A Note on Kuhn's Theorem*

Adam Brandenburger

Stern School of Business
New York University
44 West Fourth Street
New York, NY 10012, United States of America
adam.brandenburger@stern.nyu.edu

Abstract

We revisit Kuhn's classic theorem on mixed and behavior strategies in games. We frame Kuhn's work in terms of two questions in decision theory: What is the relationship between global and local assessment of uncertainty? What is the relationship between global and local optimality of strategies?

This note is a homage to Kuhn's classic theorem on the replacement of mixed by behavior strategies in games [$\text{Ku}_0 50$, $\text{Ku}_0 53$]. It reframes Kuhn's work as two results in decision theory—i.e., in the context of trees involving a decision maker and Nature. The motivation is to see the meaning of Kuhn's work at this basic level.

The decision-theoretic framing in this note is in accordance with the so-called epistemic approach to game theory. Under the epistemic approach, a game is a multi-player decision problem—more exactly, a collection of decision problems, one for each player. In line with decision theory, a player is assumed to form a (subjective) probability assessment over the strategies chosen by other players in the game, and to choose an optimal strategy under this assessment. The questions à la Kuhn are then: (a) the relationship between global and local assessments; and (b) the relationship between global and local optimality.

The epistemic approach is 'the other way round' from the traditional approach to game theory. Under the traditional approach, we talk about a mixed strategy of a player, not another player's global assessment of the first player's deterministic choice of (pure) strategy. Likewise, we talk about

* My thanks to Scott Borg, Amanda Friedenberg, Rena Henderson, Konrad Grabiszewski, Mike Moldoveanu, Ben Polak, John Pratt, Nikolaj Siggelkow, Pat Sileo, and Gus Stuart for important input. A referee made very helpful suggestions. Financial support from Harvard Business School and the Stern School of Business is gratefully acknowledged.

Johan van Benthem, Dov Gabbay, Benedikt Löwe (eds.). *Interactive Logic. Proceedings of the 7th Augustus de Morgan Workshop, London.* Texts in Logic and Games 1, Amsterdam University Press 2007, pp. 71–88.

a behavioral strategy of a player, not another player's system of local assessments about the first player. The mixed-behavioral framing is Kuhn's, of course.

In Section 5, we expand on the significance for Kuhn's Theorem of taking an epistemic perspective on games.

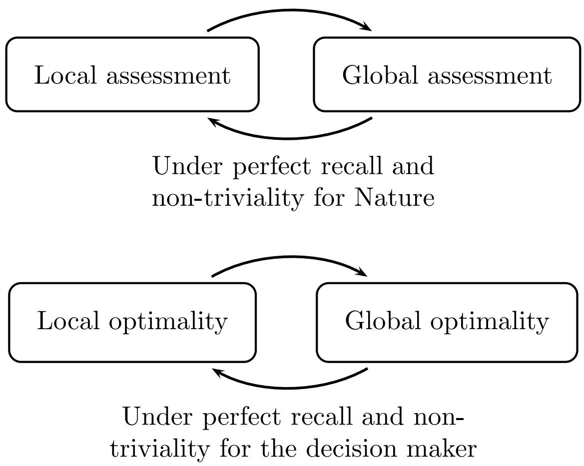

Figure 1. Summary of results

Figure 1 is a summary of the two results we cover. Each result is in two parts. For the first, we have: (i) Given a system of local probability assessments by the decision maker, i.e., an assessment over Nature's moves at each of Nature's information sets, there is a global assessment over Nature's strategies ("states") that yields the same probability of each path through the tree. (ii) If Nature has perfect recall and all chance nodes are non-trivial, then, given a global assessment by the decision maker, there is an equivalent system of local assessments. For the second result, we have: (i) If a strategy of the decision maker is locally optimal, i.e., optimal at each information set of the decision maker, then it is globally ("ex ante") optimal. (ii) If the decision maker has perfect recall and all decision nodes are non-trivial, then, if a strategy is globally optimal, it is locally optimal.

There is also a sufficiency result (which follows from part (ii) of the first result): Assume perfect recall and non-triviality for Nature. Then, it is enough to know the system of local assessments associated with any global assessment, to know which strategies are globally optimal. Putting this together with part (ii) of the second result gives: Assume perfect recall and non-triviality for both the decision maker and Nature. Then, to determine if a strategy is locally optimal, it is enough to know the system of local assessments of the decision maker.

We acknowledge that much (all?) of the contents of this note may be well known. Still, we hope that a self-contained presentation will be useful.

1 Decision Trees

A decision tree will be a two-person game in extensive form, where one player is the decision maker and the other is Nature. We now give formal definitions following Kuhn [Ku$_0$50, Ku$_0$53] (also the presentation in Hart [Ha$_4$92]).

Definition 1.1. A (**finite**) **decision tree** consists of:

(a) A set of two players, one called the **decision maker** and the other called **Nature**.

(b) A finite rooted tree.

(c) A partition of the set of non-terminal nodes of the tree into two subsets denoted N (with typical element n) and M (with typical element m). The members of N are called **decision nodes**, and the members of M are called **chance nodes**.

(d) A partition of N (resp. M) into **information sets** denoted I (resp. J) such that for each I (resp. J):

   (i) all nodes in I (resp. J) have the same number of outgoing branches, and there is a given 1-1 correspondence between the sets of outgoing branches of different nodes in I (resp. J);

   (ii) every path in the tree from the root to a terminal node crosses each I (resp. J) at most once.

Note: The focus of the well-known literature on the "paradox of the absent-minded driver" (Piccione and Rubinstein [Pi$_0$Ru$_1$97]) is on non-Kuhn trees—specifically, trees that fail condition (d.ii) above. (See also Isbell [Is57].) We consider only Kuhn trees.

For each information set I (resp. J), number the branches going out of each node in I (resp. J) from 1 through $\#I$ (resp. $\#J$) so that the 1-1 correspondence in (d.i) above is preserved.

Definition 1.2. A **strategy** (of the decision maker) associates with each information set I, an integer between 1 and $\#I$, to be called the **choice** of the decision maker at I. Let S denote the set of strategies of the decision maker. A **state of the world** (or **state**) associates with each information set J, an integer between 1 and $\#J$, to be called the **choice** of Nature at J. Let Ω denote the set of states.

Note that a pair (s, ω) in $S \times \Omega$ induces a unique path through the tree.

Definition 1.3. Fix a path p through the tree and a strategy s. Say p is **allowed under** s if there is a state ω such that (s, ω) induces p.

Definition 1.4. Fix a path p through the tree and a state ω. Say p is **allowed under** ω if there is a strategy s such that (s, ω) induces p.

Definition 1.5. Fix a node n in N and a strategy s. Say n is **allowed under** s if there is a state ω such that the path induced by (s, ω) passes through n. Say an information set I is **allowed under** s if some n in I is allowed under s.

Definition 1.6. Say the decision maker has **perfect recall** if for any strategy s, information set I, and nodes n and n^* in I, node n is allowed under s if and only if node n^* is allowed under s.

Definition 1.7. Say a node n in N is **non-trivial** if it has at least two outgoing branches.

Definition 1.8. Fix a node m in M and a state ω. Say m is **allowed under** ω if there is a strategy s such that the path induced by (s, ω) passes through m. Say an information set J is **allowed under** ω if some m in J is allowed under ω.

Definition 1.9. Say Nature has **perfect recall** if for any state ω, information set J, and nodes m and m^* in J, node m is allowed under ω if and only if node m^* is allowed under ω.

Definition 1.10. Say a node m in M is **non-trivial** if it has at least two outgoing branches.

Example 1.11. Figure 2 depicts a case of imperfect recall for the decision maker. (The circular node belongs to Nature and the square nodes belong to the decision maker.) Let s be the strategy that chooses B at information set I (and b, say, at information set I'). Then node n is allowed under s but node n^* is not.

Example 1.12. Figure 3 is the standard example of imperfect recall for Nature. Let ω be the state that chooses U at information set J (and u, say, at information set J'). Then node m is allowed under ω but node m^* is not.

Define a relation of precedence on information sets I of the decision maker, as follows: Given two information sets I and I', say that I **precedes** I' if there are nodes n in I and n' in I' such that the path from the

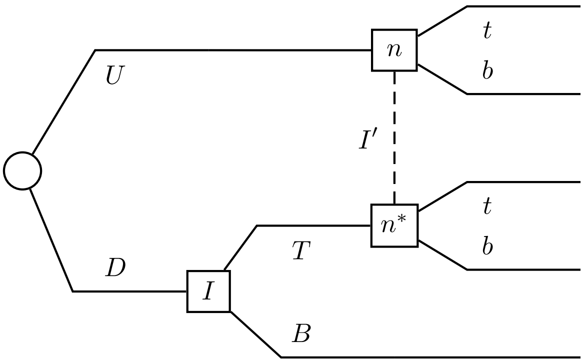

FIGURE 2.

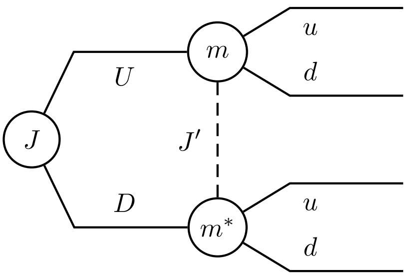

FIGURE 3.

root to n' passes through n. It is well known that if the decision maker has perfect recall and all decision nodes are non-trivial, then this relation is irreflexive and transitive, and each information set I has at most one immediate predecessor. (Proofs of these assertions can be constructed from arguments in Wilson [$\text{Wi}_0$72]. See also the appendix to this note.) Of course, the parallel statements hold for Nature.

In [$\text{Ku}_0$53], Kuhn observes that perfect recall implies that a player remembers: (i) all of his choices at previous nodes; and (ii) everything he knew at those nodes. The following two lemmas formalize these observations. (The proofs are in the appendix.) Again, parallel statements hold for Nature. (Lemma 1.13 will be used later.)

Lemma 1.13. Suppose the decision maker has perfect recall and all decision nodes are non-trivial. Fix information sets I and I', and strategies s

and s'. Suppose that I' is allowed under both s and s', and I precedes I'. Then I is allowed under both s and s', and s and s' coincide at I.

Continuation of Example 1.11. Let s choose T at I, and s' choose B at I. Then I' is allowed under both s and s', as is I. But s and s' differ at I. So Lemma 1.13 fails without perfect recall. In words, the decision maker forgets at I' whether he chooses T or B at I.

Next, write

$$[I] = \{\omega : I \text{ is allowed under } \omega\}.$$

Lemma 1.14. Suppose the decision maker has perfect recall and all decision nodes are non-trivial. Fix information sets I and I'. If I' succeeds I, then $[I'] \subseteq [I]$.

Continuation of Example 1.11. We have $[I] = \{D\}$ and $[I'] = \{U, D\}$. So Lemma 1.14, too, fails without perfect recall. In words, the decision maker knows at I that Nature doesn't choose U, but forgets this at I'.

A brief comment on the literature on these 'structural' aspects of perfect recall. Bonanno makes a nice distinction between "action" and "choice" [Bo$_0$04]. (The same action can be taken at different information sets.) He offers definitions of "knowing the actions you previously took" and "knowing what you previously knew", and shows that together these conditions characterize perfect recall. Ritzberger provides several characterizations of perfect recall [Ri$_0$99]. Van Benthem studies game trees as structures for logical languages, and, in particular, provides a dynamic epistemic logic-based axiomatization of games satisfying perfect-recall like conditions [vB01a]. Also related are the temporal logics in Halpern, van der Meyden, and Vardi [Ha$_0$vMVa04]. See [Bo$_0$04, Section 6] for further discussion of the literature.

2 Global and Local Probabilities

We now define the global and local probabilities on the tree, and then state Kuhn's Theorem.

Definition 2.1. A **global** probability measure on the tree is a probability measure on the set of states Ω.

Definition 2.2. A system of **local** probability measures on the tree associates with each information set J of Nature, a probability measure on the set of choices at J.

Fix a global measure σ, and a system of local measures $\pi(\cdot; J)$. Fix also a path p through the tree. Let J_1 be the first information set of Nature crossed by p, and let j_1 be the choice at J_1 that lies on p. Define $J_2, j_2, \ldots, J_K, j_K$

similarly, where J_K is the last information set of Nature crossed by p, and j_K is the choice at J_K that lies on p. (Note this is well defined, by condition (d.ii) of Definition 1.1. Also, we don't indicate the dependence of the indices $1, \ldots, K$ on p, but no confusion should result.)

Definition 2.3. The **global** probability of p is

$$\lambda(p; \sigma) = \sigma(\{\omega : p \text{ is allowed under } \omega\}).$$

Definition 2.4. The **local** probability of p is

$$\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = \prod_{k=1}^{K} \pi(j_k; J_k).$$

Continuation of Example 1.12. To practice these definitions, let σ assign probability $1/2$ to (U, u) and probability $1/2$ to (D, d). Also, let $\pi(U; J) = \pi(D; J) = 1/2$, and $\pi(u; J') = \pi(d; J') = 1/2$. Suppose p is the path induced by (U, u). Then the global probability of p is $\lambda(p; \sigma) = 1/2$, and the local probability of p is $\mu(p; \pi(\cdot; J), \pi(\cdot; J')) = 1/2 \times 1/2 = 1/4$.

We now state Kuhn's Theorem in the form of the following two results, and give proofs in the notation of this note.

Theorem 2.5. Fix a system of local measures $\pi(\cdot; J)$. There is a global measure σ such that for any path p,

$$\lambda(p; \sigma) = \mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)).$$

Proof. It will be convenient to write Ω as a product space $\prod_J \boldsymbol{C}(J)$, where $\boldsymbol{C}(J)$ denotes the set of choices at information set J. Given a state ω, write $\omega(J)$ for the Jth coordinate of ω, i.e., the choice ω makes at J. Set $\sigma(\omega) = \prod_J \pi(\omega(J); J)$. This is readily seen to define a probability measure on Ω.

Now fix a path p, and let $J_1, j_1, \ldots, J_K, j_K$ be defined as earlier. For $k = 1, \ldots, K$, let

$$A_k = \{\omega : \omega(J_k) = j_k\}.$$

Note that $\sigma(A_k) = \pi(j_k; J_k)$. The set of ω such that p is allowed under ω is $\bigcap_{k=1}^{K} A_k$, and, since σ is a product measure,

$$\sigma\Big(\bigcap_{k=1}^{K} A_k\Big) = \prod_{k=1}^{K} \sigma(A_k) = \prod_{k=1}^{K} \pi(j_k; J_k),$$

as required. Q.E.D.

We need one more definition, used in the proof of the next result.

Definition 2.6. Fix a global measure σ and an information set J. The **localized** probability measure at J is given by

$$\pi(j; J, \sigma) = \frac{\sigma(\{\omega : J \text{ is allowed under } \omega, \text{ and } \omega(J) = j\})}{\sigma(\{\omega : J \text{ is allowed under } \omega\})},$$

if $\sigma(\{\omega : J \text{ is allowed under } \omega\}) > 0$. If $\sigma(\{\omega : J \text{ is allowed under } \omega\}) = 0$, define $\pi(\cdot; J, \sigma)$ arbitrarily.

Theorem 2.7. Suppose Nature has perfect recall and all chance nodes are non-trivial. Fix a global measure σ. There is a system of local measures $\pi(\cdot; J)$ such that for any path p,

$$\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = \lambda(p; \sigma).$$

Proof. Fix a path p and $J_1, j_1, \ldots, J_K, j_K$ as earlier. Let

$$\begin{aligned} B &= \{\omega : p \text{ is allowed under } \omega\}, \\ C_k &= \{\omega : J_k \text{ is allowed under } \omega, \text{ and } \omega(J_k) = j_k\}, \\ D_k &= \{\omega : J_k \text{ is allowed under } \omega\}, \end{aligned}$$

for $k = 1, \ldots, K$.

We show that for $k = 1, \ldots, K-1$, $C_k \subseteq D_{k+1}$. Suppose p passes through node m in J_k. Fix ω such that J_k is allowed under ω. Then by perfect recall, node m is allowed under ω. That is, there is a strategy s such that the path induced by (s, ω) passes through m. Let (s', ω') induce the path p. Then s and s' coincide at all information sets of the decision maker crossed by p from the root to m. Indeed, we can take $s = s'$. We know that $\omega'(J_k) = j_k$. Therefore, if $\omega(J_k) = j_k$, the path induced by (s', ω) coincides with the path induced by (s', ω'), from the root to J_{k+1}. Certainly then, J_{k+1} is allowed under ω, as required.

We next show that for $k = 1, \ldots, K-1$, $D_{k+1} \subseteq C_k$. Let (s', ω') be as above, so that certainly J_{k+1} is allowed under ω'. Fix ω such that J_{k+1} is allowed under ω. Since J_k precedes J_{k+1}, Lemma 1.13 (stated for Nature) implies that: (i) J_k is allowed under ω; and (ii) ω and ω' coincide at J_k, i.e., $\omega(J_k) = j_k$.

We now have that for $k = 1, \ldots, K-1$, $C_k = D_{k+1}$. By definition, $C_k \subseteq D_k$ for each k. This shows that the C_k's are a decreasing sequence.

Given a global measure σ, define a system of local measures $\pi(\cdot; J)$ by setting $\pi(\cdot; J) = \pi(\cdot; J, \sigma)$.

Note that $C_K = B$, since J_K is the last information set of Nature crossed by p. It follows that if $\lambda(p; \sigma) = \sigma(B) > 0$, then $\sigma(C_k) > 0$ and $\sigma(D_k) > 0$

for each k. We then have

$$\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = \prod_{k=1}^{K} \frac{\sigma(C_k)}{\sigma(D_k)}. \tag{2.1}$$

But the numerator of each term in (2.1) cancels with the denominator of the next term, leaving

$$\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = \frac{\sigma(C_K)}{\sigma(D_1)}.$$

We already have $\sigma(C_K) = \sigma(B)$. Also, $D_1 = \Omega$, since J_1 is the first information set of Nature crossed by p, so $\sigma(D_1) = 1$. This establishes that

$$\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = \sigma(B) = \lambda(p; \sigma), \tag{2.2}$$

as required.

Now suppose $\lambda(p; \sigma) = \sigma(B) = 0$. If $\sigma(D_k) > 0$ for each k, we still get (2.1), from which we get (2.2), and so $\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = 0$, as required. We have $\sigma(D_1) = 1$, so the remaining case is that $\sigma(D_k) = 0$ for some $k = 2, \ldots, K$. Choose the minimum such k. Then $\sigma(D_{k-1}) > 0$. Also $\sigma(C_{k-1}) = 0$, since $C_{k-1} = D_k$. Thus $\pi(j_{k-1}; J_{k-1}) = \sigma(C_{k-1})/\sigma(D_{k-1}) = 0$, so that $\mu(p; \pi(\cdot; J_1), \ldots, \pi(\cdot; J_K)) = 0$, as required. Q.E.D.

Continuation of Example 1.12. Theorem 2.7 fails without perfect recall. To see this, let σ assign probability 1/2 to (U, u) and probability 1/2 to (D, d), as before. Then, in particular, we need $\pi(U; J) \times \pi(u; J') = 1/2$, $\pi(U; J) \times \pi(d; J') = 0$, and $\pi(D; J) \times \pi(d; J') = 1/2$, which is impossible.

3 Global and Local Optimality

Next we define global and local optimality of a strategy of the decision maker.

Definition 3.1. A **payoff function** (of the decision maker) is a map $V : S \times \Omega \to \mathbb{R}$ satisfying $V(s, \omega) = V(s', \omega')$ whenever (s, ω) and (s', ω') induce the same path.

Definition 3.2. Fix a probability measure σ on Ω. A strategy s is **globally optimal** under σ if

$$\sum_{\omega \in \Omega} \sigma(\omega) V(s, \omega) \geq \sum_{\omega \in \Omega} \sigma(\omega) V(r, \omega)$$

for every other strategy $r \in S$.

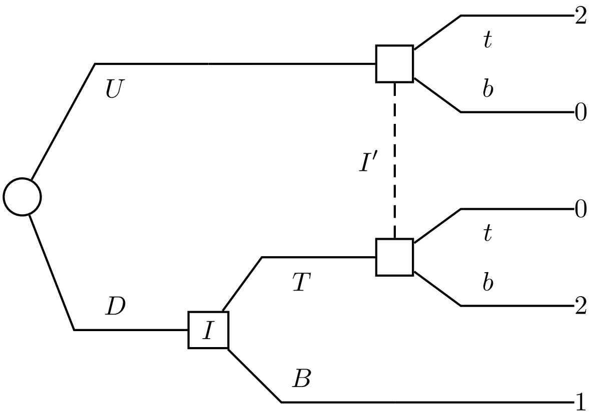

FIGURE 4.

Definition 3.3. Fix a probability measure σ on Ω. Fix also a strategy s and an information set I that is allowed under s and satisfies $\sigma([I]) > 0$. Then s is **locally optimal at** I under σ if

$$\sum_{\omega \in \Omega} \sigma(\omega|[I])V(s,\omega) \geq \sum_{\omega \in \Omega} \sigma(\omega|[I])V(r,\omega)$$

for every other strategy $r \in S$ under which I is allowed. Strategy s is **locally optimal** under σ if for every information set I that is allowed under s and satisfies $\sigma([I]) > 0$, it is locally optimal at I under σ.

In words, a strategy is globally optimal if it is expected-payoff maximizing under the (unconditional) measure σ. It is locally optimal if it is expected-payoff maximizing under each conditional measure $\sigma(\omega|[I])$ that is defined (and where I is allowed under the strategy).

Example 3.4. Figure 4 is Figure 2 with payoffs added for the decision maker. Let σ assign probability 2/3 to U and 1/3 to D. Then Tt, Bt, Tb, and Bb yield expected payoffs of 4/3, 5/3, 2/3, and 1/3, respectively—so Bt is (uniquely) globally optimal under σ.

As noted before, $[I] = \{D\}$ and $[I'] = \{U, D\} = \Omega$. So, $\sigma([I]) > 0$ and $\sigma([I']) = 1$. Also, both I and I' are allowed under all four strategies. It follows that local optimality at I' is the same as global optimality. At I, we find that Tt, Bt, Tb, and Bb yield conditional expected payoffs of 0, 1, 2, and 1, respectively—so, in fact, no strategy is locally optimal under σ.

Theorem 3.5. Fix a probability measure σ on Ω. If a strategy s is locally optimal under σ, then it is globally optimal under σ.

Proof. Partition Ω into cells $I_0, I_1, \ldots, I_L$, where $\omega \in I_\ell$, for some $\ell = 1, \ldots, L$, if I_ℓ is the first information set of the decision maker allowed under ω, and $\omega \in I_0$ if there is no information set of the decision maker allowed under ω.

Write

$$\sum_{\omega\in\Omega} \sigma(\omega)V(s,\omega) = \sum_{\ell=0}^{L} \sigma([I_\ell]) \sum_{\omega\in\Omega} \sigma(\omega|[I_\ell])V(s,\omega),$$

where we take $\sigma(\cdot|[I_\ell])$ to be arbitrary if $\sigma([I_\ell]) = 0$. Suppose s is locally optimal under σ, but not globally optimal under σ. Then there must be another strategy r such that

$$\sum_{\omega\in\Omega} \sigma(\omega)V(r,\omega) > \sum_{\omega\in\Omega} \sigma(\omega)V(s,\omega).$$

But

$$\sum_{\omega\in\Omega} \sigma(\omega)V(r,\omega) = \sum_{\ell=0}^{L} \sigma([I_\ell]) \sum_{\omega\in\Omega} \sigma(\omega|[I_\ell])V(r,\omega),$$

so there must be $\ell = 0, \ldots, L$ such that $\sigma([I_\ell]) > 0$ and

$$\sum_{\omega\in\Omega} \sigma(\omega|[I_\ell])V(r,\omega) > \sum_{\omega\in\Omega} \sigma(\omega|[I_\ell])V(s,\omega).$$

Note that in fact $1 \leq \ell \leq L$, since, on I_0, $V(\cdot,\omega)$ is independent of the strategy. Since it is a first information set of the decision maker, I_ℓ must be allowed under r. This implies that s is not locally optimal under σ at I_ℓ, a contradiction. Q.E.D.

Theorem 3.6. Suppose the decision maker has perfect recall and all decision nodes are non-trivial. Fix a probability measure σ on Ω. If a strategy s is globally optimal under σ, then it is locally optimal under σ.

Notes: (i) By finiteness, a globally optimal strategy always exists under any σ. So Theorem 3.6 implies, in particular, that under the given conditions a locally optimal strategy also exists under any σ. (ii) Kline [Kl$_1$02] contains a stronger result, based on a weakening of perfect recall.

Proof. Suppose that s is globally optimal and that, contra hypothesis, there is an information set I allowed under s and satisfying $\sigma([I]) > 0$, such that

$$\sum_{\omega\in\Omega} \sigma(\omega|[I])V(s,\omega) < \sum_{\omega\in\Omega} \sigma(\omega|[I])V(r,\omega) \tag{3.1}$$

for some other strategy $r \in S$ under which I is allowed.

Construct the strategy q that coincides with r at I and all succeeding information sets of the decision maker, and coincides with s elsewhere.

We first show that if $\omega \in [I]$, then $V(q,\omega) = V(r,\omega)$.

From $\omega \in [I]$, there is a node n_1 in I and a strategy s_1 such that the path induced by (s_1,ω) passes through n_1. Since I is allowed under s, there is a node n_2 in I and a state ω_2 such that the path induced by (s,ω_2) passes through n_2. By perfect recall, there is then a state ω_3 such that the path induced by (s,ω_3) passes through n_1. Now consider the information sets I' crossed by the path from the root to n_1. Since the paths induced by (s_1,ω) and (s,ω_3) both pass through n_1, the strategies s_1 and s must coincide at these sets. Similarly, consider the information sets J crossed by the path from the root to n_1. Since the paths induced by (s_1,ω) and (s,ω_3) both pass through n_1, the states ω and ω_3 must coincide at these sets. Therefore, the path induced by (s,ω) must pass through n_1.

We can repeat the argument with strategy r in place of strategy s, to conclude that the path induced by (r,ω) must also pass through n_1. But then, using the definition of strategy q, the paths induced by (q,ω) and (r,ω) must be the same. Thus $V(q,\omega) = V(r,\omega)$, as required.

Next, we show that if $\omega \in \Omega\backslash[I]$, then $V(q,\omega) = V(s,\omega)$. From $\omega \in \Omega\backslash[I]$, the path induced by (s,ω) does not cross I, and therefore does not cross any information set of the decision maker that succeeds I. Consider the information sets I' that are in fact crossed by the path induced by (s,ω). By construction, the strategies q and s coincide at each such I'. Thus the paths induced by (q,ω) and (s,ω) are the same, and so $V(q,\omega) = V(s,\omega)$, as required. Write

$$\sum_{\omega\in\Omega} \sigma(\omega)V(q,\omega) = \sigma([I]) \sum_{\omega\in\Omega} \sigma(\omega|[I])V(q,\omega) + \sigma(\Omega\backslash[I]) \sum_{\omega\in\Omega} \sigma(\omega|\Omega\backslash[I])V(q,\omega),$$

where $\sigma(\cdot|\Omega\backslash[I])$ is arbitrary if $\sigma(\Omega\backslash[I]) = 0$. We have

$$\begin{aligned}\sum_{\omega\in\Omega} \sigma(\omega)V(q,\omega) &= \sigma([I]) \sum_{\omega\in\Omega} \sigma(\omega|[I])V(r,\omega) + \sigma(\Omega\backslash[I]) \sum_{\omega\in\Omega} \sigma(\omega|\Omega\backslash[I])V(s,\omega) \\ &> \sigma([I]) \sum_{\omega\in\Omega} \sigma(\omega|[I])V(s,\omega) + \sigma(\Omega\backslash[I]) \sum_{\omega\in\Omega} \sigma(\omega|\Omega\backslash[I])V(s,\omega) \\ &= \sum_{\omega\in\Omega} \sigma(\omega)V(s,\omega),\end{aligned}$$

where the inequality uses (3.1) and $\sigma([I]) > 0$. But this contradicts the global optimality of s. Q.E.D.

Continuation of Example 3.4. Theorem 3.6 fails without perfect recall (for the decision maker). Indeed, we saw that only Bt is globally optimal, but it is not locally optimal at I.

4 A Sufficiency Result

We now establish a sufficiency result: Assume perfect recall and non-triviality for Nature. Then, it is enough to know the localized probabilities associated with any probability measure, to know which strategies are globally optimal.

First some notation. As in Section 2, given a path p, write J_1 for the first information set of Nature crossed by $p, \ldots, J_K$ for the last information set of Nature crossed by p (and suppress the dependence on p). Also, in this section it will be helpful to write $W(p)$ for the payoff $V(s, \omega)$, if (s, ω) induces the path p.

Definition 4.1. Fix two (global) probability measures σ and τ on Ω. Say σ and τ are **locally equivalent** if for each path p, $\lambda(p; \sigma) > 0$ if and only if $\lambda(p; \tau) > 0$, and, in this case,

$$\mu(p; \pi(\cdot; J_1, \sigma), \ldots, \pi(\cdot; J_K, \sigma)) = \mu(p; \pi(\cdot; J_1, \tau), \ldots, \pi(\cdot; J_K, \tau)).$$

In words, two measures are locally equivalent if they give rise to the same localized probability of each path that gets positive (global) probability.

Example 4.2. In the tree in Figure 5, let σ assign probability 1/2 to (U, u) and probability 1/2 to (D, d), and τ assign probability 1/4 to each of (U, u), (U, d), (D, u), and (D, d). It can be checked that σ and τ are locally equivalent.

Here is the sufficiency result:

Theorem 4.3. Suppose Nature has perfect recall and all chance nodes are non-trivial. Let σ and τ be probability measures on Ω that are locally equivalent. Then for any strategy s,

$$\sum_{\omega \in \Omega} \sigma(\omega) V(s, \omega) = \sum_{\omega \in \Omega} \tau(\omega) V(s, \omega).$$

Proof. Write

$$\sum_{\omega \in \Omega} \sigma(\omega) V(s, \omega) = \sum_{\{p:\, p \text{ is allowed under } s\}} \sum_{\{\omega:\, (s,\omega) \text{ induces } p\}} \sigma(\omega) W(p). \quad (4.1)$$

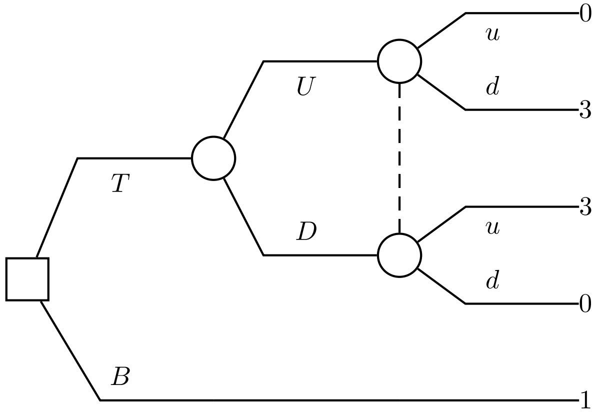

FIGURE 5.

Now, if (s,ω) induces p, then certainly p is allowed under ω. Conversely, suppose p is allowed under ω. That is, there is an s' such that (s',ω) induces p. Suppose also that p is allowed under s. That is, there is an ω' such that (s,ω') induces p. It follows (by arguing forwards along the path p) that (s,ω) must also induce p. Using the definition of $\lambda(p;\sigma)$, this establishes that (4.1) can be rewritten as

$$\sum_{\omega\in\Omega}\sigma(\omega)V(s,\omega) = \sum_{\{p:p \text{ is allowed under } s\}} \lambda(p;\sigma)W(p). \tag{4.2}$$

By the same argument we can write

$$\sum_{\omega\in\Omega}\tau(\omega)V(s,\omega) = \sum_{\{p:p \text{ is allowed under } s\}} \lambda(p;\tau)W(p). \tag{4.3}$$

Fix a path p. By the proof of Theorem 2.7,

$$\lambda(p;\sigma) = \mu(p;\pi(\cdot;J_1,\sigma),\ldots,\pi(\cdot;J_K,\sigma)), \tag{4.4}$$
$$\lambda(p;\tau) = \mu(p;\pi(\cdot;J_1,\tau),\ldots,\pi(\cdot;J_K,\tau)). \tag{4.5}$$

Fix a path p. Using local equivalence, we have either: (i) $\lambda(p;\sigma) = \lambda(p;\tau) = 0$, or (ii) $\lambda(p;\sigma) > 0$ and $\lambda(p;\tau) > 0$. In case (ii), $\lambda(p;\sigma) = \lambda(p;\tau)$, by (4.4), (4.5), and local equivalence again. Thus (i) and (ii) together establish that (4.2) and (4.3) are equal, as required. Q.E.D.

Corollary 4.4. Suppose Nature has perfect recall and all chance nodes are non-trivial. Fix probability measures σ and τ on Ω that are locally equivalent. Then a strategy s is globally optimal under σ if and only if it is globally optimal under τ.

Continuation of Example 4.2. Theorem 4.3 and Corollary 4.4 fail without perfect recall (for Nature). The measures σ and τ as above are locally equivalent. Yet T yields an expected payoff of 0 under σ, and an expected payoff of 3/2 under τ. (So B is globally optimal under σ, while T is globally optimal under τ.)

5 Discussion

Corollary 4.4 and Theorem 3.6 can be put together as follows: Assume perfect recall and non-triviality for both the decision maker and Nature. Then, to determine if a strategy is locally optimal, it is enough to know the localized probabilities of the decision maker.

For (even local) optimality, then, the analyst need only know how the decision maker sees the tree locally. We don't need to know the decision maker's global assessment of the tree.

But this does assume perfect recall and non-triviality. Perfect recall for the decision maker has a clear interpretation as a memory requirement (refer back to the end of Section 1 and also the references there). But what does perfect recall for Nature mean?[1] We'll give an answer for the game context, as analyzed under the epistemic approach.

For the application of the decision tree set-up (Definition 1.1) to a game, the decision maker is to be thought of as one player, Ann say. All the remaining players—Bob, Charlie, ...—are grouped together as Nature. This is because, under the epistemic approach, the strategies chosen by Bob, Charlie, ... are jointly uncertain as far as Ann is concerned, and so subject to joint probability assessment.

When, then, might Nature have perfect recall? One case is if there is just one other player, Bob, and he has perfect recall. The situation is different if there are two or more other players, even if each of these players has perfect recall. For example, suppose in Figure 5 that Ann chooses T or B, Bob chooses U or D, and Charlie chooses u or d. Then Bob and Charlie each has perfect recall, but if Ann assigns probability 1/2 to (U, u) and probability 1/2 to (D, d), there is no equivalent local assessment.

We could require Ann's global assessment to be a product of a global assessment of Bob's strategy and a global assessment of Charlie's strategy. Then, working with each assessment separately, we could find an equivalent local assessment. But an independence requirement like this is not in the spirit of the epistemic approach to games, which treats correlations as the norm.

[1] I am grateful to a referee for asking this question.

Of course, there will be special cases where 'overall' perfect recall still holds. (An obvious one is if the game has perfect information.) But, in general, we should work with global not local assessments by the players.

Appendix

Lemma A1. Suppose the decision maker has perfect recall and all decision nodes are non-trivial. Fix an information set I, nodes n_1 and n_2 in I, and an information set I' containing a node n_3 on the path from the root to n_1. Then there is a unique node in I' (not necessarily distinct from n_3) lying on the path from the root to n_2.

Proof. First note that by part (d.ii) of Definition 1.1, there cannot be more than one node in I' lying on the path from the root to n_2.

Now suppose, contra hypothesis, there is no node in I' lying on the path from the root to n_2. Let c denote the choice at n_3 that lies on the path to n_1. By non-triviality, there is a choice d, different from c, at I'. Construct a strategy s as follows: (i) at I', let s specify the choice d; (ii) at an information set crossed by the path from the root to n_2, let s specify the choice that lies on this path; (iii) at any other information set, let s be arbitrary. (Note that, by hypothesis, the information set I' does not fall under (ii), so s is well defined.) By construction, the node n_2 is allowed under s, while n_1 is not allowed under s. This contradicts perfect recall.

Q.E.D.

Lemma A2. Suppose the decision maker has perfect recall. Fix an information set I and nodes n_1 and n_2 in I. Fix also an information set I' containing nodes n_3 and n_4 (not necessarily distinct) where n_3 lies on the path from the root to n_1, and n_4 lies on the path from the root to n_2. Then the choice at n_3 that lies on the path to n_1 is the same as the choice at n_4 that lies on the path to n_2.

Proof. Let c be the choice at n_3 that lies on the path to n_1, and let d be the choice at n_4 that lies on the path to n_2. Suppose, contra hypothesis, that $c \neq d$. Construct a strategy s as follows: (i) at an information set crossed by the path from the root to n_2, let s specify the choice that lies on this path; (ii) at any other information set, let s be arbitrary. Note that s specifies d at I'. It follows that n_2 is allowed under s, while n_1 is not allowed under s. This contradicts perfect recall. Q.E.D.

We use Lemmas A1 and A2 in the proofs below of Lemmas 1.13 and 1.14 in the text. We also note that Lemma A1, together with part (d.ii) of Definition 1.1, easily implies the facts stated in Section 1: The precedence relation on information sets is irreflexive and transitive, and each information set I has at most one immediate predecessor.

Proof of Lemma 1.13. Since I' is allowed under s, there is a node n'_1 in I' and a state ω_1 such that the path induced by (s, ω_1) passes through n'_1. Likewise, since I' is allowed under s', there is a node n'_2 in I' and a state ω_2 such that the path induced by (s', ω_2) passes through n'_2. Since I precedes I', there are nodes n in I and n' in I' such that the path from the root to n' passes through n.

Lemma A1 then implies that there is a node n_1 in I (not necessarily distinct from n) lying on the path from the root to n'_1. That is, the path induced by (s, ω_1) passes through n_1. This establishes that I is allowed under s.

Likewise, Lemma A1 implies that there is a node n_2 in I (not necessarily distinct from n) lying on the path from the root to n'_2. That is, the path induced by (s', ω_2) passes through n_2. This establishes that I is allowed under s'.

Lemma A2 implies that the choice at n_1 that lies on the path to n'_1 is the same as the choice at n_2 that lies on the path to n'_2. Thus s and s' coincide at I. Q.E.D.

Proof of Lemma 1.14. Consider a state ω in $[I']$. By definition, there is a node n' in I' and a strategy s such that the path induced by (s, ω) passes through n'.

Since I precedes I', there are nodes n_1 in I and n_2 in I' such that the path from the root to n_2 passes through n_1.

Lemma A1 then implies there is a node n_3 in I (not necessarily distinct from n_1) such that the path from the root to n' passes through n_3. That is, the path induced by (s, ω) passes through n_3. Thus ω lies in $[I]$, as required. Q.E.D.

References

[AuHa$_4$92] R.J. Aumann & S. Hart, eds. *Handbook of Game Theory, Vol. 1.* North-Holland, 1992.

[Bo$_0$04] G. Bonanno. Memory and perfect recall in extensive games. *Games and Economic Behavior* 47(2):237–256, 2004.

[DrTu$_0$Wo$_0$57] M. Drescher, A. Tucker & P. Wolfe, eds. *Contributions to the Theory of Games, Vol. III*, *Annals of Mathematics Study* 39. Princeton University Press, 1957.

[Ha$_0$vMVa04] J. Halpern, R. van der Meyden & M. Vardi. Complete axiomatizations for reasoning about knowledge and time. *SIAM Journal on Computing* 33(3):674–703, 2004.

[$Ha_4$92] S. Hart. Games in extensive and strategic form. In [$AuHa_4$92, pp. 19–40].

[Is57] J. Isbell. Finitary games. In [$DrTu_0Wo_0$57, pp. 79–96].

[$Kl_1$02] J.J. Kline. Minimum memory for equivalence between *ex ante* optimality and time-consistency. *Games and Economic Behavior* 38(2):278–305, 2002.

[$Ku_0$50] H. Kuhn. Extensive games. *Proceedings of the National Academy of Sciences* 36(10):570–576, 1950.

[$Ku_0$53] H. Kuhn. Extensive games and the problem of information. In [$Ku_0Tu_0$53, pp. 193–216].

[$Ku_0Tu_0$53] H. Kuhn & A. Tucker, eds. *Contributions to the Theory of Games, Vol. II*, *Annals of Mathematics Study* 28. Princeton University Press, 1953.

[$Pi_0Ru_1$97] M. Piccione & A. Rubinstein. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20(1):3–24, 1997.

[$Ri_0$99] K. Ritzberger. Recall in extensive games. *International Journal of Game Theory* 28(1):69–87, 1999.

[vB01a] J. van Benthem. Games in dynamic-epistemic logic. *Bulletin of Economic Research* 53(4):219–248, 2001.

[$Wi_0$72] R. Wilson. Computing equilibria of two-person games from the extensive form. *Management Science* 18(7):448–460, 1972.