A Note on the Existence of Ratifiable Acts

Joseph Y. Halpern*
Cornell University
Computer Science Department
Ithaca, NY 14853
halpern@cs.cornell.edu
http://www.cs.cornell.edu/home/halpern

August 15, 2018

Abstract

Sufficient conditions are given under which ratifiable acts exist.

1 Introduction

Jeffrey (1983) suggested that you should choose an act that would be best for “the person you expect to be when you will have chosen”. He called an act that satisfied this property *ratifiable*. Whether one should always choose ratifiable acts is a matter of ongoing debate. Skyrms (1990b) argues that there are situations where ratifiable acts are appropriate and other situations where they are not, whereas Harper (1986) seems to suggest that they are always appropriate. Joyce (2012) gives a more recent perspective.

Although I personally do not believe that ratifiable acts are always appropriate, I do not enter into this discussion here. Rather, I examine more closely the question of when ratifiable acts exist. As Rabinowicz (1985) shows by example, in arguably natural problems, there can exist several ratifiable acts, a unique ratifiable act, or no ratifiable “pure” acts (i.e., ones where there is no randomization). Harper (1986) points out that existence of ratifiable act will require randomization. This should not be surprising. Ratifiability has always been understood as an equilibrium notion. It is well known that Nash equilibrium does not always exist in pure strategies; to show that a Nash equilibrium always exists, Nash (1951) had to use *mixed strategies* (where players can randomize over pure strategies). However, as Richter (1984) shows, when there is a cost to randomizing, ratifiable acts may not exist.[1]

While Harper (1986) is willing to ignore settings where there is no ratifiable act, saying “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”, it does not seem to me that charging for randomization is pathological. There is a cognitive cost to randomizing.

*Work supported in part by NSF under grants 1703846 and 1718108 and a grant from the Open Philanthropy Foundation.

[1]This is also a problem for Nash equilibrium; see Section 3.

In any case, there still remains the question of when ratifiable acts do exist. Skyrms (1990a, 1990b) suggests conditions under which ratifiable acts exist, but does not prove any theorem. Indeed, he says "The question of the existence of ratifiable acts ... is a delicate one which calls for careful examination" (Skyrms 1990b).

Skyrms' conditions for the existence of ratifiable acts are essentially that we allow randomization and that there is no cost for randomizing. While the intuition behind these assumptions, and the need for them, is clear, making them precise is not trivial. Moreover, it turns out that these assumptions do not quite suffice for ratifiable acts to exist. Here I show that once we formalize these assumptions appropriately then, together with an additional assumption about how conditioning works, they suffice to show that ratifiable acts exist. Not surprisingly, once we have the appropriate conditions, the proof of the existence of ratifiable acts is quite similar to that of the proof of the existence of Nash equilibrium. In the next section, I provide the details.

2 The existence of ratifiable acts

Before proving the existence of ratifiable acts, I first have to define the notion. Jeffrey (1983) does not provide a formal definition; I use Harper's (1984) formalization, also used by Skyrms (1990a, 1990b) and later authors.

I follow the standard Savage (1954) approach to decision theory. Specifically, I assume that a state space S and a set O of outcomes are given. An *act* is a function from states to outcomes. Let A denote the set of acts. For simplicity in this paper, I assume that S, O, and therefore A are finite and nonempty (although for Savage, it is important that S can be infinite; indeed, his postulates force S to be infinite).

To make sense of ratifiable acts, it has typically been assumed that there is a probability $\Pr$ on $S \times A$ and a utility function u mapping outcomes to utilities (real numbers).[2] Given an act a, let u_a be the function on states defined by setting $u_a(s) = u(a(s))$.

Definition 2.1: An act a is *ratifiable* if, for all acts $a' \in A$,

$$E_{\Pr|a}(u_a) \geq E_{\Pr|a}(u_{a'}), \quad (1)$$

where $E_{\Pr|a}(u)$ denotes the expected utility of a function u on states with respect to the probability $\Pr|a$:

$$E_{\Pr|a}(u) = \sum_{s \in S} (\Pr|a)(s)u(s).$$

∎

I remark that $\Pr|a$ is sometimes denoted $\Pr(\cdot \mid a)$, and $(\Pr \mid a)(s)$ is typically written $\Pr(s \mid a)$; I use these notations interchangeably. Intuitively, Definition 2.1 says that a is ratifiable if the

[2]Actually we do not quite need a probability $\Pr$ on $S \times A$. Rather, for each act $a \in A$, we must have a conditional probability $\Pr_a$ on S. Intuitively, $\Pr_a(s) = \Pr(s \mid a)$, the probability of s given that act a is performed. To my mind, having such a family of conditional probabilities is more reasonable than having a probability on $S \times A$. To be more consistent with the literature, I assume a probability $\Pr$ on $S \times A$ here, although the reader can easily check that all that is ever used are the conditional probability $\Pr(\cdot \mid a)$. I also implicitly assume through that $\Pr(a)$ (or, more precisely, $\Pr(S \times \{a\})$) is positive, so that conditioning on a is well defined.

agent would be at least as happy with a as with any other act a' conditional on a having been performed.

Next, I make precise the notion that there is "no cost for randomizing". Given A (which a view as a set of pure—deterministic—acts), let A^* be the *convex closure* of A, so that if $A = \{a_1, \ldots, a_n\}$, then A^* consists of all acts of the form $\alpha_1 a_1 + \cdots + \alpha_n a_n$, where $\alpha_1 + \cdots + \alpha_n = 1$ and $\alpha_i \geq 0$ for $i = 1, \ldots, n$. The act $\alpha_1 a_1 + \cdots + \alpha_n a_n$ is interpreted as "perform a_1 with probability α_1 and ... and perform a_n with probability α_n". Similarly, let O^* be the convex closure of O. Extend u to a function $u^* : O^* \to I\!R$ in the obvious way, by taking

$$u^*(\alpha_1 o_1 + \cdots + \alpha_n o_n) = \alpha_1 u(o_1) + \cdots + \alpha_n u(o_n). \tag{2}$$

The assumption that there is no cost for randomization is captured by assuming that

$$(\alpha_1 a_1 + \cdots + \alpha_n a_n)(s) = \alpha_1 a_1(s) + \cdots + \alpha_n a_n(s) : \tag{3}$$

the outcome of performing the act $\alpha_1 a_1 + \cdots + \alpha_n a_n$ in state s is, with probability α_1, the outcome of performing act a_1 in state s and ... and with probability α_n, the outcome of performing a_n in state s. That is, performing a convex combination of acts in state s leads to the obvious convex combination of outcomes in O^*; there are no "untoward" outcomes. By way of contrast, if there were a penalty for randomizing, both $a_1(s)$ and $a_2(s)$ might give the outcome \$1,000, although $(.5a_1 + .5a_2)(s)$ gives an outcome of \$0. Of course, this is inconsistent with (3). In combination with (2), (3) ensures that the utility of performing the act $\alpha_1 a_1 + \cdots + \alpha_n a_n$ is the appropriate convex combination of the utilities of performing $a_1, \ldots, a_n$.

Assumption (3) does not suffice to show the existence of ratifiable acts. To explain why, I first introduce the additional assumption that is needed, which involves conditioning on randomized acts. More precisely, we need to extend Pr to a probability $\Pr^*$ on $S \times A^*$. Again, as I pointed out above, all we really need to know is $\Pr(\cdot \mid a^*)$ for acts $a^* \in A^*$. Supposed that we know $\Pr(\cdot \mid a)$ and $\Pr(\cdot \mid a')$ for $a, a' \in A$. What should $\Pr^*(\cdot \mid \alpha a + (1-\alpha)a')$ be? One approach to defining this is to use Jeffrey's rule and take

$$\Pr^*(\cdot \mid \alpha a + (1-\alpha)a') = \alpha \Pr(\cdot \mid a) + (1-\alpha) \Pr(\cdot \mid a'). \tag{4}$$

This seems reasonable: with probability α, a will be played, in which case $\Pr(\cdot \mid a)$ describes the conditional probability; and with probability $(1-\alpha)$, a' will be played, in which case $\Pr(\cdot \mid a')$ describes the conditional probability. While this is reasonable and what I will assume to get the result, note that it is a nontrivial assumption.

To understand the issues involved, consider a variant of Newcomb's problem: There are two boxes, a red box and a blue box, and an extremely accurate predictor has put \$1,000 in one of them and nothing in the other. The agent must choose one of these boxes; that is, we take $A = \{r, b\}$. We can take O to consist of two outcomes: getting \$1,000 and getting nothing. If the predictor predicts that the agent will choose the red box, then he puts the \$1,000 in the blue box; if he predicts that the agent will choose the blue box, he puts the \$1,000 in the red box. Thus, conditional on choosing the red box, the agent would prefer to choose the blue box; conditional on choosing the blue box, the agent would prefer to choose the red box. So neither r nor b is ratifiable.

Now suppose that we extend A to A^*. To make this concrete, suppose that we allow a randomizing device; the agent can set the device to a number α between 0 and 1, and will then play the act $\alpha r + (1-\alpha)b$. Assuming (3) and (4), a straightforward argument shows that choosing $\alpha = 1/2$ gives the unique ratifiable act. However, suppose that we change the game so that if the predictor predicts that the agent randomizes, that is, if the predictor predicts that the agent will set α strictly between 0 and 1, then the predictor will put \$1,000 in the red box.

Let the state space S consist of two states; in s_1, the predictor puts \$1,000 in the red box, and in s_2, the predictor puts \$1,000 in the blue box. We can assume that (3) holds: the outcome of randomizing over acts is the appropriate convex combinations of outcomes. On the other hand, (4) does not hold; randomizing is strong evidence for state s_1.

Assuming (3) and (4) allows us to prove the existence of a ratifiable act in A^*; that is, there are ratifiable acts once we allow randomization.

Theorem 2.2: *If (3) and (4) hold, then there is always a ratifiable act in A^**

Proof: The argument proceeds very much like the argument for the existence of a Nash equilibrium, using Kakutani's (1941) fixed-point theorem. I first state Kakutani's theorem, and then explain all the terms in the statement of the theorem by showing how they hold in the setting of interest.

Kakutani's Fixed-Point Theorem: If X is a non-empty, compact, and convex subset of $\mathbb{R}^n$, the function $f : X \to 2^X$ has a closed graph, and $f(x)$ is non-empty and convex for all $x \in X$, then f has a fixed point, that is, there exists some $x \in X$ such that $x \in f(x)$. ∎

We take the X in Kakutani's theorem to be A^*. If $|A| = n$, then we can identify the element $\alpha_1 a_1 + \cdots + \alpha_n a_n \in A^*$ with the tuple $(\alpha_1, \ldots, \alpha_n) \in \mathbb{R}^n$. With this identification, A^* is a closed subset of $\mathbb{R}^n$ ("closed" just means that if $x_1, x_2, \ldots$ is a sequence of element in A^* that converges to y, then $y \in A^*$); by construction it is convex (if $x, y \in A^*$, then so is $\alpha x + (1-\alpha)y$ for $\alpha \in [0,1]$); it is nonempty, since it includes A and A is non-empty; and it is clearly a bounded subset of $\mathbb{R}^n$ (since all components in a tuple in A^* are between 0 and 1). Since A^* is closed and bounded, by standard results, it is compact.

Let the function f in the theorem be the best-response function. More precisely, given $x \in A^*$, let $g_x : A^* \times A^* \to \mathbb{R}$ be defined by taking

$$g(x,y) = \sum_{s \in S} \Pr{}^*(s \mid x) u^*(y(s));$$

that is, $g(x,y)$ describes how the agent feels about y given that she actually played x. Assumptions (3) and (4), together with the definition of u^*, guarantee that g is a continuous function. Now define

$$f(x) = \{y : \forall z \in A^*(g(x,y) \geq g(x,z))\}.$$

That is, $y \in f(x)$ if the agent feels that playing y would have been one of her best choices, if she actually plays x. Note that x is ratifiable precisely if $x \in f(x)$. The fact that $f(x)$ is a convex set follows from our assumption on the utility function: if $y, y' \in f(x)$, then $g(x,y) = g(x,y') = g(x, \alpha y + (1-\alpha)y')$, so $\alpha y + (1-\alpha)y' \in f(x)$ for all $\alpha \in [0,1]$. Since A^* compact and g is continuous, $g(x,\cdot)$ takes on a maximum in A^*, so $f(x)$ is nonempty. The function f has a closed graph; that is, whenever $x_n \to x$, $y_n \to y$, and $y_n \in f(x_n)$, then $y \in f(x)$. That f has a closed graph follows easily from the fact that g is continuous. Since the assumptions of Kakutani's fixed-point theorem hold, f has a fixed point. As I observed above, this fixed point is a ratifiable act.[3] This completes the proof. ∎

3 Discussion

I have provided sufficient conditions for the existence of ratifiable acts. The key conditions are the assumption that there is no cost for randomizing, which is captured by requiring that the outcome of performing a convex combination of deterministic acts is the convex combination of the outcomes of the individual deterministic acts, and an assumption that can be viewed as saying that randomizing has no evidential value, which is captured by requiring that conditioning on a convex combination of deterministic acts is equivalent to the convex combination of conditioning on the individual deterministic acts.

Interestingly, the fact that charging for randomization can affect the existence of equilibrium has also been noted in a game-theoretic context. Halpern and Pass (2015) consider a setting where, associated with each action in a normal-form game, there may be a cost. They show that, in this setting, the analogue of Nash equilibrium may not exist. Interestingly, their example involves a cost for randomization. Consider first the standard rock-paper-scissors game where, as usual, rock beats scissors, scissors beats paper, and paper beats rock. If a player's choice beats the other player's choice, then he gets a payoff (utility) of 1 and the other player gets -1; if both players make the same choice, they both get 0. As is well known, this game has a unique Nash equilibrium, where players randomize, choosing each action with probability $1/3$.

Now suppose that we modify the game slightly, and charge ϵ for randomizing. That is, if a player randomizes, no matter how he randomizes, his payoffs are decreased by $\epsilon > 0$ (so, for example, he gets $1 - \epsilon$ if his choice beats the other player's choice and $-\epsilon$ if there is a draw). The intuition here is that randomization is cognitively expensive. In any case, with this utility function, it is easy to see that there is no Nash equilibrium, no matter how small ϵ is (as long as it is positive). For suppose that we have a Nash equilibrium (a_1, a_2). If a_1 involves randomization, then player 1 can do better by deterministically playing the best response to whichever action player 2 puts the highest probability on. (If player 2 puts equal probability on several actions, player 1 chooses one of them and plays a best response to that.) This deterministic choice gets at least as high a payoff as a_1 ignoring the ϵ cost of randomizing, and thus must be strictly better than a_1 when the cost of randomization is taken into account. Thus, player 1 does not randomize in equilibrium. A similar argument shows that player 2 does not randomize in equilibrium either. But it is clear that there is no equilibrium for this game where both players use deterministic strategies.

Halpern and Pass also show that if randomization is free (and a few other technical conditions hold), then the analogue of Nash equilibrium does exist in their framework. Very roughly speaking, "randomization is free" means that the cost of a randomized act of the form $\alpha_1 a_1 + \cdots + \alpha_n a_n$ is the sum of α_1 times the cost of a_1 and ... and α_n times the cost of a_n. The similarity to (3) should be clear. Note that this condition is violated in the rock-paper-scissors example.

[3] Exactly the same arguments show that, under the assumptions above, Skyrms' (1990a) deliberation process, which is intended to model which strategy to play in a game or which act to choose in a decision problem, also converges.

But what should we do if this condition does not hold? Charging extra for randomization, as in the rock-paper-scissors example, does not seem so unreasonable, given the cognitive costs of randomizing. Interestingly, as observed by Halpern and Pass (2015), there are rock-paper-scissors tournaments (indeed, even a rock-paper-scissors world championship), and books written on rock-paper-scissors strategies (Walker and Walker 2004). Championship players are clearly not randomizing uniformly (they could not hope to get a higher payoff than an opponent by doing this). This leaves open the question of what an appropriate solution concept would be in such a situation. This question also applies in the case of ratifiability, but is perhaps less serious in that context. I agree with Joyce (2012) that the agent should choose the best act according to the prescriptions of causal decision theory. If this results in some regret after the fact, then so be it—it was still the right choice. But for those who believe that ratifiability is a normative requirement, this is an important question that must be addressed.

References

Halpern, J. Y. and R. Pass (2015). Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory 156*, 246–268.

Harper, W. (1984). Ratifiability and causal decision theory: comments on Eells and Seidenfeld. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Vol. 2*, pp. 213–228.

Harper, W. (1986). Mixed strategies and ratifiability in causal decision theory. *Erkenntnis 24*(1), 25–36.

Jeffrey, R. C. (1983). *The Logic of Decision*. Chicago: University of Chicago Press. second edition.

Joyce, J. M. (2012). Regret and instability in causal decision theory. *Synthese 187*(1), 123–145.

Kakutani, S. (1941). A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal 8*(3), 457–459.

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics 54*, 286–295.

Rabinowicz, W. (1985). Ratificationism without ratification: Jeffrey meets Savage. *Theory and Decision 19*(2), 171–200.

Richter, R. (1984). Rationality revisited. *Australasian Journal of Philosophy 62*(4), 392–403.

Savage, L. J. (1954). *Foundations of Statistics*. New York: Wiley.

Skyrms, B. (1990a). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.

Skyrms, B. (1990b). Ratifiability and the logic of decision. *Midwest Studies in Philosophy XV*, 44–56.

Walker, G. and D. Walker (2004). *The Official Rock Paper Scissors Strategy Guide*. New York: Simon & Schuster, Inc.