

Do the EPR Correlations Pose a Problem for Causal Decision Theory?

Adam Koberinski, Lucas Dunlap, William L. Harper

Abstract

We argue that Causal Decision Theory (CDT) is no worse off than Evidential Decision Theory (EDT) in handling entanglement, regardless of one’s preferred interpretation of quantum mechanics. In recent works, Ahmed (2014) and Ahmed and Caulton (2014) have claimed the opposite; we argue that they are mistaken. Bell-type experiments are not instances of Newcomb problems, so CDT and EDT do not diverge in their recommendations. We highlight the fact that a Causal Decision Theorist should take all lawlike correlations into account, including potentially acausal entanglement correlations. This paper also provides a brief introduction to CDT with a motivating “small” Newcomb problem. The main point of our argument is that quantum theory does not provide grounds for favouring EDT over CDT.

1 Introduction

In two recent works (Ahmed 2014; Ahmed and Caulton 2014), Arif Ahmed argues that Causal Decision Theory (CDT) faces a serious challenge from quantum mechanics. He argues that the EPR correlations—correlations predicted by quantum mechanics to obtain between entangled pairs of particles, which have been experimentally verified to an extremely high degree of precision by the so-called Bell experiments¹—can be used to develop a decision problem in which CDT and Evidential Decision Theory (EDT) give diverging answers. As such, Ahmed characterizes this as a version of Newcomb’s problem. Newcomb problems are typically defined as problems for which one’s actions are evidentially relevant to their credences, but not causally relevant.

Ahmed and Caulton (2014) set up betting scenarios based on two devices, A and B. Each has three settings, 1, 2, and 3, and outputs for each setting are binary

¹See Giustina et al. 2015 for a notable recent example.

(‘Yes’ and ‘No’). The observationally deduced facts pertaining to the setup are as follows:

1. whenever the switches on A and B are on the same setting, the devices display the same output;
2. when the switches on A and B are on different settings, they display the same output 25% of the time, and different outputs 75% of the time; and
3. for each individual device, the output on any setting is random—50% ‘Yes’ and 50% ‘No’ (Ahmed and Caulton 2014, cf. p. 4317).

This sort of setup can be realized with entangled quantum systems. According to three possible acausal interpretations of entanglement—corresponding to acausal denials of the no conspiracy assumption (response (A2) in (Ahmed and Caulton 2014)), outcome independence (B2), and parameter independence (C2), respectively—Ahmed and Caulton claim that betting scenarios can be constructed such that CDT and EDT diverge, and CDT and quantum theory diverge. They conclude that this is reason to reject CDT, since it gives differing recommendations to seemingly identical theories in the form of differing interpretations of quantum nonlocality.²

In this paper, we will analyze the details of Ahmed’s position, and argue that his conclusion is not sound. First, Bell experiments are not Newcomb-type problems: one’s actions are neither evidentially nor causally relevant to experimental outcomes. In cases like this, EDT and CDT recommend the same actions. We introduce CDT and EDT in Section 2, and provide an example of a true Newcomb problem. In attempting to make the EPR setup a Newcomb-type problem, Ahmed assumes a strict understanding of CDT for which acausal correlations cannot be used. Violations of Bell inequalities have been experimentally verified, and we argue in Section 3 that both evidential and causal decision theorists should use the experimental evidence to set their credences. Finally, Ahmed imposes constraints on certain interpretations of quantum theory which lead to scenarios in which an agent bets against Bell inequality violations. In Section 4 we show that these constraints are sufficient to limit the set of possible credence functions to Bell-factorizable ones. Anybody who accepts Ahmed’s constraints—including evidential decision theorists—would set their credences against the experimental evidence. We end with a brief conclusion.

²Cavalcanti (2010) makes a similar argument, though his conclusion is slightly different. He claims that the failure of CDT plus acausal interpretations of entanglement in EPR-type betting scenarios means that either CDT is false, or that CDT is capable of making a practical distinction between what were previously thought to be empirically equivalent metaphysical theses.

2 CDT and Newcomb–Type Problems

CDT was first proposed by Robert Stalnaker in correspondence with David Lewis (Stalnaker 1980), and was first popularized by Allan Gibbard and William L. Harper (1978). It was developed as a way of reconciling decision theory with a decision problem first introduced by Robert Nozick, called Newcomb’s Problem (Nozick 1969).³ CDT is an approach to rational decision problems. Gibbard and Harper contrasted CDT with Evidential Decision Theory. EDT seeks to maximize the expected utility EU of an action A , construed as the welcomeness of the news that one is about to perform the act. The utility w of the action–outcome pair is weighted by one’s credence in the outcome, conditional on the action A . Formally,

$$EU(A) = \sum_i Cr(S_i|A) w(A \& S_i). \quad (1)$$

CDT, on the other hand, weights the utility w by one’s credences in the *subjunctive conditionals* $A \rightarrow S_i$, to be read as “If I were to perform A then S_i would obtain.” Formally,

$$EU(A) = \sum_i Cr(A \rightarrow S_i) w(A \& S_i). \quad (2)$$

In cases where one’s actions are evidentially irrelevant to the outcome, $Cr(S_i|A) = Cr(S_i)$, and where the outcome does not depend on one’s actions, $Cr(A \rightarrow S_i) = Cr(S_i)$. Then the expected utility reduces from (1) or (2) to

$$EU(A) = \sum_i Cr(S_i) w(A \& S_i), \quad (3)$$

in the former and latter cases, respectively. Importantly, if one’s actions are *both evidentially irrelevant and independent of the outcome*, then EDT and CDT make the same recommendations: maximize expected utility according to (3).

A fundamental constraint on an agent’s credence in the subjunctive conditional is that it equal the agent’s evaluation of the chance of outcome S_i conditional on A .⁴ That is, the agent must, when possible, set his credence in the subjunctive conditional $Cr(A \rightarrow S_i)$ to the chance that S_i will obtain, given A : $Ch(S_i|A)$. How the agent evaluates the objective conditional chance $Ch(S_i|A)$ will depend on the circumstances. In the EPR experiments at issue, there is now overwhelming evidence that the objective chances agree with those predicted by quantum theory. This point will be relevant for our critique of Ahmed’s argument.

³Nozick credits the physicist William Newcomb with the original formulation of the problem.

⁴See e.g. (Skyrms 2013). Chance here can be viewed as objective and prior to the agent’s credences, or epistemic and built out of a systematization of credences.

2.1 Small Newcomb Problem

The original version of Newcomb’s problem has been extensively discussed in the decision theory literature.⁵ The details of the problem are rather contrived. As such, there is interest in other, more realistic Newcomb–type problems which share the essential decision-theoretic features of Newcomb’s Problem.

Consider this “small” Newcomb–type problem, where the payoffs are small enough for ordinary people to have clear intuitions, and where the “predictor” is not some unexplained, mysterious being, perhaps with *supernatural* predictive powers (as in the original Newcomb Problem), but only a personality tester whose predictions are backed up by a track record which you know.⁶ You are faced with two boxes; one is open and you can see that it contains a fifty-dollar bill. The other is closed, but you are informed that it contains either a thousand dollars or nothing. Two days ago, you completed the questions on a personality test developed by the personality tester. You now know that if, on the basis of your test performance, the tester predicted you would take only the closed box then the thousand dollars was put in it and if the prediction was that you would take both boxes then nothing was put in the closed box.

	$p1B$	$p2B$
$1B$	\$1,000	\$0
$2B$	\$1,050	\$50

Figure 1: Payoff structure for the Small Newcomb Problem.

Suppose the track record of the personality tester, which you know, would support your assigning rational conditional credences $Cr(p1B|1B)$ and $Cr(p2B|2B)$ as both at least as high as .6. If you are using EDT, this would lead to your choosing $1B$.⁷ However, even if your choice of $1B$ would give you evidence that you had been scored as a one-boxer, it would give you no reason to suppose that you could increase your chance of obtaining the \$1,000 by foregoing the \$50.⁸ This dominance argument supports CDT’s sure thing argument for $2B$.

⁵See (Lewis 1981) and (Weirich 2016) for accessible introductions.

⁶This is a version of an example that appeared in (Harper 1993, pp. 84–85).

⁷If utilities are linear with these amounts of money, as they often are for modest amounts like this, then evidential decision theory will recommend one box so long as $Cr(p1B|1B) - Cr(p1B|2B)$ is greater than $50/1000$, or .05.

⁸This would clearly be what Cusbert (2017) would identify as a case of ordinary ‘forwards’ causation.

3 CDT and Causation

It is clear that what is relevant in the example above is that the choice of $1B$ or $2B$ could have no causal influence on whether the predictor put the \$1,000 in the closed box. But CDT doesn't rely on *causation* in any metaphysically significant sense.⁹ Ahmed seems to read too much into the name "Causal Decision Theory." We note that Gibbard and Harper (1978) did not use the term. It was coined and popularized by Lewis (1981).

Although causal talk has been present since the beginning of this discussion, we argue that it is best understood as a proxy for talk of *dependencies that exist in nature*. In short, what is relevant for CDT is not the presence of *causation* in any metaphysically significant sense, but the presence of *physical dependencies*. Of course, cases where there is a causal connection between an action and an outcome will be cases where such dependencies exist. But one of the most fundamental lessons of quantum mechanics is that there are physical dependencies that exist in nature, which can be perfectly well described by our best physical theories, which may nonetheless be acausal.

As shown in equation (2), the relevant quantity to evaluate is one's credence in the subjunctive conditional, and these are evaluated in terms of the *conditional chance* $Ch(S_i|A)$.¹⁰ Skyrms (2013) has shown that the chance view of subjunctive conditionals agrees with the possible world semantics, and helpfully generalizes past strict causation by using conditional chance. These chances can be viewed as objective or epistemic chances. In the context of quantum theory, an objective view of the chances would start with ("bare") quantum theory to generate the conditional chances, while an epistemic view would take the experimental tests of Bell inequality violations as evidence for the correct chances. The loophole-free tests (Giustina et al. 2015) should be enough to convince most that the chance of violating the Bell inequalities is approximately one, so we end up with practically indistinguishable chances as when we start from quantum theory. In either case, causation is unnecessary for getting at the chances predicted by quantum theory, and so something broader than causal relationships may be used in CDT.

As we noted above, the central prescription of CDT is that one should set

⁹In fact, when analyzing counterfactuals for which $P(A \rightarrow S) \neq P(S|A)$, Gibbard and Harper (1978, p. 127) make reference only to worlds that "obey physical laws", not to worlds in which causation plays any role.

¹⁰As mentioned in (Ahmed and Caulton 2014, fn. 21)—and quoted below—Ahmed grants that CDT should deal with conditional chance in this way, and it thus appears clear that our interpretations of what CDT is do not differ drastically. However, here he views conditional chance as revealing only causal dependencies: "the extent to which the conditional chance of Y on X , $Ch(Y|X)$, exceeds the unconditional chance $Ch(Y)$ of Y , reflects the extent to which the occurrence of X *causally promotes* the occurrence of Y " (Ahmed and Caulton 2014, p.4328, emphasis original).

one’s credences in accord with what one believes the chances to be. In the EPR experiments, we now have extraordinary empirical evidence that quantum theory gets the chances right. So CDT will prescribe that we set our credences in accord with these chances. In the original formulation of the position, “cause” was chosen as the most robust indicator for the existence of physical dependencies. But as we’ve learned from quantum nonlocality, cause cannot be the final word on where dependencies exist in nature.

As we’ll describe in more detail below, while it’s true that the notion of a controllable causal connection between the settings or outcomes on the one side of an entangled pair and the outcomes on the other is not a viable explanation of the EPR correlations, on pain of signaling,¹¹ this ‘acausality’ is not as powerful an assumption as Ahmed seems to be supposing. There are robust statistical dependencies in these setups that have been identified and described by quantum mechanics. These are relevant for reasoning using CDT.

To reiterate, CDT *by itself* imposes no particular structure on the conditional chances. In the context of Bell-type experiments, agents familiar with quantum mechanics should use the theory to generate the relevant chances. One can hold that the relevant chances are the quantum ones, either by rejecting the assumption that spacelike separated events are causally independent (as the de Broglie-Bohm theory or any other deterministic theory would have it) or by accepting that there are lawlike correlations that are *not causal*. Though this latter move may seem counterintuitive, it has been defended as the best way to reconcile nonlocality and relativity (see, for example, Shimony (1984), Redhead (1987), Ghirardi (2010), and Myrvold (2016)). Predictions of quantum theory are not interpretation-dependent. Given the robust experimental evidence for violations of the Bell inequalities, acceptance of quantum theory is not even necessary in order to set one’s credences in a non-Bell-factorizable way. Supposing only that the agent is aware of the evidence, their credences should overwhelmingly favour violations of the Bell inequalities in EPR-type setups. On this view, the evidence is interpreted as providing frequency data as evidence for the underlying chances.

4 CDT and the EPR Correlations

As we’ve argued above, EDT and CDT do not diverge in Bell-type experiments. In trying to separate the two, Ahmed relies on a misconstrual of the constraints

¹¹Causal connections are allowed, so long as the agent does not have epistemic access to the detailed state of the world. In the de Broglie–Bohm theory, for example, the choice of experiment on one particle has a causal influence on the distant system. However, an agent cannot have precise knowledge of the true state of the world, and must therefore average over all compatible states. In effect, this averaging washes out any causal connection present in the agent’s credences.

on subjunctive conditionals in CDT. Both causal and evidential decision theorists ought to assume that the relevant chances in Bell-type setups are those predicted by quantum theory. This is sufficient to undermine Ahmed’s argument against CDT; EDT and CDT recommend taking the same bets. Ahmed and Caulton (2014) further impose certain “theoretical assumptions” on acausal interpretations of nonlocality. To understand how this argument works, we will need to describe Ahmed and Caulton’s argument in some detail.¹²

The strategy of their paper is to show that on three of the eight “theoretical assumptions” identified by the authors as reasonable explanations for the existence of the EPR correlations, the EU-score generated by CDT will differ from that generated by EDT. It is this feature of the EPR correlations that the authors use as the justification for the claim that the EPR correlations represent a Newcomb-type problem. Then, the authors claim, if one gives any non-zero credence to the possibility that any one of these theoretical assumptions is the true explanation for the EPR correlations, then they can construct a betting scenario in which CDT will recommend betting against the well-established probabilistic predictions of quantum theory.

In each of the three theoretical assumptions in question, the explanation for the EPR correlations is “acausal”. They are characterized by the authors as follows:

A2 (deny the *no conspiracy* assumption): “[T]he correlation between the setting on the receivers and the prior state of the particles is *acausal*...” (Ahmed and Caulton 2014, p. 4319)

B2 (deny *outcome independence*): “[T]here is a *non-causal* correlation between the hidden variables associated with the particles [...] there is no *causal* connection underlying this correlation, so that any reading ready on any receiver is in fact causally independent of any reading on any other receiver.” (Ahmed and Caulton 2014, p. 4320)

C2 (deny *parameter independence*): “Also to be included in (C2) is any view that denies causation for at least some instances of the EPR experiment. Like (B2), the statistical correlation is treated as brute, i.e. non-causal.” (Ahmed and Caulton 2014, p. 4322)

The authors then rely on this assumption of acausal correlation in each of these three interpretations of the EPR correlations to invoke the transition from

¹²A similar argument appears in Cavalcanti 2010, where Bell factorizability is imposed on the structure of CDT (cf. (Cavalcanti 2010, Section 3.1, eq (11))). Though the details may differ, we think our argument in this section should largely apply there as well. We think it unfair to impose factorizability on CDT, and to do so one must impose parameter independence, outcome independence, *and* the no conspiracy assumptions.

our equation (2) from CDT, to the independent equation (3) (e.g. see Ahmed and Caulton, p. 4316). There is an argument in their Section 4 that is supposed to apply to CDT regardless of which of (A2), (B2) or (C2) one views as the best acausal interpretation. They devise a betting scenario in which two outcomes—where you bet the devices will display the same output when both set to the same setting—will pay $1 - z$, $0 \leq z \leq 1$, while the other two—when the outputs will differ—will pay $-z$. Now, given the facts of the setup (conditions 1-3 in Section 1 above, (4) and (5) in (Ahmed and Caulton 2014, p.4317)), they claim that in calculating the expected utility using EDT, one should give zero credence to the outcomes paying $-z$, while the causal decision theorist cannot rule these outcomes out entirely based solely on the facts of the setup. Thus, depending on the exact values of the payouts and credences, there will be a scenario in which EDT beats CDT in this setup, because EDT reflects the facts of quantum theory in a way that CDT somehow cannot.

Insofar as EDT is licensed to give zero credence to these outcomes based on the past evidence, CDT should do the same.¹³ The experimental evidence indicates that Bell-inequalities are in fact violated, and Bell (1964) proved that these correlations cannot be causally connected in a local manner. Therefore, a causal decision theorist with an acausal interpretation of entanglement should accept the correlations as is, since no plausible (to them) causal mechanism would render these correlations spurious. Ahmed and Caulton concede that one may well give these options zero credence, and therefore argue against each of (A2), (B2), and (C2) individually as well. We turn to these arguments next.

In Section 3, Ahmed and Caulton provide a hidden variables setup to show that an acausal denial of the no conspiracy assumption leads to a losing setup with CDT. However, the setup they provide is a noncontextual hidden variables theory, and the Bell-Kochen-Specker theorem rules out noncontextual hidden variable theories for reproducing quantum entanglement (Bell 1966; Kochen and Specker 1967). CDT does not enter the picture, since the fault lies with noncontextuality. For their argument regarding (B2) and (C2) interpretations, they then impose four constraints on what a causal decision theorist must presume of the chances. Since we will need to refer to the first three in detail, they are quoted fully below.¹⁴

- (i) It is surely absurd to suppose that the choice of bet between ‘het’ and ‘hom’ makes any difference to the reading on either receiver once we are given their settings. We could in any case impose this condition

¹³On a Bayesian view of setting credences, these outcomes should never receive exactly zero credence, but the past body of evidence can be made sufficiently large such that $Cr(AB, het) \approx 0$.

¹⁴Ahmed and Caulton (2014) use ‘hom’ and ‘het’ to indicate cases in which the outcomes on two receivers show the same reading (homogeneous) and different readings (heterogeneous), respectively.

by brute force: i.e., by requiring that you choose the kind of bet on any given run (i.e. between ‘hom’ and ‘het’) after the run is over but before you have had a chance to see the relevant readings. . .

(ii) Since on (B2) the readings on the receivers are causally independent of one another, the chance of either reading on either receiver is independent of the reading on the other receiver, even given the settings on both receivers. . .

(iii) The reading on either receiver is causally independent of the setting on the other receiver given its own setting. (This follows from the assumption that there is no prior instruction set.) (Ahmed and Caulton 2014, pp. 4331-2)

From point (i), the first action they take to be causally irrelevant (and, incidentally, evidentially irrelevant) for the outcome of the betting scenario is the agent’s choice of bet. So far we are in agreement. In both CDT and EDT, we can define, for example,

$$Ch(y_{AyB}|1_A2_B) := Ch(y_{AyB}|1_A2_B, het) = Ch(y_{AyB}|1_A2_B, hom), \quad (4)$$

the chance of outcome ‘Yes’ at A and ‘Yes’ at B given that A was set to 1 and B was set to 2, as independent of the agent’s bet. By point (ii), we have reached a condition on the chances that violates the assumptions of a (B2) type solution. Condition (ii) states that, when coupled with CDT, the chances for the outcomes at A and B should factorize, since there is no causal relationship between the two, for example:

$$Ch(y_{AyB}|1_A2_B) = Ch(y_A|1_A2_B)Ch(y_B|1_A2_B). \quad (5)$$

But the B-type response is a denial of outcome independence, and (B2) is just the stipulation that the outcome dependence is acausal. So point (ii) would be rejected by someone taking a (B2) interpretation, before considerations of EDT versus CDT come into play. Effectively, Ahmed and Caulton are claiming that a causal decision theorist who denies outcome independence in quantum theory must accept outcome independence.

Point (iii) contradicts the assumptions of the (C2) solution.¹⁵ Like point (ii) just mentioned, point (iii) amounts to an assertion of parameter independence, while option (C2) is a denial of parameter independence. These three assumptions are jointly unsatisfactory for at least two of the three interpretations to which Ahmed and Caulton intend them to apply.

¹⁵Though these conditions are brought up in the context of B-type solutions, they claim that the same argument applies equally to (C2) solutions: “the foregoing argument of course applies as well to them as it does to (B2)-type theories: CDT and EDT will give conflicting advice. . . to any agent who accepts (C2)” (Ahmed and Caulton 2014, p. 4335)

Points (i)–(iii) are sufficient to derive a Bell-inequality, which we know to be violated by quantum systems. Neither causal nor evidential decision theorists should accept these constraints. Further, if one does accept the constraints, then CDT and EDT both recommend betting against the evidence. The other assumption is that the randomization procedure responsible for the varying experimental settings on repeated runs of the experiment screens off these settings from the conditions at the source. That is, the probability distribution over experimental settings is independent of common causes or hidden variables. This is the free will, no conspiracy, or no retrocausality assumption (Eq 6, Ahmed and Caulton 2014).

With these assumptions in place, it follows that the expectation values of statistical correlations satisfy Bell-type inequalities. These assumptions can be tested by performing the experiment sufficiently many times that observed frequencies will, with high probability, be close to their expectation values. As mentioned, violations of Bell-type inequalities are well attested in experiment, and it is widely accepted that the assumptions that jointly entail satisfaction of these inequalities are not all true.

In both Ahmed (2014) and Ahmed and Caulton (2014), we find a footnote asserting that this conclusion can be evaded; though the conditions imposed on chances suffice to derive a Bell-type inequality, it is claimed that it does not follow that the assumptions are at odds with the predictions of quantum mechanics.

Note that points (i)–(iii) suffice to derive a Bell inequality for the Ch function. This does not contradict the predictions of quantum mechanics so long as the chances given by Ch do not reflect long-run relative frequencies. It is permitted if, for example, Ch represents single-case chances that vary from case to case. And that is what Ch should represent if conditional chance matters to Causal Decision Theory: for CDT is supposed to be sensitive to the tendency of a setting to causally promote this or that outcome in the particular decision situation to which you are applying it. The situation here is similar to that in Newcomb’s problem, where, even though there is a long-run correlation between one’s choosing one box and this having been predicted, the latter is, on any occasion, conditionally independent of the former with respect to the appropriately causal chance function. This is consistent with the claim that chances control long run frequencies if, as in the Newcomb case and as here, these one-off conditional chances vary from one occasion to the next. (Ahmed 2014, p. 151 fn. 3; Ahmed and Caulton 2014, p. 4332 fn. 21).

If, indeed, the conclusion could be evaded so easily, this would completely undermine the interest in experimental tests of Bell-type inequalities, as the impetus

of the experiments is to rule out theories satisfying the *prima facie* plausible assumptions that lead to those inequalities. The reasons why the conclusion cannot be evaded so easily should be apparent from what has already been said. It is not an assumption of the analysis that the chance-functions $Ch_\lambda(y_{Ay_B}|AB)$ not vary from case to case, and, indeed, expositions of Bell's theorems usually explicitly mention that the extreme instance of this variation, the deterministic case, is included as a special case. In the deterministic case, different outcomes in each trial of the experiment mean that case-by-case chance functions must vary. What is assumed is that, whatever variation there is of conditions at the source, and consequent variation of $Ch_\lambda(y_{Ay_B}|AB)$, the experimental settings are effectively independent of this variation, so that, if one looks at the subset of runs corresponding to a given pair of settings, these constitute a fair sample of conditions at the source. This has the consequence that the observed long-run frequencies can be expected to approximate the average value of $Ch_\lambda(y_{Ay_B}|AB)$, where the average is taken over the distribution of various conditions at the source. If, for every λ , the chance-function Ch_λ generates expectation values that satisfy a Bell-type inequality, then so does any average of these various chance-functions on any distribution over λ .

As mentioned in Section 3, we suspect the root issue here is that Ahmed assumes that for CDT, all causally independent events must have independent chance distributions. This is not the case generically, even in the classical setting. In cases of incomplete information, events that do not causally influence one another may nevertheless be correlated. What is especially significant about quantum correlations is that conditioning on common causes is not sufficient to screen correlations. In effect, Ahmed's argument relies on the unwarranted assumption that the causal decision theorist is constrained to use chance functions attributable to a common cause of the outcomes. This appears to be a denial of the possibility of lawlike (i.e., governing single case chances) acausal correlations. If the causal decision theorist were truly constrained in this way, then this would indeed be a good argument against accepting causal decision theory. However, the causal decision theorist should use the chances predicted by quantum theory, and these chances—at least insofar as they relate to the empirically accessible quantities—are independent of any particular choice of interpretation. Both CDT and EDT will make the same predictions in these experimental setups; since one's actions are independent of and evidentially irrelevant to the outcomes, both will use (3) to calculate expected utility. An agent should set their credences equal to what they take the chances to be, and the empirical evidence overwhelmingly suggests that these are given by quantum theory. This is true regardless of questions about interpretation.

5 Conclusion

Experimental tests of Bell-inequality violations are not situations in which evidential decision theory (EDT) and causal decision theory (CDT) diverge. In trying to make them diverge, Ahmed imposes a strong requirement of statistical independence on causally independent events. Perhaps the most fundamental lesson of quantum mechanics is that this is false. As we have argued, this line of reasoning ignores the fact that quantum correlations, while potentially understood as acausal, are nonetheless lawlike dependencies, which should be included in both CDT and EDT. Nonlocality is a relationship of statistical dependence, even on interpretations where the wings of the Bell-type experiment are causally isolated.

One of the lessons of quantum mechanics is that causation itself is not the ultimate arbiter of when lawlike dependencies exist in nature. When one's actions are independent of the set of outcomes, and have no evidential import, expected utility is equal for CDT and EDT. An agent should set their credences in line with what she believes the objective chances to be. The experimental evidence overwhelmingly suggests that the objective chances are in agreement with the predictions of quantum theory.

Entanglement phenomena make clear that there are lawlike correlations that are not straightforwardly causal. The fact that the structure of the physical dependencies in the world is more complex than we knew prior to the discovery of the EPR correlations, does not entail that there *are not* relevant lawlike dependencies for the application of CDT. Causal decision theory is a framework within which to make decisions, and does *not* impose constraints on the chances that the world may provide one with. Given the success of quantum theory, a rational agent *should* take the chances it predicts into account when betting on Bell-type experiments, and causal decision theory does not prohibit one's doing so.

Acknowledgements

The authors would like to thank the philosophy of physics reading group at Western University for helpful feedback on early drafts, Wayne Myrvold for helpful discussions, and two anonymous reviewers for their comments, which helped to strengthen the argument of this paper. This research was supported by a Joseph-Armand Bombardier Doctoral CGS Award from the Social Sciences and Humanities Research Council of Canada (Adam Koberinski).

References

- Ahmed, A. (2014). *Evidence, Decision, and Causality*. Cambridge University Press.
- Ahmed, A. and A. Caulton (2014). Causal decision theory and EPR correlations. *Synthese* 191(18), 4315–4352.
- Bell, J. S. (1964). On the Einstein-Podolsky-Rosen paradox. *Physics* 1, 195–200.
- Bell, J. S. (1966). On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics* 38(447–452).
- Cavalcanti, E. G. (2010). Causation, decision theory, and Bell’s theorem: a quantum analogue of the Newcomb problem. *The British Journal for the Philosophy of Science* 61(3), 569–597.
- Cusbert, J. (2017). Backwards causation and the chancy past. *Mind*, fzw053.
- Ghirardi, G. (2010). Does quantum nonlocality irremediably conflict with special relativity? *Foundations of Physics* 40(9), 1379–1395.
- Gibbard, A. and W. L. Harper (1978). Counterfactuals and two kinds of expected utility. In C. Hooker, J. Leach, and E. McClennen (Eds.), *Foundations and Applications of Decision Theory*, Volume 13 of *Western Ontario Series in Philosophy of Science*, pp. 125–162. D. Reidel Publishing Company.
- Giustina, M., M. A. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, and A. Zeilinger (2015). Significant-loophole-free test of Bell’s Theorem with entangled photons. *Physical Review Letters* 115(25), 250401.
- Harper, W. L. (1993). Causal and evidential expectations in strategic settings. *Philosophical Topics* 21(1), 79–97.
- Kochen, S. B. and E. Specker (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics* 17, 59–87.
- Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy* 59(1), 5–30.
- Myrvold, W. (2016). Lessons of Bell’s Theorem: Nonlocality, yes; action at a distance, not necessarily. In M. Bell and S. Gao (Eds.), *Quantum Nonlocality and Reality: 50 Years of Bell’s Theorem*, pp. 238–260. Cambridge University Press.

- Nozick, R. (1969). Newcomb's problem and two principles of choice. In N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel*, pp. 114–146. D. Reidel Publishing Company.
- Redhead, M. (1987). *Incompleteness, Nonlocality, and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics*. Oxford University Press.
- Shimony, A. (1984). Controllable and uncontrollable non-locality. In S. Kamefuchi (Ed.), *Foundations of Quantum Mechanics in Light of New Technology*, pp. 225–230. Physical Society of Japan.
- Skyrms, B. (2013). The core theory of subjunctive conditionals. *Synthese* 190, 923–928.
- Stalnaker, R. (1972 [1980]). Letter to David Lewis. In W. L. Harper, R. Stalnaker, and G. A. Pearce (Eds.), *Ifs*, Volume 15 of *Western Ontario Series in Philosophy of Science*, pp. 151–152. D. Reidel Publishing Company.
- Weirich, P. (2016). Causal decision theory. *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition).