988) provided two triviality results that seemingly makes it hard – or at least harder – to be an expressivist. The first result was that for any three propositions *A*, *B* and *C* in a space of propositions, all but trivial probability measures are excluded if one imposes the constraint that $P(C) = P(B|A)$. In particular, only trivial probability measures will allow that $P(A \rightarrow B) = P(B|A)$, where the conditional $A \rightarrow B$ is in the space of propositions. The second result showed that all but trivial pairs of probability measures and value measures are excluded if one imposes the constraint that for any given *A* there is a proposition A° such that the *expected utility* or *desirability* of *A*, $V(A)$, is identical to (or is a positive affine transformation of) the probability $P(A^\circ)$ and that this identity is stable under conditionalization (exactly how A° is to be understood will be discussed, but for heuristic purposes it can be read 'The state of affairs that *A* is good', or '*A* is good' for short).

The first result is troubling for the expressivist who holds that the assertion of a conditional expresses that one takes the consequent to be assertible *given* the antecedent, and so that the degree of assertability of a conditional is the conditional probability of the consequent given its antecedent, and is the same as its probability. This is troubling as it is a direct assault on the *Ramsey Test*, a leading expressivist model for the interpretation of conditionals, but it is also the focus of a substantial part of the expressivist literature on conditionals.¹

¹ Expressivists who regard the conditional as an exceptional construction take Lewis' first triviality result to be a central argument for the exceptional status of the conditional; on

The second result, however, might at first sight not seem at all troubling for the expressivist. Lewis presented the result as a problem for the non-Humean *desire-as-belief* thesis that seeks to subsume the motivational power of desire under belief. On the desire-as-belief thesis, rational desires track judgements of objective good and bad, and this thesis would appear to be the exact opposite of the expressivist *evaluative-belief-as-desire* thesis in which judgements of goodness are taken to be mere expressions of one's motivational state (preferences/desires). The second triviality result might at first seem to hit the desire-as-belief thesis harder as there is a pressure on the desire-as-belief theorist to account for the whole motivational structure of a rational agent in terms of the agent's degrees of belief in evaluative propositions. By contrast, the expressivist is only under pressure to account for how degrees of belief in evaluative propositions – to the extent that there are any such propositions – reflect the agent's motivational structure, with no presumption that the whole motivational structure can be thus revealed. Nevertheless, it will be argued that even this more modest ambition is affected by Lewis' triviality result.

What is ultimately at stake is whether expressivists can account for the fact that agents can *reason* with value judgements when value judgements (e.g. *x is good*) are given an expressive interpretation. A sophisticated agent must have the capacity to form a coherent set of less than certain value judgements, whether these are given an expressive interpretation or not. The most established theory of coherent combinations of less than certain judgements is probability theory. By insisting that value judgements do no more than express the agent's preferences, the expressivist – much like the desire-as-belief theorist – takes value judgements and preferences to be conceptually entangled. Lewis' second triviality result is a direct challenge to the cogency of the very idea of such entanglement. In Lewis' (1996, p. 308) words, the triviality result is directed at the view that speaks of "beliefs as if they were desires; or of states that occupy a double role, being at once beliefs and desires".

Consider an example. Jane, who is in Stockholm, has placed a bet: if it is raining in both Paris and London, or if it is raining neither in Paris nor in London she will win one million; if it is raining in Paris but not in London, or if it is raining in London but not in Paris, she will win nothing. Winning and losing this bet is all she cares about. Calibrating the desirability measure that we take to describe her preferences we can assign the value 1 to winning the bet and the value 0 to losing the bet. Jane is willing to make value

such a view Lewis' triviality result is not in itself a problem. However, the triviality result is problematic for those that hold that conditionals can embed in complex sentences (negated conditionals, etc.), see discussion in McGee (1989); Edgington (1995); McDermott (1996); Kölbel (2000); Edgington (2000).

judgements like *Winning the bet is good* and *Not winning the bet is bad* and we take her thereby to express that her degree of desirability for winning the bet is 1 and her degree of desirability for losing the bet is 0. Jane doesn't know whether or not it is raining in Paris, so her degree of desirability for (the expected utility of) *It is raining in Paris* (let this be A), $V(A)$, is somewhere between 0 and 1. If, say, she takes there to be an even chance that it is raining in Paris, an even chance that it is raining in London, and takes the probability of rain in Paris to be independent of rain in London, her degree of desirability for rain in Paris, $V(A)$, is 0.5.

Rain in Paris might, by Jane's own lights, be good and might be bad. But what, exactly, should be Jane's degree of belief in the proposition that rain in Paris is good? I want to explore – and ultimately to argue for – the hypothesis that Jane's degree of belief that rain in Paris is good should track the desirability of rain in Paris, that is, that in the situation described we should have the identity

$$V(A) = P(A \text{ is good}) \quad (1)$$

and that this identity should be retained when conditionalising under new evidence (that doesn't alter the value structure of the situation, e.g. by introducing new bets) so that the desirability of A and the degree of belief in the value judgement *A is good* become truly entangled. Why this identity? First of all, there is the expressivist contention that value judgements are nothing else than expressions of one's underlying motivational state, here represented by a desirability measure V ; so Jane's degree of belief that *A is good* should track *some* aspect of her desirability measure. Second, the structure of the situation is very simple, there is only one good thing that can be achieved – winning the bet – so one could expect that Jane's degree of belief that rain in Paris is good should follow her degree of belief that she will win the bet *given* that it is raining in Paris. That is, one would expect the following identity:

$$P(A \text{ is good}) = P(G | A), \quad (2)$$

where G is the proposition that holds if and only if she wins the bet (G can be read "The present situation is good"). Given, in addition, that winning the bet is all Jane cares for, this identity should remain stable under conditionalisation.² As $V(G) = 1$ and $V(\sim G) = 0$ we have, by the additivity of value:

$$V(A) = V(G)P(G | A) + V(\sim G)P(\sim G | A) = P(G | A). \quad (3)$$

² Compare Kanger's (1971) use of a proposition to the effect all normative requirements have been met).

Our identity $V(A) = P(A \text{ is good})$ follows from (2) and (3). Third, there is an intuitive appeal to the identity $V(A) = P(A \text{ is good})$; it makes sense in the initial situation in which there seems to be an even chance that rain in Paris is good. Furthermore, it makes intuitive sense to hold that as the desirability of A rises or lowers with new evidence, the probability that A is good should rise or lower proportionately: whatever evidence that would change the desirability of rain in Paris (without radically altering the situation as described) should change the probability that rain in Paris is good accordingly; this gives us *prima facie* reason to hold that the identity $V(A) = P(A \text{ is good})$ should be stable under conditionalisation.

Now, the intuitions appealed to – we shall soon see – run headlong into a massive obstacle in the shape of two technical impossibility results. And one could leave it at that and dismiss the intuitions. But one should keep in mind that a technical result is very sensitive to the formal framework in which it is presented, and its philosophical significance is sensitive to whether the formal framework is rich enough to model the situation to be analyzed. I will argue that the formal framework in which the impossibility result is framed is not rich enough to do justice to the intuitions involved and that the roadblock can be removed.

Admittedly, intuitions can be unstable and the technical manoeuvre of letting value range in the interval 0 and 1 could seem suspect. Indeed it seems easy to construct intuitive counterexamples to the thesis that $V(A) = P(A \text{ is good})$ even in simple situations. Say that we are in a situation where if A is true then it is highly likely that one will *win* some small amount, however there is also, if A is true, a small probability that one will *lose* a large amount. In such a case it would seem highly likely that A is good (as it is highly likely that one will win) even though the expected utility of A is low. Put some numbers to this. Let $P(\text{win } \text{€}1 | A) = 0.9$ and $P(\text{lose } \text{€}10\,000 | A) = 0.1$. The expected monetary value $EMV(A)$ of A would then be $-\text{€}999.1$ which is ‘low’, yet it seems highly likely (0.9) that A is good. So a ‘low’ utility for A is coupled with a ‘high’ probability that A is good; is this not a direct counterexample to the thesis that the probability that A is good tracks the desirability of A ? No! One can see this by re-calibrating the value scale using the following positive linear transformation of EMV : $V'(B) = (EMV(B) + 10\,000)/10\,001$ for all B . This transformation compresses the scale so that all values come in the interval $[0, 1]$. Now the ‘value’ of losing $\text{€}10\,000$ is 0 while the value of winning $\text{€}1$ remains 1; the expected value of A , $V'(A)$, is 0.9 which on this new compressed scale is ‘low’ or at least ‘bad’ (as it corresponds to an expected monetary loss of $\text{€}999.1$), yet now we have $V'(A) = P(A \text{ is good}) = P(\text{win } \text{€}1 | A)$. That is, the substantive claim is not that $V(A) = P(A \text{ is good})$ regardless of the scale in which value is measured, the claim is that $P(A \text{ is good})$ should *stably track* $V(A)$, so that there is *some* positive linear transformation V' of

V for which the identity $V'(A) = P(A \text{ is good})$ is stable under conditionalization.

So much for intuitions. The claim is that we have a situation in which there is a proposition A such that (for any E such that $P(E \& A) > 0$):³

$$V_E(A) = P_E(A \text{ is good}).$$

But David Lewis (1988) showed that there can be no non-trivial probability measure that allows for such a stable relationship in any situation. One way to see this (and this is one way in which Lewis presented his second triviality result) is to appeal to Lewis' first triviality result. For we have (for any E such that $P(E \& A) > 0$):

$$V_E(A) = P_E(A \text{ is good}) = P_E(G | A).$$

As $A \text{ is good}$ is a proposition with a probability $P(A \text{ is good})$ that tracks the conditional probability $P(G | A)$ under arbitrary conditionalization, this is subject to Lewis' first triviality result: if $P(A) > 0$ then $P(G) = 1$.

So we have at least one situation in which apparently reasonable expressivist intuitions collapse into triviality. Is this really a problem? One might object that the original example, which only involves two levels of value, is too simplistic to be of any relevance. So refine the situation and assume that if it is raining in Rome all other bets are called off, but if it is also raining in Paris and London Jane will still win a half a million. We now need a more fine grained vocabulary to distinguish the different outcomes. Jane judges that winning half a million is *merely ok* and her normalized value for such an outcome is 0.5. So now rain in Paris might be good, might be bad, and might be merely ok, and so now the desirability of rain in Paris can no longer be expected to follow the probability that rain in Paris is good, for the desirability of rain in Paris must also track the probability that rain in Paris is merely ok. Would we not, however, expect the following more complex relationship to hold: $V(A) = P(A \text{ is good}) + P(A \text{ is merely ok}) \times .5$? Again, Lewis' triviality result (in its extended form) shows that the requirement that such a relationship is stable under conditionalisation collapses into triviality.

The result is quite general. If the vocabulary for making value judgements (*good*, *merely ok*, etc.) is sufficiently fine grained to distinguish between the desirability of (bundles of) the basic objects of desire (in this case, the desirability of winning one million, winning half a million, and winning nothing), the desirability of a proposition A cannot be correlated with the probabilities that A will have one of these values in the way that

³ Notation: $P_E(A) = P(A | E) = P(A \& E) / P(E)$, while $V_E(A) = V(A \& E) \times P(A | E)$.

one – as an expressivist – would expect. The worry is that there might be something fundamentally flawed with the idea that pure value judgements (e.g. *A is good*) track desirability *and nothing else*; the result seemingly implies that value judgements and preferences cannot be conceptually entangled in the way that the expressivist would require.

2. What to do?

Price (1989) suggests that one can take the expressivist to be committed to no more than the identity $V(A) = P(A \text{ is good} \mid A)$, an identity that is not by itself susceptible to Lewis' triviality result. Price notes that the triviality result would only reappear if we impose the invariance requirement that $P(A \text{ is good} \mid A) = P(A \text{ is good})$ and suggests that we should not impose such a requirement. Now, it is true that if the identity $P(A \text{ is good} \mid A) = P(A \text{ is good})$ does not hold then we can avoid the triviality result. But such a fix does not explain what went wrong in our original example. Note that the identity $V(A) = V_A(A)$ is always stable: the desirability of *A* is never affected by learning that *A* is true. So one would at the very least require some explanation for why the identity $P(A \text{ is good} \mid A) = P(A \text{ is good})$ can never (except in trivial situations) be stable, an explanation not by appeal to some structural probabilistic result, but in terms of the meaning of *A is good*. As long as the sole reason for questioning the stable identity $P(A \text{ is good} \mid A) = P(A \text{ is good})$ is to avoid the triviality result we have not dispelled the worry that there is something fundamentally wrong about taking value judgements and preferences to be conceptually entangled. Lewis' triviality result is only avoided if in every non-trivial situation there is some evidence that could make the question whether *A* is good become probabilistically dependent on whether *A* is true. Until we can account for what such evidence could be – in terms amenable to the expressivist – we are not out of harms way.

The expressivist has a strategy available that escapes both triviality results: deny that conditional sentences and evaluative sentences are in the domain of the probability and desirability measures, on the ground that they lack propositional content. For instance, Adams (1975) takes the conditional probability $P(B \mid A)$ to determine the degree of *assertability* of the conditional $A \rightarrow B$ but leaves conditional propositions outside the domain of the probability function itself. This respects the expressivist intuition that conditional and evaluative sentences are exceptional, but the strategy is problematic; for when a measure of degree of assertability is extended to cover sentences of arbitrary logical complexity one would expect a probability-like structure to emerge and the danger is that Lewis' triviality results – that apply to certain mathematical structures regardless of how these structures are interpreted –

would reappear but now directed at degrees of assertability. Some have argued on this basis that the conditional doesn't meaningfully embed in sentences of arbitrary logical complexity (e.g. Edgington (1995)). Such a move is controversial (see e.g. McDermott (1996); Kölbel (2000)) but would block Lewis' triviality result for conditionals, however no one denies that *evaluative* sentences embed; so evaluative expressivism still must face up to the challenge.

Bradley and List (2009) argue that Lewis' triviality result can be blocked if we restrict the desire-as-belief thesis to what they call 'purely non-evaluative' propositions. They propose that we can think of a world as consisting of an ordered pair (ϕ, ψ) where ϕ is the non-evaluative state of the world and ψ is the evaluative state of the world. A purely non-evaluative proposition is one whose truth value at a world depends only on the non-evaluative state of the world and a purely evaluative proposition is one whose truth value at a world depends only on the evaluative state of the world. Bradley and List argue that if we take all purely non-evaluative propositions to be probabilistically independent of all purely evaluative propositions, then it is possible to let $V(A) = P(A \text{ is good})$ for all purely non-evaluative propositions A , thus blocking Lewis' triviality result for a non-trivial space of propositions. However, neither the desire-as-belief theorist nor the expressivist evaluative-belief-as-desire theorist can rest content with this. The problem is the exceedingly strong assumptions needed to avoid the triviality result. Typically, an item of purely non-evaluative news can be a strong indicator (evidence for) that something good or bad has happened; that is, typically, purely non-evaluative propositions aren't probabilistically independent of purely evaluative propositions. Indeed it is difficult to think of a situation where such independence – which would be stable under conditionalisation – would obtain. In any case, it does not help us explain what happened in the original example. 'Rain in Paris is good' is not probabilistically independent of 'It is raining in London'. On acquiring evidence that it is raining in London, the putatively purely non-evaluative 'It is raining in Paris' is not probabilistically independent of the putatively purely evaluative 'The current situation is good', which means that the situation does not satisfy Bradley and List's constraint of complete independence; so we have no explanation of what went wrong in the example.

Hájek and Pettit (2004) have put forward a semantic analysis of *A is good* that avoids the triviality thesis in a way that (they suggest) could be amenable to the expressivist. Their suggestion is that *A is good* contains a hidden indexical and that its truth conditions should be relativized to an agent relative expected utility measure: the truth value of an assertion of '*A is good*', just like the truth value of an assertion 'I am hungry' depends on who makes the utterance. They take aim on the fact that Lewis' triviality result establishes that there cannot be a *single* proposition with a probability

that stably tracks the expected utility of A under conditionalization. This leaves it open that there might be a class of propositions such that when conditionalizing on E there is some proposition in the class that has a probability that coincides with the expected utility of A given E . Hájek and Petit point out that their suggestion has the desired consequence: there will be non-trivial combinations of probability and expected utility measures in which one can have $V_E(A) = P_E(A \text{ is good})$ for arbitrary E , as the proposition expressed by ' $A \text{ is good}$ ' will vary with the evidence E . However their analysis comes at a high cost.

First of all, Hájek and Petit do not specify how the proposition expressed by $A \text{ is good}$ is supposed to vary with the expected utility measure; all they point to is the *possibility* that for any E there will be some proposition X such that $P_E(X) = V_E(A)$. They do not specify truth conditions for $A \text{ is good}$ that would deliver such an X for each desirability measure or explain how it is supposed to be related to A . In fact, the existence of such an X is *not* guaranteed;⁴ so their suggestion contains the implicit substantial assumption that value judgements involve something more than the preferences of the agent.

Second, Hájek's and Petit's proposal forces expressivism into a doubly-relativist position. The run-of-the-mill expressivist may concede that the truth value of $A \text{ is good}$ can vary with the agent's fundamental or non-instrumental values (more on this below). For instance, the expressivist will hold that the winner of a bet can accept a sentence – 'This is a good outcome' – that the loser rejects, even though both winner and loser agree on all factual matters and even though they are both linguistically competent speakers. Indeed the expressivist may even concede that the truth value of 'This is a good outcome' – to the extent that it makes sense to speak of the truth value of an evaluative sentence – varies with the fundamental (non-instrumental) values of the speaker. But Hájek and Petit go further and relativize the truth of $A \text{ is good}$ to the agent's degrees of belief (which makes it 'doubly' relativist). This goes beyond what I take to be the basic thrust of meta-ethical expressivism.

Here's why. In our example, rain in Paris is at best an instrumental good, and then only if it is raining in London. Thus two speakers with the same

⁴ Consider the three element model $U = \{u_1, u_2, u_3\}$. Let the power set of U be the set of propositions, and let $P(\{u_1\}) = 0.6$, $P(\{u_2\}) = 0.1$, $P(\{u_3\}) = 0.3$, and, for $X \subseteq U$, $P(X) = \sum_{u \in X} P(\{u\})$. Let $V(\{u_1\}) = V(\{u_3\}) = 1$ and $V(\{u_2\}) = 0$. Then $V(\{u_1, u_2\}) = 6/7$ and $V(\{u_2, u_3\}) = 1/4$. However, there is no proposition in this space that has probability $6/7$ or $1/4$. Furthermore, take any positive linear transformation V' of V ; if there is some proposition X in the space such that $P(X) = V'(\{u_1, u_2\})$, then there is *no* proposition Y such that $P(Y) = V'(\{u_2, u_3\})$ (the proof is somewhat tedious as one has to go through all combinations of propositions X and Y – 64 cases in all – but quite straightforward). In this model there just aren't enough propositions and probabilities to go around.

non-instrumental values (e.g. two speakers who both want Jane to experience the pleasure of winning) – but with different evidence for rain in London – can disagree about whether rain in Paris is good or not. In a case like this the run-of-the-mill expressivist would hold that the disagreement can be fully explained by the disagreement over factual (non-evaluative) matters. Thus – holding steady the shared non-instrumental values – the relativisation of the truth of ‘Rain in Paris is good’ to the epistemic state of speakers is not motivated by the contention that there is no fact-of-the-matter whether rain in Paris is good or not; for given the shared non-instrumental values there *is* a fact of the matter whether rain in Paris is good or not (i.e. rain in Paris is good if it is raining in London; rain in Paris is not good if it isn’t raining in London). Thus one cannot motivate the specifically *epistemic* relativisation of the truth of ‘Rain in Paris is good’ on purely expressivist grounds.⁵

So: what to do? A place to start is to explore the fact that the two triviality results are related; for as they are related, they to some extent stand or fall together. If we find a strategy for avoiding one triviality result, one can expect this strategy to be helpful in avoiding the other triviality result. And there is a strategy for avoiding the first triviality result: treat conditionals as *gappy* propositions (McGee (1989); McDermott (1996); Bradley (1998)). On such an analysis – and this has independent linguistic motivation – the conditional $A \rightarrow B$ has the same truth value as B when A is true, and *lacks* truth value when A is false. If one takes the probability of a gappy proposition to be the probability that it is true *given that it has a truth-value*, we find that (when A is non-gappy and $P(A) \neq 0$)

$$P(A \rightarrow G) = P(G | A).$$

Consider what this would mean for our desirability tracking proposition A *is good*. We saw above that it will track the conditional probability $P(G | A)$, so we will have:

$$V(A) = P(A \rightarrow G).$$

⁵ There may be other grounds. Hájek’s and Petit’s proposal is similar in structure to accounts of the indicative conditional that relativize the proposition expressed by an indicative conditional to the epistemic state of the speaker (e.g. van Fraassen (1976)). But whereas the indicative conditional is generally thought to have a clear *epistemically* expressive flavour – so if they are to be speaker-relativized it is natural to relativize them to the epistemic state of the speaker – the introduction of epistemic relativity in evaluative judgments is more difficult to motivate directly. However, if (as is argued below) ‘Rain in Paris is good’ lacks truth-value when it isn’t raining in Paris, there is a sense in which there is a possibility that there might not be a fact of the matter – even on the presumption of shared fundamental values – whether Rain in Paris is good or not (namely, when it isn’t raining in Paris). This could give grounds for an epistemic relativisation of the truth of ‘Rain in Paris is good’. Whether this, in the end, is just a variant of the analysis to be presented below would require a deeper analysis than can be given here.

That is, the proposition whose probability tracks the desirability of *A* (in our present example) is: ‘If *A*, then the current situation is good’. So the expected value or desirability of *It is raining in Paris* is the probability of the conditional *If it is raining in Paris, the current situation is good*.

If the gappy analysis is correct this means that ‘Rain in Paris is good’ is gappy; ‘Rain in Paris is good’ is *true* if it is raining in Paris and the current situation is good, *false* if it is raining in Paris and the current situation is not good, and *lacks truth value* if it isn’t raining in Paris.

As is shown in the appendix, this analysis can be generalized to languages with arbitrarily fine-grained value distinctions (e.g. ‘the current situation is good to degree *r*’ for any real valued *r*) allowing that an agent’s value for *A* is massively entangled with judgements of the degree of goodness of the current situation. The analysis thus allows for a systematic way of avoiding both of Lewis’ triviality results by drawing on an approach to avoid Lewis’ first triviality result in order to avoid Lewis’ second triviality result.

An essential feature of this analysis is that propositions can be ‘gappy’. This is by no means an unproblematic move. Probability measures on a space of propositions that allow for gaps have different formal properties than standard probability measures (see Cantwell (2006) for a complete axiomatisation). They can however be given a rigorous interpretation in terms of betting quotients on conditional bets (a bet on a gappy proposition is won if the proposition is true, lost if the proposition is false, and canceled if the proposition lacks truth value) and so present a perfectly respectable interpretation of degrees of belief in gappy propositions in terms of betting dispositions (see also McGee (1989) for a more sophisticated betting interpretation). Furthermore, introducing gappy propositions does not incur any structural cost; for (see the appendix) the requisite gappy propositional structure is completely parasitic on its non-gappy core. Given a standard non-gappy structure of propositions and a probability measure and a desirability measure on this structure, the extension of these structures to the gappy case is uniquely determined by the base structures.

But does it even make sense for an expressivist to speak of the *proposition* expressed by ‘The current situation is good’? And why would we think that ‘*A* is good’ expresses a gappy proposition? These questions will be addressed in the remaining two sections. The concluding appendix treats the main claims in a more formally rigorous manner.

3. ‘The current situation is good’ – an Expressivist Interpretation

Consider a set of possible states and let the subsets of the set of states space denote propositions (‘gappy’ propositions have not yet entered the scene). For a given utility measure one can identify the set of possible states that

have a utility over a certain threshold. The threshold for when a possible state has a high enough utility to classify as 'good' will in many cases be vague, but let us ignore this complication and consider only situations where the differences in utility are discrete enough to warrant a clear subdivision of states into those that have a high-enough utility and those that don't. Let us say, provisionally, that 'The current situation is good' is true in the set of possible states that are above the threshold, and is false in the set of possible states that are below the threshold.

On the suggested truth-conditions the question whether the current situation is good does not depend on what the agent believes about the world; *the current situation is good* is an *epistemically objective* proposition. Thus the judgement *the current situation is good* is to be distinguished from the judgement *the current situation has a high expected utility*; one can be ready to make the latter judgement even though one is aware that the current situation may have a low utility. For instance, say that I find myself in a situation where I have a big chance of winning a substantial amount of money and a small risk of losing a moderate amount. I would take such a situation to have a high degree of expected utility, but there is an obvious sense in which I do not yet know whether the situation is good or bad, as I do not yet know whether I will win or lose: the current situation's *expected* utility is high, its *actual* utility is unknown. On the proposed analysis, the proposition *the current situation is good* pertains to actual, not expected, utility, and is in this sense objective. I can be rather confident that the current situation is good (by my own lights), yet be wrong.

Epistemic objectivity is consistent with a form of value subjectivity, for want of a better word. The truth-set of 'The current situation is good' is defined in relation to a utility function and so can be relativized to an agent. This is potentially problematic as it suggests that evaluative claims are either given an indexically subjectivist interpretation (e.g. 'The current situation is good' simply means 'The current situation is good *according to my values*'), or a relativist interpretation (e.g. the proposition expressed by 'The current situation is good' is *true for me* but perhaps not for others), with the consequent problems of accounting for what it is for two agents to genuinely disagree about evaluative claims. However, the expressivist can shun both subjectivism and relativism, at least if one takes these to involve commitment to the meaning-theoretical stance that meaning is determined by truth conditions.

A thorough analysis of the meaning theoretical issues faced by the expressivist cannot be given here. But a few brief comments can bring out the flavour of what an expressivist account could be. The expressivist takes the meaning of a sentence to be determined by the mental state that it is used to express. As 'The current situation is good' is used by speakers to express that the speaker takes the current situation to have a high enough utility

then *that* – the speakers appreciation of the current situation – is what determines the meaning of the sentence. So, for instance, the expressivist could hold (relative to the present framework) that the meaning of ‘The present situation is good’ is determined by the norm that an agent should hold the sentence to be acceptable to degree x iff x is the agent’s degree of belief that the world is in a state that is desirable. The semantically primitive notion here is *acceptability*, not *truth*. Such a norm of acceptance provides the mental state expressed by ‘The present situation is good’; it is a norm that neither entails nor excludes the possibility that evaluative judgements are objective, but invokes neither objective nor subjective notions of truth. As degree of acceptability is the primitive semantic notion, we have to understand agreement and disagreement by other means than by appeal to the truth value of the proposition about which speakers disagree. For instance, the expressivist can take fundamental disagreement about whether the current situation is good or not – disagreement that remains when all factual issues have been settled – to involve a conflict about what one *should* value. In asserting ‘The current situation is good’ one is recommending the audience to adopt the attitude thereby expressed; there is no need to invoke subjective or relative truths in the meaning theoretical picture.

Still, norms of use will determine an agent-relative ‘truth-set’ for ‘The present situation is good’, and this has a number of theoretically useful consequences: it allows one to speak of the (agent relative) *proposition* expressed by the sentence, and to speak of the agent as having some particular *degree of belief* in that proposition, but these are now defined in terms that do not violate the basic expressivist contention that normative evaluative claims express motivational attitudes and that it is this expressive quality that determines their meaning. In the spirit of Blackburn’s (e.g. Blackburn (1993)) quasi-realism, central concepts like *proposition*, *truth* and *belief* have been reshaped so as to fit the mold of expressivism.

One would, of course, need to tell a longer story here. For instance, there is the worry that the lack of an objectively determined truth value for evaluative claims will undermine the interpretation that gives meaning to the numerical assignments of degrees of belief and utility – the starting point of the current account. There is a strong tradition, following Ramsey, de Finetti, Savage, etc., that interprets ascriptions of subjective probability and value in behavioural terms, and on at least some versions of such interpretations, the notion of agent independent truth seemingly plays a role. So, for instance, one important interpretation of subjective probabilities is the *betting interpretation* – one’s subjective probability for A should be the betting quotient at which one is willing to accept bets on or against A . Now, if a bet is to be settled, the people who engage in the bet must eventually be able to agree on whether a bet has been won or lost. But if there is a sense in which one agent can legitimately hold that the current situation is good while at the same time another agent can legitimately deny that the

current situation is good, then there is no objective basis on which to settle a bet on whether the current situation is good or not. So it would seem that the betting interpretation breaks down when applied to propositions such as *the current situation is good*.

There is, I think, a straightforward retort to this particular worry. The betting interpretation is an operationalization of the concept of a degree of belief, based on the intuition that beliefs are action guiding. The operative term is how beliefs and desires *cohere*, not how they fit with an objective reality. An agent is coherent if and only if the set of betting quotients that the agent finds acceptable cannot be subjected to a *Dutch Book*. That is, if an agent is incoherent he or she can be offered a set of bets, that separately are acceptable for the agent, but that jointly will lead to a sure loss *by the agent's own lights*. What is important in this operationalization is that the settlement conditions for a bet mirrors the agent's disposition to act on that proposition when its truth or falsity is linked to something of value to the agent. There is no fundamental problem in holding that in order to measure the agent's degree of belief that, say, the present situation is good, a bet on whether the present situation is good or not should be settled relative to what the agent considers good or bad. Remember: the proposition *the current situation is good* is epistemically objective, the agent can have a high degree of belief that the current situation is good even though the situation, by the agent's own standards of goodness, is not good; so an agent can lose a bet on whether the current situation is good or not even by the agent's own standards of what is good and bad. This concession – to let the settlement conditions of bets be determined by the agent whose degrees of belief are to be measured – is for measurement purposes only; a 'real' bet on whether the present situation is good, a bet in which both agents can lay equal claim to being determinants of the settlement conditions, will presumably end in conflict or in the participants agreeing to call the bet off.

4. Edging Closer to an Analysis of 'A is good'?

Consider a straightforward evaluative sentence like 'The toss of the coin had a good outcome'. This sentence will be (speaker relatively) *true* in any possible state where the coin was tossed and the outcome was such that it had a (agent relatively) good value; it will be *false* in any possible state where the coin was tossed and the outcome was such that it had a (agent relatively) bad value. But what about the states where coin wasn't tossed? In these cases the definite descriptions 'The toss of the coin' and 'The outcome of the toss of the coin' denote nothing. On at least one influential account of the meaning of definite descriptions we in such cases we have *presupposition failure*; to say that the toss of the coin had a good outcome is to *presuppose*, not to *say*, that the coin was tossed and had an outcome. When there is presupposition

failure (on at least one dominant analysis) the sentence lacks truth value. So 'The toss of the coin had a good outcome' turns out to be true if the coin was tossed and the outcome was good, false if the coin was tossed and outcome was bad, and to lack truth value if the coin wasn't tossed. That is, it is equivalent to 'If the coin was tossed, then the outcome of the toss was good' when the conditional is given a gappy interpretation.

Propositions, being abstract entities, exist in all possible states; however, being abstract objects propositions are not the kind of things that are good or bad. We can attribute goodness and badness to actions, events and states of affairs – things that we take to be part of the makeup of our world – not to abstract entities. I would suggest that a sentence like 'Rain in Paris is good' is best interpreted as attributing goodness to the state of affairs that obtains when it is raining in Paris, it doesn't predicate goodness to the *proposition* that it is raining in Paris. That is, the formal (ungrammatical) construction 'It is raining in Paris is good' or (the grammatical) 'That it is raining in Paris is good' should be interpreted as (the grammatical) 'The state of affairs that it is raining in Paris is [a] good [state of affairs]'. The state of affairs that obtains if and only if a proposition is true will obviously not obtain when the proposition is false. So if it isn't raining in Paris, the construction 'The state of affairs that it is raining in Paris' fails to denote a state of affairs: we have presupposition failure. This explains why 'Rain in Paris is good' turns out to be equivalent to 'If it is raining in Paris, then this [the state of affairs that obtains when it is raining in Paris] is good' under the gappy interpretation of the conditional. Presupposition failure thus explains both the truth-gappiness of 'Rain in Paris is good' and its connection to the conditional construction.

There is a fundamental problem remaining however. To judge that rain in Paris is good is naturally seen as making a *partial* assessment of the situation as a whole, but in the present paper 'Rain in Paris is good' has in effect been paraphrased as 'If it is raining in Paris, then the current situation is good *as a whole*'. This paraphrase is apt only in contexts where if anything is good, the situation as a whole is good, that is, in contexts in which there is only one relevant dimension of value (such as in a one-bet betting situation where one only values the money lost or gained).

One can give a reasonable interpretation of '*A* is good' that respects the fact that it is a partial assessment (this is given a more formal treatment in the appendix): *A* is good if *A* is true and the state of the world would have been worse (less desirable) had *A* been false.⁶ In terms of assertability conditions:

⁶ According to this analysis one can hold that *A* is good even if the world would have been only a tiny tiny bit worse had *A* been false. A different analysis would be required if one wants to capture the idea that the state of affairs that *A* contributes some absolute good.

'*A* is good' is acceptable to degree *x* iff *x* is the agent's degree of belief that [*A* is true and the state of the world would have been less desirable (by the agent's own lights) had *A* not been true].

Or, in terms of agent relative truth conditions (see the appendix for a more rigorous analysis):

'*A* is good' is true (agent relatively) iff *A* is true and the state of the world would have been less desirable (by the agent's own lights) had *A* not been true.

'*A* is good' is false (agent relatively) iff *A* is true and the state of the world would have been no less desirable (by the agent's own lights) had *A* not been true.

Clearly, under this analysis it will not in general hold that '*A* is good' is equivalent to 'If *A*, then the present situation is good as a whole' (e.g. it can happen that the present situation is bad, but would have been worse if *A* had not been true, and so *A* is good even though the situation as a whole is bad). However, assuming that whether or not it rains in Paris is causally independent of whether or not it rains in London (which, clearly, need not in general be the case), i.e. assuming that it would have rained in London even if it hadn't rained in Paris, it follows that *Rain in Paris is good* is true (by Jane's values) iff it is raining in both Paris and London (and so Jane wins the bet), and is false iff it is raining in Paris but not in London (and so Jane loses the bet). That is, in the specific circumstances of our original example, '*Rain in Paris is good*' is true (false) iff 'If it is raining in Paris, then the present situation is good as a whole' is true (false). So we will, in our original example, have the stable identity $V(A) = P(A \text{ is good})$.

As argued, if it doesn't rain in Paris it makes no sense to say that the state of affairs that it is raining in Paris is good. However, if it rains in London it is clear that rain in Paris *would be* good, regardless of whether it is actually raining in Paris or not. Note the switch here from the indicative '*is good*' to the subjunctive '*would be good*'. The subjunctive '*Rain in Paris would be good*' does not presuppose that it raining in Paris; there is no presupposition failure and so no truth gap.

'*A would be good*' is true (agent relatively) iff the state of the world would be more desirable⁷ were *A* to be true than were *A* to be false.

'*A would be good*' is false (agent relatively) iff the state of the world would be no more desirable were *A* to be true than were *A* to be false.

⁷ Perhaps 'more desirable' should be replaced by '*significantly* more desirable' or some construction involving a threshold of value to distinguish *A would be good* from *A would be better*.

Just as there is a link between the indicative ‘Rain in Paris is good’ and the indicative ‘If it is raining in Paris, the present situation is good’, one can conjecture that there is a corresponding link between the subjunctive ‘Rain in Paris would be good’ and the subjunctive conditional ‘If it were to rain in Paris, the present situation would be good’. Just as the probabilities of the two indicatives should go hand in hand, the probabilities of the two subjunctives should go hand in hand; the subjunctives, however, would not track the desirability ($= V(A)$ = the *evidential* expected utility) of rain in Paris, rather it would track the *causal* expected utility of rain in Paris (e.g. see Byrne and Hájek (1997)).⁸

In brief. One can show – *contra* Lewis – that it is possible to have a systematic correlation between the desirability of A and degrees of belief towards evaluative propositions of the form *If A , then the current situation is good to degree r* . Furthermore, while we do not in general have $V(A) = P(A \text{ is good})$ on a natural interpretation of ‘ A is good’, there are non-trivial situations in which – *contra* Lewis – the identity holds and is stable under conditionalization. Lewis’ triviality result is a potential threat to the idea that there is a deep entanglement between desires and evaluative judgments – a potential threat that should worry the expressivist – but this threat, I have argued, can be avoided in a way amenable to the expressivist.

Appendix: A Formal Model

Let U be a set of states $\{u_1, u_2, \dots\}$; let p (a probability mass on U) be a real valued function on U such that $\sum_{u \in U} p(u) = 1$, and let μ (a utility function on U) be a real valued function on U . The probability mass p is taken to be a representation of the agent’s epistemic state, while the value function is taken to be a representation of how the agent values basic the

⁸ I take the objections of Weintraub (2007) towards taking expected utility (desirability) as a measure of goodness to be, aptly, directed at the difficulty of construing a holistic measure such as expected utility in terms of how the goodness of any given state of affairs contributes to the goodness of the situation as a whole. Daskal (2010) elaborates on this idea and presents a measure of the ‘absolute’ value of propositions or states of affairs (roughly: the value that each proposition or state of affairs contributes to the whole situation). This analysis pertains to the semantics of ‘good’ but does not really address the triviality result that besets the desire-as-belief thesis: Lewis quite deliberately did not commit himself to holding that it should be ‘ A is good’ (on a natural interpretation of such a sentence) that tracks the desirability of A , any evaluative proposition would do. In any case Daskal’s measure of the absolute value of propositions is probably best understood as the basis for an analysis of the subjunctive *A would be good* rather than *A is good* as Daskal factors the ‘contribution’ of A ’s goodness when A is false into the analysis. This, of course, does not make the analysis any less interesting, indeed the present analysis of ‘ A is good’ borrows an element of counterfactuality from Daskal’s analysis, but only for situations in which A is true.

basic bundles of goods associated with a possible state. This is a standard ‘non-gappy’ representation of degrees of belief and desirability on an outcome space.

Let $T(A)$ denote the *truth set* of the sentence A (so $T(A) \subseteq U$) and $F(A)$ denote the *falsity set* of A (so $F(A) \subseteq U$). Together the pair $(T(A), F(A))$ can be said to constitute the *proposition* expressed by A . We want to allow for the possibility that there may be ‘gappy’ propositions, so there is no general requirement that $T(A) \cup F(A) = U$, but it will be assumed that the truth set and the falsity set do not overlap: $T(A) \cap F(A) = \emptyset$.

The probability of A will be defined as the probability that the proposition expressed by A is true divided by the probability that it has a truth-value:

$$P(A) = \frac{\sum_{u \in T(A)} p(u)}{\sum_{u \in T(A) \cup F(A)} p(u)}.$$

The *desirability* of A is defined in a standard manner:

$$V(A) = \frac{\sum_{u \in T(A)} p(u) \times \mu(u)}{\sum_{u \in T(A)} p(u)}.$$

Assume that the language contains the conditional \rightarrow and the conjunction $\&$, with the following truth and falsity sets:

$$\begin{aligned} T(A \rightarrow B) &= T(B) - F(A), & F(A \rightarrow B) &= F(B) - F(A). \\ T(A \& B) &= T(A) \cap T(B), & F(A \& B) &= F(A) \cup F(B). \end{aligned}$$

THEOREM 1

For any non-gappy A and B such that $P(A) > 0$:⁹

$$P(A \rightarrow B) = \frac{P(A \& B)}{P(A)}$$

⁹ When non-gappy propositions are involved the general definition of the conditional probability would be:

$$P(B|A) = \frac{\sum_{u \in T(B) - F(A)} p(u)}{\sum_{u \in (T(B) \cup F(B)) - (F(A) - F(A))} p(u)}.$$

In the non-gappy cases this simplifies to $P(B|A) = \frac{\sum_{u \in T(B) \cap T(A)} p(u)}{\sum_{u \in T(A)} p(u)} = P(A \& B)/P(A)$.

Proof:

$$\begin{aligned}
 P(A \rightarrow B) &= \frac{\sum_{u \in T(A \rightarrow B)} p(u)}{\sum_{u \in T(A \rightarrow B) \cup F(A \rightarrow B)} p(u)}, \\
 &= \frac{\sum_{u \in T(B) - F(A)} p(u)}{\sum_{u \in (T(B) - F(A)) \cup (F(B) - F(A))} p(u)}, \\
 &= \frac{\sum_{u \in T(B) \cap T(A)} p(u)}{\sum_{u \in (T(B) \cap T(A)) \cup (F(B) \cap T(A))} p(u)}, \\
 &= \frac{\sum_{u \in T(B) \cap T(A)} p(u)}{\sum_{u \in T(A)} p(u)}, \\
 &= \frac{P(A \& B)}{P(A)}.
 \end{aligned}$$

□

The identity holds for all probability measures, so is stable under conditionalization.

Assume that for each real number r , there is a sentence G^r such that $T(G^r) = \{u \in U \mid \mu(u) = r\}$ and $F(G^r) = \{u \in U \mid \mu(u) \neq r\}$. G^r can be read “The current situation is good to degree r ”.

THEOREM 2

For any non-gappy A such that $P(A) > 0$:

$$V(A) = \sum_{r \in \mathbb{R}} P(A \rightarrow G^r) \times r.$$

Proof: Assume that $P(A) > 0$. We then have:

$$\begin{aligned}
 \sum_{r \in \mathbb{R}} P(A \rightarrow G^r) \times r &= \sum_{r \in \mathbb{R}} \frac{P(A \& G^r)}{P(A)} \times r, \\
 &= \sum_{r \in \mathbb{R}} \frac{\sum_{u \in T(A) \cap T(G^r)} p(u)}{\sum_{u \in T(A)} p(u)} \times r, \\
 &= \frac{\sum_{u \in T(A)} p(u) \mu(u)}{\sum_{u \in T(A)} p(u)}, \\
 &= V(A).
 \end{aligned}$$

□

So, in particular: if $\mu(u) \in \{0, 1\}$ for all $u \in U$, then $V(A) = P(A \rightarrow G^1)$.

The theorem also shows what happens if we allow for degrees of goodness. If, say, $\mu(u) \in \{0, 0.5, 1\}$ for all $u \in U$ then we have $V(A) = P(A \rightarrow G^1) + P(A \rightarrow G^{0.5}) \times 0.5$.

For the analysis of ‘ A is good’ we need to add an element that allows for counterfactual comparisons in the model. Let \circ be an operator that for each state u and set of states X picks out the state in X that is most similar to u (when X is empty $u \circ X$ is undefined). For instance, $u \circ F(A)$ takes one to the state closest to u in which A is false.

Define (for A such that $F(A) \neq \emptyset$, the case when A is a tautology will be left open):

$$T(A \text{ is good}) = \{u \in T(A) \mid \mu(u) > \mu(u \circ F(A))\},$$

$$F(A \text{ is good}) = \{u \in T(A) \mid \mu(u) \leq \mu(u \circ F(A))\}.$$

A sentence A is *outcome essential* if $\mu(u) \neq \mu(u \circ F(A))$ whenever $u \in T(A)$ (that is, A is outcome essential if things would have been either better or worse had A been false).

OBSERVATION 1

For any non-gappy outcome essential A such that $F(A) \neq \emptyset$, if $\mu(u) \in \{0, 1\}$ for each $u \in U$ and $P(A) > 0$ then:

$$V(A) = P(A \text{ is good}).$$

Proof: Note that as A is outcome essential: $u \in T(A \text{ is good})$ iff $u \in T(A)$ and $\mu(u) = 1$. Note also that $A \text{ is good}$ has a truth-value iff A is true. So:

$$P(A \text{ is good}) = \frac{\sum_{u \in T(A)} (p(u) \times \mu(u))}{\sum_{u \in T(A)} p(u)} = V(A). \quad \square$$

The restriction to situations where there are only two levels of value (e.g. *good* and *bad*) and to outcome essential propositions is important. If there are more than two levels of value, or if a proposition can hold in a good state without contributing to the goodness of that state, the identity $V(A) = P(A \text{ is good})$ will not in general obtain. These are strong restrictions, but they are satisfied in our original example.

Let:

$$U = \{u_{PL}, u_{P \sim L}, u_{\sim PL}, u_{\sim P \sim L}\},$$

$$\mu(u_{PL}) = \mu(u_{\sim P \sim L}) = 1,$$

$$\mu(u_{P \sim L}) = \mu(u_{\sim PL}) = 0,$$

(in u_{PL} it is raining in both Paris and London, in $u_{P\sim L}$ it is raining in Paris but not in London, and so on). If $u_{PL} \circ F(\text{It is raining in Paris}) = u_{\sim PL}$ and $u_{P\sim L} \circ F(\text{It is raining in Paris}) = u_{\sim P\sim L}$, then ‘It is raining in Paris’ is outcome essential and so, as long as $P(\text{It is raining in Paris}) > 0$, $V(\text{It is raining in Paris}) = P(\text{Rain in Paris is good})$.

We can obtain a more general result if we allow sentences of the form ‘[the state of affairs that] A contributes r [utils]’, with truth-conditions:

$$\begin{aligned} T(A \text{ contributes } r \text{ utils}) &= \{u \in T(A) \mid \mu(u) - \mu(u \circ F(A)) = r\}, \\ F(A \text{ contributes } r \text{ utils}) &= \{u \in T(A) \mid \mu(u) - \mu(u \circ F(A)) \neq r\}. \end{aligned}$$

That is, A contributes r utils if A is true and the difference in value had A been false is r .

A sentence A is *strongly outcome essential* if $\mu(u \circ F(A)) = \min_{v \in U} \mu(v)$ whenever $u \in T(A)$ and $\mu(u) > \mu(u \circ F(A))$ (that is, A is strongly outcome essential if whenever A is true and A makes some positive contribution in utility, then A makes the whole contribution in utility).

THEOREM 3

For any non-gappy outcome essential A such that $F(A) \neq \emptyset$, and $P(A) > 0$, if $\min_{v \in U} \mu(v) = 0$, then:

$$V(A) = \sum_{r \in \mathbb{R}} P(A \text{ contributes } r \text{ utils}) \times r.$$

Proof: Note that as A is strongly outcome essential: $u \in T(A \text{ contributes } r \text{ utils})$ iff $u \in T(A)$ and $\mu(u) = r$. Note also that A contributes r utils has a truth-value iff A is true. So:

$$\begin{aligned} \sum_{r \in \mathbb{R}} P(A \text{ contributes } r \text{ utils}) \times r &= \sum_{r \in \mathbb{R}} r \times \frac{\sum_{u \in T(A) \& \mu(u)=r} p(u)}{\sum_{u \in T(A) \& \mu(u)=r} p(u)} \\ &= \frac{\sum_{u \in T(A)} p(u) \times \mu(u)}{\sum_{u \in T(A)} p(u)} \\ &= V(A). \end{aligned}$$

□

The restriction to strongly outcome essential propositions is important. If the state of affairs that A only contributes some of the value to a complete state then the expectation value of A can differ from the expected contribution of the state of affairs that A .

Consider, finally, the extension to the original example where we factor in the possibility of rain in Rome (R). We then get eight possibilities:

$$U = \{u_{RPL}, u_{RP\sim L}, u_{R\sim PL}, u_{R\sim P\sim L}, u_{\sim RPL}, u_{\sim RP\sim L}, u_{\sim R\sim PL}, u_{\sim R\sim P\sim L}\}.$$

From the description of the situation we have:

$$\begin{aligned}\mu(u_{\sim RPL}) &= \mu(u_{\sim R\sim P\sim L}) = 1, \\ \mu(u_{RPL}) &= \mu(u_{RP\sim L}) = 0.5, \\ \mu(u_{RP\sim L}) &= \mu(u_{R\sim PL}) = \mu(u_{\sim RP\sim L}) = \mu(u_{\sim R\sim PL}) = 0.\end{aligned}$$

Assuming that the three cities are weatherwise causally independent we find that the proposition that it is raining in Paris is strongly outcome essential and so, as long as it has a probability greater than 0:

$$\begin{aligned}V(\text{It is raining in Paris}) &= P(\text{Rain in Paris contributes 1 utile}) \\ &\quad + P(\text{Rain in Paris contributes 0.5 utiles}) \times .5.\end{aligned}$$

John CANTWELL

References

- [1] ADAMS, E. W. (1975). *The Logic of Conditionals*. Dordrecht: Reidel.
- [2] BLACKBURN, S. (1993). *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- [3] BRADLEY, R. (1998). A Representation Theorem for a Decision Theory with Conditionals. *Synthese* 116, 187–229.
- [4] BRADLEY, R. and C. LIST (2009). Desire-as-belief revisited. *Analysis* 69(1), 31–37.
- [5] BYRNE, D. and A. HÁJEK (1997). David, Hume, David Lewis, and Decision Theory. *Mind* 106(423), 411–428.
- [6] CANTWELL, J. (2006). The laws of Non-Bivalent Probability. *Logic and Logical Philosophy* 15, 163–171.
- [7] DASKAL, S. (2010). Absolute value as belief. *Philosophical Studies* 148(2), 221–229.
- [8] EDGINGTON, D. (1995). On Conditionals. *Mind* 104, 235–329.
- [9] EDGINGTON, D. (2000). General conditional statements: A reply to Kölbel. *Mind* 109, 109–116.
- [10] HÁJEK, A. and P. PETTIT (2004). Desire beyond belief. *Australasian Journal of Philosophy* 82(1), 77–92.
- [11] KANGER, S. (1971). New Foundations for Ethical Theory. In R. Hilpinen (Ed.), *Deontic Logic: Introductory and Systematic Readings*, pp. 36–58. Dordrecht: Reidel.
- [12] KÖLBEL, M. (2000). Edgington on compounds of conditionals. *Mind* 109, 97–108.

- [13] LEWIS, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *Philosophical Review* 85, 297–315.
- [14] LEWIS, D. (1988). Desire as Belief. *Mind* 97(387).
- [15] LEWIS, D. (1996). Desire as Belief II. *Mind* 105(418).
- [16] McDERMOTT, M. (1996). On the truth conditions of certain ‘If’-sentences. *Philosophical Review* 105, 1–37.
- [17] MCGEE, V. (1989). Conditional probabilities and compounds of conditionals. *Philosophical Review* 98, 485–542.
- [18] PRICE, H. (1989). Defending Desire as Belief. *Mind* 98(389), 119–127.
- [19] VAN FRAASSEN, B. (1976). Probabilities of Conditionals. In W. L. Harper and C. A. Hooker (Eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, pp. 261–300. D. Reidel.
- [20] WEINTRAUB, R. (2007). Desire as Belief, Lewis Notwithstanding. *Analysis* 67(2), 116–122.