

Supposition and Choice: Why ‘Causal Decision Theory’ is a Misnomer

John Collins
Columbia University

February 7th, 1996

This paper has as its topic two recent philosophical disputes. One of these disputes is internal to the project known as decision theory, and while by now familiar to many, may well seem to be of pressing concern only to specialists. It has been carried on over the last twenty years or so, but by now the two opposing camps are pretty well entrenched in their respective positions, and the situation appears to many observers (as well as to some of the parties involved) to have reached a sort of stalemate. The second of these two disputes is, on the other hand, very much alive. While it has been framed in decision theoretic terms, it is definitely not a dispute internal to that enterprise. It is, rather, a debate about the very coherence of the notion of objective value, and as such touches on issues of central importance to, for example, meta-ethics and moral psychology.

I think that this second debate deserves more attention from the general philosophical public than it has so far received. That is one of the motivating reasons for this paper. If this debate has received less attention than it deserves, that may be because it centers around a certain rather technical looking argument presented within a formal framework, on a topic that has not traditionally been approached in that way. In fact I see it as a striking example of the right way to use formal methods in philosophical argumentation: to clarify rather than to obscure, to present a particular argument more succinctly and more precisely than could be done informally, and to do so in a way that lays bare the presuppositions and premises of the argument for everyone to see.

In what follows I shall sketch and explain the argument and examine those assumptions. I am going to suggest that the right response to the argument is not to accept its conclusion, but rather to reject one of those presuppositions. The interesting result then will be that the dropping of that assumption allows us to see the first (apparently stalemated) debate in a new and revealing light. That is the second of my aims in writing this piece.

1 ‘Evidential’ and ‘Causal’ Decision Theory

The first of my two disputes—the one internal to decision theory, the apparently stalemated one—is the debate between the proponents of the so-called ‘evidential’ and ‘causal’ versions of decision theory. This debate, you’re probably familiar with it, usually focusses on a certain wacky, science-fiction-like story about a remarkably accurate predictor of human choice and the million dollars that she may or may not have placed yesterday in an opaque box there on the table in front of you. The story in question strikes some as a reliable intuition pump, and others (it’s hard to improve here on Dick Jeffrey’s memorable phrase!) as no more than a “Prisoner’s Dilemma for space cadets”. Whatever the case may be, I’m not going to discuss Newcomb’s Problem here. That there’s no real need to do so is part of the point of this paper.

But let me at least describe what is at stake. Decision theorists (all of them) agree that desirability, i.e. subjective value, is to be explicated as expected value. They agree that insofar as decision theory is a normative enterprise, it prescribes that one choose so as to maximize expected value. Furthermore, they agree that the notion of expected value involves, at least implicitly, the notion of revision of belief. The expected value of an option A is the probability-weighted average of the ways in which A might be true. The probabilities involved are just those that result when the agent revises to accommodate the proposition that A is the choice made. In other words, to calculate the expected value of A , the agent should revise to accommodate A , and see how good things would be then. We can write down the following definition of expected value that is neutral between the two versions of the theory:

Expected value of A (Neutral version):

$$V(A) = \sum_{w \in A} v(w) \cdot P_A(w)$$

The sum here is over all the A -worlds, P_A is the probability function that results when the agent's subjective probability function P is revised so as to accept A , and v is the agent's *value measure*, i.e her subjective value function restricted to point propositions. (For the sake of simplicity I am abusing notation by not distinguishing worlds from point propositions, and assuming here that there are only finitely many of them).

Some further terminology will be useful later. We can see from the above definition that an expected value function can be thought of determined by two factors, the probability function P and the value measure v . It follows that a change in the agent's expected values can come about in one of two different ways. Following Jeffrey I shall call two different sorts of value change: *doxastic changes* and *valuational changes* in expected value. A doxastic change in expected value is the result of a change in belief, i.e. a change in the probability function P . A valuational change in expected value results from a change in the underlying value measure v .

Note that this neutral definition involves the revised probability function P_A , recall: P revised to accommodate the truth of A . This is what I meant above when I said that the notion of expected value involves at least implicitly the notion of revision of belief. The dispute between the two versions of decision theory amounts to no more than a disagreement about what sort of revision is appropriate here. There are two rival suggestions about revision method, hence two candidates for the role of expected value, and hence two versions of decision theory. If revision in the above neutral definition is taken to be revision by conditionalization, then we get evidential decision theory:

Evidential expected value:

$$EEV(A) = \sum_{w \in A} v(w) \cdot P(w/A)$$

If the revision in question goes by an alternative method known as “imaging,” then we get the causal theory:

Causal expected value:

$$CEV(A) = \sum_{w \in A} v(w).P\#A(w)$$

Conditionalization requires no explanation here. Imaging is a less familiar process. To get a better feeling for the way imaging works, and for the way it differs from the method of conditionalization, it will help to think of both processes in terms of what Bas van Fraassen has called “Muddy Venn Diagrams.”¹ Think of logical space as consisting of all the points on some two-dimensional plane surface, like a table-top. Propositions and belief states, being sets of worlds, can be thought of as regions of this plane surface. We might indicate the belief state of a particular agent at a particular time by spreading black silty mud over the appropriate part of the table. The mud covers all and only the worlds that are live epistemic possibilities for the agent. The thickness of the mud at a particular point represents the agent’s subjective probability for that state of affairs. The probability of a proposition is then given by finding the volume of mud covering the region corresponding to that proposition and representing this as a fraction of the total amount of mud. What does belief revision look like on this picture of things? Revising the agent’s belief state to accommodate some proposition A , amounts to rearranging the mud on the table in some way so that after the change, the only mud that is left is located on the region of the table that represents the proposition A . Here are two physical processes that would do the job:

1. Wipe away all the mud that lies outside the region A , and leave the rest of the mud where it is.
2. Leave any mud that is already in the region A where it is. Then for each point w outside A , slide any mud that is lying on w over to the closest point (or points) to w within the region A .

The first of these processes amounts to conditionalization; the second is analogous to imaging.

Less picturesquely: consider each $\neg A$ -world to which the function P assigns non-zero probability. For each such world w find the “closest” or

¹See Bas van Fraassen [?] p.161.

“most similar” A -world to w , call it ‘ $w\#A$ ’, and map the probability P assigned to w over to $w\#A$. The sense of “closest” or “most similar” here is intended to be that familiar from the Lewis–Stalnaker semantics for the subjunctive conditional. Thought of this way, it is not hard to see that $P\#A(w)$ is none other than the prior probability $P(A\Box\rightarrow w)$ assigned by the agent to the subjunctive conditional $A\Box\rightarrow w$. Since such conditionals are widely regarded to be relevant to a proper analysis of the causal relation, this definition has come to be thought of as yielding a notion of expected value sensitive to the causal structure of the situation in question. Hence the name “causal decision theory.” That’s (one version of) the standard account, though as I suggested earlier, I think it is the wrong way to look at things. I’ll describe what I take to be the right way later.

2 Objective Value and Desire-as-Belief

The second of the two disputes on which this paper centers concerns the possibility of giving a coherent cognitivist or anti-Humean account of motivation. Our ordinary explanations of behavior make reference to two kinds of mental state, which we call belief and desire. It is worth asking whether these kinds are necessarily distinct, or whether it might be possible to construe desire as belief of a certain sort—belief, say, about what would be good. Decision theory formalizes our notions of belief and desire, treating each as a matter of degree. In this context the thesis that desire is belief might amount to the claim that the degree to which an agent desires any proposition equals the degree to which the agent believes that it would be good if that proposition were true.

In [?] David Lewis presented an argument against this anti-Humean proposal that desire is belief. Lewis proved that, on pain of triviality, the Desire-as-Belief Thesis cannot be added to the axioms of decision theory.

If the Desire-as-Belief Thesis were true, it would be possible to do away with all reference to desire in ordinary explanations of behavior. Talk of desire could systematically give way to talk about belief in objective goodness. Two aspects of this are worth noting:

1. the proposal would reduce an apparently non-cognitive attitude (desire) to one that is clearly cognitive (belief);

2. the cognitive attitude offered is an attitude towards something objective (goodness), whereas the original attitude (the desire) was something subjective.

Seen in this light, the Desire-as-Belief Thesis appears as but one of a whole slew of similar proposals sharing with it one or both of those features. Here are some examples:

- (a) Credence as belief about objective chance.
- (b) Epistemic possibility as belief about objective possibility.
- (c) Conditional belief as belief about objective conditionals.
- (d) Desire as belief about objective goodness.
- (e) Hope as belief about objective hopefulness.
- (f) Terror as belief about objective ghastliness.

And so on. Readers may extend this list for themselves *ad libitum*. Some of these proposals are pretty clear non-starters. Take, for example, (a), which might be formulated as follows, using P for subjective probability (or credence), and Ch for objective chance. We suppress, for the sake of simplicity, reference to the particular agent and time.

Credence-as-Belief: $P(A) = x$ iff $P(Ch(A) = x) = 1$

This thesis falls to the “Integrating Out Argument”² Suppose that the agent’s credence is divided between various chance hypotheses. For simplicity, let’s assume it to be divided evenly between just two hypotheses: the hypothesis that the chance of A is 0.9, and a contrary hypothesis that the chance of A is 0.1. But then:

$$\begin{aligned}
 P(A) &= 0.5 \times P(A/Ch(A) = 0.9) + 0.5 \times P(A/Ch(A) = 0.1) \\
 &= (0.5 \times 0.9) + (0.5 \times 0.1) \\
 &= 0.5
 \end{aligned}$$

²See Savage [?] p.58, where this strategy is used against the suggestion that there might be higher-order subjective probabilities. Savage attributes the argument to Max Woodbury.

where we have appealed to the plausible principle:

$$P(A/Ch(A) = x) = x$$

By the Credence-as-Belief Thesis, this entails that $P(Ch(A) = 0.5) = 1$, which contradicts our initial assumption that the agent's subjective probability is divided between two distinct chance hypotheses. The Credence-as-Belief Thesis cannot be maintained, as it implies complete opinionation about chance. Exactly the same reasoning prevails against thesis (b).

Now in fact the integrating out argument can also be used against a simple formulation of the thesis that desire is belief, according to which the agent desires A to degree x just in case the agent fully believes that A 's degree of objective goodness is x . We suppose that the function V gives the agent's degrees of desire, and that there is a function u that assigns to each proposition A its degree of objective goodness. As before, the agent is taken to have full belief in X iff X receives subjective probability 1.

Desire-as-Full-Belief: $V(A) = x$ iff $P(u(A) = x) = 1$

The *reductio* proceeds, as before, from the assumption that the agent entertains two rival hypotheses about the goodness of some proposition. We see that this first version of the DAB Thesis entails that any rational agent must be completely opinionated about matters of objective goodness. That is an absurd conclusion.

But the cognitivist may object that this first version of the Desire-as-Belief Thesis, and corresponding first versions of our other theses, are misformulations. Belief, as well as goodness, admits of degree. Two dimensions of degree allow a three-way ambiguity in sentences like:

To what extent do you believe that a tightening of the money supply would adversely affect the stock market?

On the first interpretation, the question takes for granted full belief in the proposition that the effect of monetary tightening on the markets would be adverse, and is asking you to judge the extent of the damage. On the second reading, the question is asking you to judge your degree of belief in the proposition that stock prices would fall if the Fed cut back reserves. On the third reading, the word 'extent' serves to invoke simultaneously consideration

of degree of belief and of the magnitudes of the possible shifts in the market index. In this third case, an appropriate answer to the question might involve something like an expectation: a probability-weighted average of the magnitudes of the changes deemed possible.

There is a corresponding range of possible formulations of each of the above theses. Our first version of the Desire-as-Belief Thesis, in which only goodness is taken to be a matter of degree, corresponds to the first reading of the question. Corresponding to the second interpretation of the question is a second version of DAB, in which belief rather than objective goodness is allowed to admit of degree.

$$V(A) = x \quad \text{iff} \quad P(u(A) = 1) = x$$

Since we are now taking objective goodness to be an all or nothing matter, we may suppose for the purposes of reformulating the second version of the DAB Thesis that to each proposition A there corresponds a proposition A° (pronounced “ A -halo”: the proposition that A is objectively good. A° may be defined in terms of the objective value measure u as follows:

$$A^\circ \text{ is true at } w \quad \text{iff} \quad u(w) = 1$$

and the second version of the DAB Thesis simplifies to:

Desire-as-Degree-of-Belief: $V(A) = P(A^\circ)$

The third version of the DAB Thesis allows that both belief and objective goodness come in degrees. Degree of desire is then equated with expectation of goodness; it is represented as a probability weighted average of all the possible degrees of objective goodness. Call this the :

Desire-as-Expectation Thesis: $V(A) = \sum_x x.P(u(A) = x)$

The DAE Thesis appears to be the most general of the three versions we have considered, but it turns out our second version is already general enough. By this I mean that there is really nothing to be gained by moving from the second version to the third. In this paper I will follow David Lewis in taking the second Desire-as-Degree-of-Belief version of the Thesis to be our official statement of DAB. There will be no loss of generality since the

discussion could be straightforwardly (though tediously) adapted to cover the Desire-as-Expectation formulation instead. Given that that is the case, the discussion will be much simpler if we focus on the second version rather than the third.

In his 1988 paper Lewis presented an argument against this anti-Humean proposal. Lewis proved that the DAB thesis cannot be added to the axioms of decision theory without trivializing the theory. But the proof Lewis gave in the original paper was unnecessarily complicated. Two recent papers have clarified matters considerably. One is a paper I co-authored with Horacio Arló Costa and Isaac Levi entitled ‘Desire-as-Belief Implies Opinionation or Indifference’.³ In that paper we presented a considerably simpler proof of Lewis’s result. Lewis has since responded with an even more streamlined version of the argument.⁴ Here I will just give a quick sketch of Lewis’s latest version of the argument.

The first step is to see that DAB is equivalent to the conjunction of the following two principles

Desire-as-Conditional-Belief: $V(A) = P(A^\circ/A)$

Independence: $P(A^\circ/A) = P(A^\circ)$


Proof: Since DAB holds throughout a set of $\langle P, V \rangle$ closed under conditionalization, we may conditionalize by A on both sides of the equation. This yields DACB. IND then follows directly from DACB and DAB. The implication in the other direction is immediate.

Desire-as-Belief Theorem: Desire-as-belief implies opinionation or indifference.

Proof: If the agent is neither opinionated nor indifferent, we may choose a proposition A such that the agent’s probability function P assigns positive probability to the three regions marked with a ‘*’ on the diagram below.

³Arló Costa, Collins, and Levi [?].

⁴See Lewis [?].



Now if IND holds for P , it will fail after conditionalization on $A \vee A^\circ$, since this updating raises the probability of A° , but not the conditional probability $P(A^\circ/A)$. QED.

This is a surprising result. How should we respond to it? What we might call the “straight response” is to take the DAB Theorem as establishing that the simplest version of a rationalist, anti-Humean account of motivation is actually incoherent. I think that we may reasonably doubt whether this is really something that could be demonstrated *a priori*.

Christine Korsgaard has drawn a distinction between two kinds of skepticism about the rationalist position.⁵ A weaker non-cognitivist view, which Korsgaard calls *content skepticism*, finds fault with particular, specific principles that have been advanced as candidate norms of practical reason. A more ambitious Humean strategy aims at establishing what Korsgaard calls *motivational skepticism*. According to motivational skepticism, we can see *a priori* that there can be no such thing as pure practical reason, even before coming to consider any specific rationalist proposals. The straight response would be to take the Desire-as-Belief Theorem as an *a priori* argument for the Humean theory, as an argument for the Humean position based only on a consideration of the logic of desire and belief, or, in other words, as a conclusive argument for motivational skepticism.

⁵Korsgaard [?].

Korsgaard rejects motivational skepticism, as, presumably, do all those who propose specific norms for the operation of pure practical reason.⁶ I share their doubts about the possibility of this kind of strongly *a priori* proof of the Humean theory of motivation. But that is just to say that I believe we have reason to resist what I have been calling the ‘straight response.’

I shall proceed in the next section to investigate what it would take to deny the conclusion of the Desire-as-Belief Theorem. In other words, let us assume that cognitivism is not a position that can be ruled out on purely on the basis of the logic of belief and desire. Let’s hold the Desire-as-Belief Thesis fixed for the moment and see if we can reconstrue Lewis’s argument as a reduction against one of that argument’s premises or presuppositions.

Call a probability revision method ‘#’ *linear* if it satisfies the following condition:

Linearity: If $P = a.P' + (1-a).P''$, then $P\#A = a.(P'\#A) + (1-a).(P''\#A)$.

According to this linearity condition the operations of revising a probability function and taking a mixture of probability functions commute. In other words, the linearity condition is saying that the revision of a mixture of two probability functions is the mixture of the revisions.

In the finite case this condition may be equivalently expressed as:

Linearity: $P\#A(B) = P(\{w : w\#A \in B\})$

As either way of expressing this condition should make clear, the demand for linearity is the demand that the action of the probability revision method be completely determined by the action of that method on each point proposition to which the probability function to be revised assigns non-zero probability. When a probability revision method has this property of linearity, i.e. of having its action completely determined by its action on point propositions, it is harmless and helpful to abuse notation a little further by conflating each world w with the opinionated probability function that assigns probability 1 to world w and takes the value 0 elsewhere. I shall further simplify matters by restricting our attention to what I shall call *sharp* probability revision methods. A revision method is said to be sharp if it

⁶For a general overview of this debate see R. Jay Wallace [?]. I am indebted to Jay wallace for helpful discussions of these issues.

“preserves opinionation,” i.e. if it maps any opinionated probability function to another opinionated function. All of this simplification and notation abuse is intended just to enhance the readability of the following important result:

Linearity Theorem: If an agent’s degrees of desire are given by a sharp, linear revision method, then they satisfy the Desire-as-Belief Thesis.

Proof: Suppose the agent’s probability function is P , and that the agent’s value measure is v . Let $G = \{w : v(w) = 1\}$. Let $A^\circ = \{w : w\#A \in G\}$. Then:

$$\begin{aligned}
 V(A) &= \sum_{w \in A} v(w).P\#A(w) \\
 &= \sum_{w \in G} P\#A(w) \\
 &= P\#A(G) \\
 &= P(\{w : w\#(A) \in G\}) \\
 &= P(A^\circ) \quad \text{QED}
 \end{aligned}$$

This is an interesting result indeed! It tells us that the imposition of the linearity condition on the revision method implicit in the definition of expected value is sufficient to achieve the anti-Humean aim of coherently construing desire as a cognitive attitude toward an objective feature of the world.

How does this fit in with the earlier proof that Desire-as-Belief implies opinionation or indifference? Remember that that proof assumed that degree of desire or expected value was defined the way the evidential decision theorist says it should be, in other words with the agent’s probabilities revised by conditionalization. Since conditionalization is clearly not a linear revision method, there is no conflict between the earlier result and what we have just proved. Still, it would be a useful exercise to go back and check exactly where Lewis’s proof of the Desire-as-Belief result breaks down (as it must!) when the revision method is linear.

The answer to that exercise is that if the revision method is linear, the argument breaks down at the very first step, where Desire-as-Conditional-Belief was derived from Desire-as-Belief. That step appealed to an *Invariance Assumption*, i.e. to the assumption that $V_A(B) = V(A\&B)$, to obtain the

result that $V_A(A) = V(A)$. Now the Invariance Assumption holds for revision by conditionalization, but not for doxastic changes in value that are the result of linear revisions of the agent's degrees of belief.

Theorem: $\#$ is an imaging function iff $\#$ is a sharp, linear probability revision method that satisfies the following three conditions:

1. Success: $w\#A$ is an A -world;
2. Centering: if w is an A -world, then $w\#A = w$;
3. Minimality: $w\#(A \vee B)$ is either $w\#A$ or $w\#B$.

Proof: Omitted.

3 Updating and Supposing

The straight response to these results would be to accept them as establishing on the one hand that a conditional belief cannot be belief in a conditional, and on the other hand that desire cannot simply be belief about what would be good. Should we accept this straight response? I think that the answer is no. I shall argue here that the correct response to these results is to take them as evidence that there are really two kinds of rational belief revision, suited to two different roles. I will refer to these two kinds of belief revision as *updating* and *supposing*.⁷

This alternative response is to claim that there are really two kinds of rational belief revision, each serving a particular purpose, one of which goes by conditionalization and one of which does not. It is the revision method that does not involve conditionalization which allows the Desire-as-Belief Thesis to hold.

It is not hard to give a purely formal characterization of two revision methods that would do the job. But merely describing the formal properties of two revision methods that between them can do all the work we require from an account of rational belief revision does not in itself solve our problem, it just provides a convenient way of restating it. We have good reasons to think that rational belief revision goes by intersection, and we also have good

⁷I borrow this terminology from Brian Skyrms [?].

reasons to say that the rational way to revise is by imaging. The problem remains: on any particular occasion, which is the correct method?

To claim that it is sometimes appropriate to revise by conditionalization, and sometimes appropriate to revise by imaging, as I want to claim here, will seem to be a rather *ad hoc* move unless we can find some independent way of identifying the situations to which each of these methods applies—some way of making the required distinction without reference to the conflicting criteria. I think that this can be done. We have referred to the two methods of belief revision as updating and supposing. Updating is the process by which one genuinely revises belief to account for new information actually received. Supposing, on the other hand, is a process of mere hypothesizing. It is a genuine revision method, but a method one uses not to accommodate information actually received, but to entertain the truth of a proposition—to determine hypothetically what would be the case, were a certain piece of information true. A satisfactory account of belief revision will have to be a two-part account.

It may at this point strike you as just a little strange that David Lewis, author of the original argument against the possibility of Desire-as-Belief, and champion of all things Humean, should be the very same same David Lewis who defends a version of decision theory that deserves to be called “anti-Humean” at least as much as it deserves to be called “causal”? I certainly find it puzzling.

Lewis is clearly not unaware of the issues I have discussed here, but I am not sure why he remains so unperturbed. The sort of line I have been pushing here today is dismissed rather quickly in a parenthetical comment I shall quote in full: Lewis says:

(A famous difficulty need not concern us here. Suppose a certain action would serve as an effective means to your ends, yet at the same time it would constitute evidence—evidence available to you in no other way—that you are predestined inescapably to some dreadful misfortune. Should you perform the action?—Yes; your destiny is not a consideration, since that is outside your control. Do you desire to perform it?—No, you want good news, not bad. Since our topic here is not choiceworthiness but desire, and since the two diverge, we adopt an ‘evidential’ conception of expected value, on which the value of the useful action that brings bad

news is low. Choiceworthiness is governed by a different ‘causal’ conception of expected value.)

This strikes me as a rather weak line to take against the kind cognitivist strategy I have outlined here. Having distinguished newsworthiness from choiceworthiness, is it really as clear as Lewis thinks that the ordinary folk-psychological notion of desire goes with the former rather than the latter? It is crucial to Lewis’s interpretation of the formal result that this be the case, else the Desire-as-Belief Theorem so-called would really have nothing at all to tell us about desire. Is degree of desire really degree of newsworthiness as Lewis so confidently claims? I was at first inclined to think quite the opposite. I would have thought that one’s desires and beliefs were first and foremost revealed by one’s dispositions to choose, and to act. I still think that desire is conceptually pretty tightly connected to the notion of choice. But I wouldn’t want to insist on that here. In fact I suspect that it is simply an indeterminate matter whether the ordinary folk-psychological concept of desire lines up with the notion of newsworthiness or rather with choiceworthiness. The fact is that the newsworthy/choiceworthy distinction is not one that the folk are likely to make, and notoriously when it is forced upon the attention of the uninitiated, for example by asking them to contemplate some bizarre scenario like that in Newcomb’s problem, there is no univocal folk response. That a community of intelligent folk who share a common folk psychological scheme can split so dramatically and irreconcilably on the subject of whether to choose one box or two boxes in the Newcomb problem suggests to me pretty strongly that the matter is one of conceptual indeterminacy. Similarly it seems to me that there are plenty of those who defend anti-Humean accounts of motivation, plenty of cognitivists, plenty of defenders of the notion of objective value and plenty of neo-Kantian rationalists on the topic of pure practical reason, who share, pretty much, a common folk psychological scheme, and who are not simply guilty of some kind of conceptual confusion.

Lewis has written that:

Decision theory (at least if we omit the frills) is not an esoteric science, however unfamiliar it may seem to an outsider. Rather it is a systematic exposition of the consequences of certain well-chosen platitudes about belief, desire, preference and choice. It

is the very core of our common-sense theory of persons, dissected out and elegantly systematised.

I doubt whether it is really a “platitude,” really a part of the “core of our common-sense theory of persons” that degree of desire is to be explicated by a notion of expected value that involves conditionalization rather than a linear probability revision method. Perhaps the whole evidential/causal debate is to count as “frills” rather than “core.” But if so, one can hardly adopt a firm stance on such an issue in the course of giving an argument that is supposed to be a strong *a priori* refutation of cognitivism.

But let me conclude with an argument *ad hominem*. I think that there is no small irony in the fact that the greatest contemporary defender of Humeanism is also one of the co-founders of the only version of decision theory that has the resources to make sense of the notion of objective value and the cognitivist, anti-Humean account of desire.

References

- [1] Arló Costa, Horacio, John Collins, and Isaac Levi. ‘Desire-as-Belief Implies Opinionation or Indifference,’ *Analysis* (1995) pp. 1–5.
- [2] Collins, John. ‘Belief, Desire, and Revision,’ *Mind* (1988) pp.333–342.
- [3] Korsgaard, Christine. ‘Skepticism About Practical Reason,’ *The Journal of Philosophy* (1986) pp.5–25.
- [4] Lewis, David. ‘Desire as Belief,’ *Mind* (1988) pp.333–342.
- [5] Lewis, David. ‘Desire as Belief II,’ *Mind* (1996).
- [6] Savage, Leonard J. *The Foundations of Statistics*, John Wiley and Sons (1954).
- [7] Skyrms, Brian. ‘Updating, Supposing, and MAXENT,’ *Theory and Decision* (1987) pp.225–246.
- [8] van Fraassen, Bas. *Laws and Symmetry*, Clarendon Press, Oxford (1989).
- [9] Wallace, R. Jay. ‘How to Argue About Practical Reason,’ *Mind* (1990).