

MIXED STRATEGIES AND RATIFIABILITY
IN CAUSAL DECISION THEORY

One of the perceived difficulties with causal decision theory has been its commitment, in one version or another, to such apparently spooky apparatus as counter-factual conditionals (Gibbard and Harper, 1978), objective chance hypotheses (Sobel, 1978), or hypotheses about what you can or cannot causally influence by your choice (Lewis, 1981; Skyrms, 1980, 1984). Mixed strategies in game theory provide for objective chance hypotheses that are not very spooky and that game theory is already committed to anyway. Here it is shown how these relatively uncontroversial chance hypotheses support a natural application of causal decision theory to strategic reasoning in normal form games.

One apparently bizarre feature of causal decision theory is the possibility of unstable choices. This is seen to be of a piece with the instability of non-equilibrium strategies under best response strategic reasoning. I argue that such examples are clear cases of unratifiability in the sense of Jeffrey's (1983, p. 18) intuitive idea and show that this idea has a natural explication as stability of choice in causal decision theory. I claim that, under this explication of ratifiability, only ratifiable alternatives are admissible as rational choices. In many situations the only appropriate ratifiable alternatives will be mixed strategies. I suggest that this is no more a defect in the ratifiability requirement than the corresponding need for mixed strategies in zero-sum games is a defect of the requirement that solutions be equilibria.

1. INDEPENDENCE IN NORMAL FORM GAMES

Suppose you are in a two person normal form game. The strategies available to the other player provide the relevant set of chance hypotheses for you to use in evaluating your own options. Where y is the probability vector corresponding to the mixed strategy $(y_1 B_1 \dots y_n B_n)$, let K_y be the chance hypothesis corresponding to the assumption that the other player executes that strategy. This chance hypothesis determines for each pure strategy B_i an objective chance

$Cy(Bi) = yi$ that it ends up getting realized. The independence assumption built into normal form games requires that all these chances be independent of your choice of strategy. One natural way to represent this requirement is to have each chance hypothesis Ky determine not just the unconditional chances $Cy(Bi) = yi$, but also determine conditional chances of the Bi 's on your strategies so that

$$Cy(Bi|A) = yi = Cy(Bi)$$

for any strategy A (pure or mixed) you might choose. This represents the independence assumption built into the characterization of normal form games as requiring that the Bi 's be stochastically independent of your available choices of strategy in each of the chance hypotheses corresponding to his mixed strategies.

One way to formulate causal decision theory is to take chance hypotheses specifying objective conditional chances of relevant outcome determining states on yours acts as primitive. The chance hypotheses corresponding to the other player's mixed strategies afford such a representation where the outcome determining states are that player's pure strategies. The causal utility is evaluated as a double sum.

$$U(A) = \text{SUM}_i \int_y P(Ky) Cy(Bi|A) dp \cdot U(A, Bi),$$

where the inner sum is an integral representing your epistemic expectation over the relevant alternative chance hypotheses of the objective conditional chance of Bi on A , and $U(A, Bi)$ is your expected utility for the outcome of executing strategy A when the other player ends up performing the pure strategy Bi . Where A is the mixed strategy $(x1A1, \dots, xmAm)$

$$U(A, Bi) = \sum_{j=1} xj \cdot u(Aj, Bi),$$

which is the objective expectation over your pure strategies $A1 \dots Am$ of the utility $u(Aj, Bi)$ you assign to the outcome of the pure strategy pair (Aj, Bi) . Your pure strategy Aj is just a mixed strategy where $xj = 1$, so that the corresponding objective expectation reduces to your utility for the outcome assigned to the pair of pure strategies.

Under the independence assumption built into normal form games the inner sum reduces to your unconditional epistemic probability for the proposition that the other player ends up doing Bi . Therefore, the

double sum formulation reduces to

$$U(A) = \text{SUM}_i P(Bi) \cdot U(A, Bi).$$

To see that this is so note that your epistemic probability for the proposition that the other player ends up doing Bi is just your epistemic expectation over the relevant chance hypotheses of its objective unconditional chance.

$$P(Bi) = \int_y P(Ky) \cdot Cy(Bi) dp.$$

The inner sum in the double sum formulation is your epistemic expectation over these same chance hypotheses of the objective conditional chance of Bi given A . Each of the relevant chance hypotheses makes Bi stochastically independent of A ,

$$Cy(Bi|A) = Cy(Bi) \text{ all } Ky;$$

therefore,

$$\int_y P(Ky) Cy(Bi|A) dp = \int_y P(Ky) Cy(Bi) dp,$$

no matter how your epistemic probability may be distributed over the Ky 's.

Causal utility allows evaluation of an alternative A' from the evidential point of view corresponding to the assumption that A is chosen.

$$U_A(A') = \text{SUM}_i \int_y P(Ky|A) \cdot Cy(Bi|A') dp \cdot U(A', Bi).$$

The inner sum here is your epistemic conditional expectation on A of the objective conditional chance of Bi given A' . Your hypothetical reasoning relative to the assumption that you will execute A requires a conditional redistribution of your epistemic probabilities over the Ky 's, but it cannot alter the chances specified in the Ky 's themselves. The independence built into these chance hypotheses insures

$$\int_y P(Ky|A) \cdot Cy(Bi|A') dp = \int_y P(Ky|A) \cdot Cy(Bi) dp.$$

But,

$$\int_y P(Ky|A) \cdot Cy(Bi) dp = P(Bi|A),$$

since your epistemic conditional probability of Bi on A is just your conditional epistemic expectation on A of the objective chance of Bi . Therefore the double sum formulation again reduces to a simpler one,

$$U_A(A') = \text{SUM}_i P(Bi|A) \cdot U(A', Bi).$$

This gives a very simple way to represent decision-theoretically the sort of strategic reasoning used in game theory.

Strategic reasoning generates non trivial conditional epistemic probabilities players can entertain about each other's strategies. The foregoing results make it clear that the causal independence built into normal form games will not be upset by even extreme epistemic dependence. Consider the following game.

	$B1$	$B2$
$A1$	[(3, 10)	(0, 1)]
$A2$	[(10, 1)	(2, 10)]

Suppose you assign $P(B1|A1) = 1 = P(B2|A2)$. You are sure your opponent will choose some best response to what you choose.¹ Evidential decision theory will recommend $A1$,

$$\begin{aligned} V(A1) &= \text{SUM}_i P(Bi|A1) \cdot u(A1, Bi) = 3 \\ V(A2) &= 2. \end{aligned}$$

But, causal decision theory will recommend $A2$, because

$$\begin{aligned} U(A1) &= P(B1) .3 + P(B2) .0 \\ U(A2) &= P(B1) .10 + P(B2) .2 \end{aligned}$$

so that dominance goes through. What makes it correct that dominance does go through here is the independence built into each of the alternative chance hypotheses. This independence cannot be upset by any redistribution of your epistemic probability over these chance hypotheses, not even by the extreme epistemic redistributions corresponding to $P(B1|A1) = 1 = P(B2|A2)$.

We have learned that evidential decision theory gets the wrong answer here, even when you are certain that the $(A1, B2)$ and $(A2, B1)$

cases will not arise. In order to have the reduction of the matrix

$$\begin{matrix} A1 & \left[\begin{matrix} 3 \\ 2 \end{matrix} \right] \\ A2 & \left[\begin{matrix} 3 \\ 2 \end{matrix} \right] \end{matrix}$$

suggested by these conditional certainties you would need to know that the objective conditional chance of *B1* on *A1* and *B2* on *A2* are both 1, but this is ruled out by the chance independence built into normal form games.²

Another lesson is that if you want to use your evaluation of the subjunctive conditional “If I were to choose *A1* the other player would end up doing *B1*” in order to guide your choice of strategy so that

$$U(A1) = \text{SUM}_i P(A1 \square \rightarrow Bi) \cdot U(A1, Bi),$$

then you had better interpret the conditional in the non-backtracking manner advocated by Gibbard and Harper. Unless you have counterfactual independence so that

$$P(A1 \square \rightarrow Bi) = P(Bi) = P(A2 \square \rightarrow Bi)$$

your conditional will fail to respect the independence assumption built into normal form games. If you use the backtracking interpretation suggested by some writers (e.g., Horgan, 1981) and make

$$P(A1 \square \rightarrow B1) = P(B1|A1)$$

in this game you will not get a recommendation compatible with the game theoretic solution.

A third lesson is that game theory uses both kinds of hypothetical reasoning. Backtracking reasoning is essential to the epistemic dependence represented in a strategic reasoning prior. This is how it is that you hypothetically concentrate all your epistemic probability on the chance hypothesis *K1* that the other player chooses his pure strategy *B1* when you assume you choose *A1*. When however you evaluate the expected utility of the alternative *A2* relative to the epistemic position corresponding to this assumption that you will choose *A1*, you want non-backtracking reasoning that holds fixed the assumption that the other player does *B1*. Thus, it is that

$$U_{A1}(A1) = 3 < U_{A1}(A2) = 10$$

so that choice of *A1* is unstable. The game theoretic solution pair (*A2*, *B2*) represents the only stable pair under the strategic reasoning

assumption that each player will make a best response to what the other does.

2. UNSTABLE CHOICE

It is illuminating to formulate the Death in Damascus problem (Gibbard and Harper, 1978, p. 156) as a zero-sum game. Your opponent is Death. Your pure options are to stay in Damascus $A1$ or to go to Aleppo $A2$. His options are to seek you in Damascus $B1$ or to seek you in Aleppo $B2$. If you are in the place where he seeks you, you die. If not, you get a reprieve. Let -100 be assigned as your utility for meeting Death and 0 as your utility for a reprieve.

	$B1$	$B2$
$A1$	$[-100$	$0]$
$A2$	0	$-100]$

You believe Death is very good at predicting your choice. You assign the following epistemic conditional probabilities,

$$P(B1|A1) \approx 1 \approx P(B2|A2)$$

You also believe that Death doesn't cheat. His choices are made independently of yours, perhaps already made on the basis of his accurate reading of your character and circumstances. The independence assumptions built into normal form games are met.

Under these assumptions both of your pure strategies are unstable. Suppose you decide to stay in Damascus. This gives you evidence that Death will seek you there. Therefore, deciding to stay puts you in an evidential position where going to Aleppo is evaluated as the better option. But, the same problem comes up in reverse if you decide to go to Aleppo instead.

According to evidential decision theory there is no paradox, you are just in the unfortunate situation of being faced with two equally bad alternatives. If the problem is modified to skew your utilities somewhat toward Damascus (say by adding $+5$ to each $A1$ outcome to reflect the idea that staying would give you an opportunity to visit with your mother), then the tie will be broken and evidential decision theory will recommend staying. With causal decision theory, however, choosing to stay will continue to count as irrational because it will still put you in an evidential position from which the other alternative is evaluated as

better. Reed Richter (1985) has argued that this shows a clear superiority of evidential decision theory. He provided the following interesting example to drive this point home.

	<i>B1</i>	<i>B2</i>	
<i>A1</i>	[(10, 10)	(100, 100)]
<i>A2</i>	[(100, 100)	(- 1000, - 1000)]

You expect your partner to choose the same option you do, $P(B1|A1) \approx 1 \approx P(B2|A2)$. Causal independence also holds, as for Death in Damascus. According to Richter, you ought to choose *A1* to avoid the heavy penalty of the (*A2*, *B2*) outcome, but causal decision theory will not endorse this since *A1* is unstable.

The problem can be made more vivid by considering yourself an agent who uses causal decision theory to reason by the sort of deliberation dynamics recently studied by Brian Skyrms (1982, 1984a). Imagine you are faced with the unmodified Death in Damascus problem. You begin to dither. You tentatively incline toward staying – shifting your degree of belief in the proposition that you will stay to, say .6. This new epistemic input generates a corresponding change in your epistemic probability of *B1* the proposition that he seeks you in Damascus. You now recalculate the causal utilities of *A1* and *A2* using these new epistemic probabilities for *B1* and *B2*. The result is that the alternative of going to Aleppo looks better, so you begin to incline toward it. Skyrms has provided schemes specifying how the amount of inclination depends on earlier stages.³ All of them have the feature that in this problem you would not be able to reach a commitment to either pure option, but would end up hung up in a deliberation fixed point where the probability of performing *A1* is $\frac{1}{2}$. In the modified Death in Damascus, you would still be unable to reach commitment to stay in Damascus, and you would continue to get hung up in the same $\frac{1}{2}$ *A1* deliberation fixed point. In Richter's problem the fixed point you would get stuck on assigns .9244 as your probability for *A1*.

Consider an ideally rational deliberator who has reasoned himself into the deliberation fixed point $P(A1) = \frac{1}{2}$ in Death in Damascus. As soon as he believes he is tending toward *A1* with either more or less than probability $\frac{1}{2}$ he will reason himself back into the fixed point. If his reasoning is effective in forming his intentions then he will have reasoned himself into becoming a chance device which if activated

would have objective chance of $\frac{1}{2}$ of ending up choosing $A1$. If he were forced to choose or were to choose by whim then he would execute the mixed strategy $\langle \frac{1}{2}A1, \frac{1}{2}A2 \rangle$ without needing any external chance devices to guide his choice. His own reasoning would generate and sustain the tendency that made the objective chance equally divided between $A1$ and $A2$.

I assume that Death cannot predict the outcome of the chance device (internal or external) that an agent uses to execute a mixed strategy even though he can predict which mixed strategy gets chosen. According to standard zero-sum game theory the mixed strategy $\langle \frac{1}{2}A1, \frac{1}{2}A2 \rangle$ is the best thing you can do in this terrible situation. It guarantees you an expected utility of -50 whatever strategy pure or mixed Death might choose. Similarly, in the modified Death in Damascus case ($+5$ added to each $A1$ outcome) the same $\frac{1}{2}A1$ mixed strategy is still best. It guarantees you -47.5 whatever strategy Death might choose. Any departure from $\frac{1}{2}$ in either problem will get you into trouble. If you play a strategy $\langle xA1, (1-x)A2 \rangle$ where $x > \frac{1}{2}$ ($x < \frac{1}{2}$) you can assume that Death will play $B1$ ($B2$). Under this assumption your expected utility for the mixture $xA1$ is worse than your expected utility for the $\frac{1}{2}A1$ mixture recommended by game theory. The instability of the non-equilibrium strategies in either version of Death in Damascus is just an example of standard game-theoretic strategic reasoning.

Consider Richter's problem. He assumes you believe your partner will do what you do. I assume this extends to mixed strategies so that for each mixture $\langle x, 1-x \rangle$ you assign $P(\langle xB1, (1-x)B2 \rangle | \langle xA1, (1-x)A2 \rangle) = 1$. If you play $.9244A1$ you assume he plays $.9244B1$ and expect utility 16.8 . If you play $xA1$ where $x \neq .9244$ you assume he plays $xB1$. Relative to this assumption the expected utility of $xA1$ is less than 16.8 , but the expected utility of $.9244A1$ will still be 16.8 . The instability of the pure strategies and of all strategies other than $.9244A1$ is in no way pathological. The stable mixture is exactly the right thing to do.

3. RATIFIABLE CHOICE

Richard Jeffrey (1983, p. 18) has proposed ratifiability as a requirement on rational decisions. His basic idea is that an act A is *ratifiable* just in case no alternative has a higher expected utility on the assumption that A is chosen than A itself has. Explicating this idea requires some way to

represent the evaluation of an alternative A' from the epistemic point of view corresponding to the assumption that A is performed. Causal decision theory provides a natural framework for this, as we have seen (section 1). This supports the following explication of Jeffrey's intuitive idea.⁴

A is ratifiable iff $U_A(A) \geq U_A(A')$ all A' .

The problem with the pure strategies in *Death in Damascus* is that they are unratifiable in the sense of this explication of Jeffrey's idea.

I submit that Ratifiability, in this sense, is a necessary condition on non-pathological applications of causal decision theory. The basic recommendation is:

Choose from among your ratifiable alternatives (if there are any) one which maximizes unconditional causal utility.

I regard cases where no act is ratifiable as genuinely pathological and have no qualms about allowing that causal utility theory makes no recommendations in them.

The radical element in this proposal is its restriction to ratifiable alternatives when you maximize unconditional causal utility. This is radical because it can clash with the recommendation to maximize causal utility over all options when no overall maximal option is ratifiable. I think it is quite correct to have ratifiability override unrestricted maximization in this way, because choosing an unratifiable option cannot be a non-pathological application of causal decision theory.⁵ If A is unratifiable then as soon as you commit yourself to A you will give yourself evidence which shows that this commitment is irrational.

Consideration of mixed strategies shows that this requirement is often less radical than it might otherwise seem. Consider the following example Skyrms (1983).

You are to choose one of three shells, and will receive what is under it. A very good predictor has predicted your choice. If he predicts you take shell 1 (A_1) he puts 10¢ under it and nothing under the others; if he predicts that you pick shell 2 (A_2) he puts \$10 under it and \$100 under shell 3. If he predicts you will take shell 3 (A_3) he puts \$20 under it and \$200 under shell 2. In the latter cases he puts nothing under shell 1.

Assume the predictions are causally independent of your choices and that your utilities are linear with these amounts of money.

	<i>B1</i>	<i>B2</i>	<i>B3</i>
<i>A1</i>	.10	0	0
<i>A2</i>	0	10	200
<i>A3</i>	0	100	20

You are to assume that $P(B1|A1) = 1 = P(B2|A2) = P(B3|A3)$. Therefore *A1* is the only ratifiable pure strategy. But, unless $P(B1) > .9986$ at least one of the other alternatives will have higher unconditional utility.

Suppose your initial assignment to $P_0(B1) < .9986$ and you begin to deliberate by Skyrmsian deliberation dynamics, at each stage updating your probability assignments to the *Bi*'s by Jeffrey's rule on inputs consisting of your new assignments to the partition of your acts. You will end up in the deliberation fixed point ($\frac{2}{3}A2, \frac{1}{3}A3$). The corresponding mixed strategy will be ratifiable, and it looks like the intuitively rational choice. For any problem of this structure there will be such a good ratifiable mixed strategy, whether you reach it by deliberation dynamics or consider it directly, so long as you assign conditional probabilities on it by mixing your conditional probabilities on your pure strategies by the mixture weights.

Sometimes conditional probabilities on a mixed strategy might fail to equal the corresponding mixture of the conditional probabilities on the pure strategies. This might happen if you fill in appropriate utilities for the predictor's matrix and apply best response reasoning, even though you keep $P(B1|A1) = 1 = P(B2|A2) = P(B3|A3)$. Various things can happen depending on how you fill in the matrix, but no such assignment will generate either of the non-equilibrium pure strategies as a reasonable solution.

NOTES

¹ In some ways this game is more interesting than the Prisoners' Dilemma, because best response reasoning supports the conditional probability assignments $P(B1|A1) = 1 = P(B2|A2)$ that make epistemic utility legislate the non-solution outcome (*A1, B1*). In a Prisoners' Dilemma

	<i>B1</i>	<i>B2</i>
<i>A1</i>	(3, 3)	(0, 10)
<i>A2</i>	(10, 0)	(2, 2)

the only conditional probabilities consistent with best response strategic reasoning are

$P(B2|A1) = 1 = P(B2|A2)$, which will make evidential decision theory agree with causal decision theory in recommending $(A2, B2)$.

² This shows that the much bandied argument (Levi, 1975; Horgan, 1981; Seidenfeld, 1985) that perfect prediction reduces Newcomb's problem to a choice between a million and a thousand is not valid (see Harper, 1985a for more on this).

³ One of his most recent suggestions (Skyrms, 1984a) would have $P_{n+1}(A_i) = k \cdot P_n(A_i) + \text{COV}(A_i)/k + \text{SUM}_j \text{COV}(A_j)$ where $\text{COV}(A_i) = \text{Max}\{0, U(A_i) - \text{SUM}_j P(A_j) \cdot U(A_j)\}$ and k is an index of caution – the higher k is the more slowly the decision maker moves in the direction of a decision that looks attractive at the time. This implicitly evaluates the mixed strategy $\langle P1A1 \dots PnAn \rangle$ as the expectation by the P 's of the utilities of the A_j 's so that, when positive, $U(A_i) - \text{SUM}_j P(A_j) \cdot U(A_j)$ is a measure of how much better A_i is evaluated than the mixed strategy corresponding to getting stuck at the deliberation point which assigns $P1A1 \dots PnAn$.

⁴ Jeffrey's explication of his intuitive idea is different. See Harper (1985a) for a discussion of this explication and of why the present explication is better.

⁵ See Eells (1985) and Weirich (1985) for additional support for this view. See Rabinowicz (1985) for an excellent and quite fair discussion supporting unconditional utility. In the present paper I am mainly concerned to deal with such apparently counter intuitive examples as Skyrms's shell game. The heart of Rabinowicz's argument is a similar example. In Harper (1985b), I show that the recommendation I propose can make best response reasoning support commitment to weak Nash equilibria. This answers an argument by McClennen (1978) designed to show that game theory is inconsistent with its decision theoretic foundations. It also opens the way to what may turn out to be interesting extensions of the solution concept for some non-cooperative games.

REFERENCES

- Eells, Ellery: 1985, 'Causal Decision Theory', *PSA* (1984).
- Gibbard, Allan and William Harper: 1978, 'Counterfactuals and Two Kinds of Expected Utility', in Hooker et al. Vol. I, pp. 125–262.
- Harper, William: 1985a, 'Ratifiability and Causal Decision Theory Comments on Eells and Seidenfeld', *PSA* (1984), Vol. 2.
- Harper William: 1985b, 'Causal Decision Theory and Game Theory', manuscript.
- Hooker, Clifford A., James J. Leach, and Edward F. McClennen (eds.): 1978, *Foundations and Applications of Decision Theory*, Western Ontario Series, Number 13, Reidel, Dordrecht, Holland.
- Horgan, Terence: 1981, 'Counterfactuals and Newcomb's Problem', *Journal of Philosophy* **68**, 331–356.
- Jeffrey, R. C.: 1983, *The Logic of Decision*, 2nd edn., University of Chicago Press, Chicago and London.
- Levi, Isaac: 1975, 'Newcomb's Many Problems', *Theory and Decision* **6**, 161–175. Reprinted in Hooker et al. (eds.), 1978, pp. 369–383.
- Lewis, David: 1981, 'Causal Decision Theory', *Australasian Journal of Philosophy* **59**, 5–30.
- McClennen, Edward F.: 1978, 'The Minimax Theory and Expected Utility Reasoning', in Hooker et al. (eds.), Vol. I, pp. 337–359.

- Seidenfeld, Teddy: 1985, 'Comments on Eells', *PSA* (1984).
- Skyrms, Brian: 1980, *Causal Necessity*, Yale Press, New Haven, Connecticut.
- Skyrms, Brian: 1982, 'Causal Decision Theory', *Journal of Philosophy* **79**, 695-711.
- Skyrms, Brian: 1983, 'A Problem for Ratifiability', circulated note.
- Skyrms, Brian: 1984, *Pragmatics and Empiricism*, Yale Press, New Haven, Connecticut.
- Skyrms, Brian: 1984a, 'Deliberational Fixed Points', circulated note.
- Sobel, Jordan Howard: 1978, 'Probability, Chance and Choice: A Theory of Rational Agency', manuscript.
- Sobel, Jordan Howard: 1982, 'Expected Utilities of Disjunctions, Decision Matrices and Dominance Arguments', *Synthese* (forthcoming).
- Rabinowicz, Włodzimerz: 1985, 'Ratificationism without Ratification: Jeffrey Meets Savage', *Theory and Decision* **19**, 171-200.
- Richter, Reed: 1984, 'Rationality Revisited', *Australian Journal of Philosophy* **62**, 392-400.
- Weirich, Paul: forthcoming, 'Decision Instability', *Australasian Journal of Philosophy*.

Manuscript received 8 May 1985

Dept. of Philosophy
University of Western Ontario
London, N6A 3K7
Canada