Degrees of indeterminacy of games

by

Andreas Blass (Ann Arbor, Mich.)

We propose to define a pre-ordering of the class of all infinite two-person zero-sum games with perfect information and two-valued payoff function (henceforth simply called games) by placing one game below another if the latter is easier for player 1 to win than the former. (We use 0 and 1 as names for the two players.) Thus, $A \leqslant B$ means that there is a "strategy" by means of which 1 can win B if he is permitted to consult an "oracle" capable of winning A as player 1. A game for which 1 has a winning strategy is at the top of the ordering, because he can simply use this strategy and ignore the oracle. A game for which 0 has a winning strategy is at the bottom of the ordering, for there can be no oracle for such a game. The undetermined games, those for which neither player has a winning strategy, are distributed between these extremes.

It is, of course necessary to replace this talk about oracles with a precise mathematical definition of the ordering. We shall give two definitions, corresponding to two conventions as to how player 1 is to interrogate the oracle. For one ordering (the strong one), he interrogates the oracle by acting as player 0 in a play of A against the oracle, occasionally interrupting this game to make a move in B, using what he has learned from the oracle to select this move. For the weak ordering, he proceeds similarly except that he may, from time to time, abandon a play of A and start over, asking the oracle about a different opening in A. He may later return to a previously abandoned play of A and continue it. The oracle is obliged to play consistently in that, as long as it is confronted with the same moves in two of the plays of A, it answers the same way in both plays.

It is clear, even from this informal description, that the weak ordering will require a more complex definition than the strong one. Also, the weak ordering lacks the following symmetry property enjoyed by the strong ordering: If B is easier for player 1 than A, then A is easier for 0 than B.

On the other hand, the weak ordering yields (if we identify equivalent games) Brouwerian lattices, while the lattices obtained from the strong ordering are not even modular. More importantly, the weak ordering seems to correspond more accurately than the strong one to the intuitive concept "easier to win".

1. Games. We use the usual conventions that an ordinal is the set of smaller ordinals and a cardinal is an initial ordinal; in particular, $\omega = \aleph_0$ is the set of natural numbers. For any set X, the set of infinite sequences of elements of X is ${}^{\omega}X$, while the set of finite sequences is $\mathrm{Fin}(X)$. The set ${}^{\omega}X \cup \mathrm{Fin}(X)$ of all sequences (of length $\leqslant \omega$) from X is $\mathrm{Seq}(X)$. If $s \in \mathrm{Seq}(X)$ and $i \leqslant \mathrm{length}(s)$, then $\bar{s}(i)$ is the initial segment $\langle s(0), \ldots, s(i-1)\rangle$ of s of length i.

A *game* on a set X is a function A from $\mathrm{Seq}(X)$ to $\{0, 1\} = 2$. For nonempty X, we think of a game A on X as being played as follows. There are two players, 0 and 1. The moves occur in an ω-sequence, and at each move one of the players selects an element of X. Thus, an $s \in {}^{\omega}X$ is produced. At move n, the sequence $\bar{s}(n)$ has already been produced and is known to both players; $s(n)$ is chosen by player $A(\bar{s}(n))$. When the play s is completed, player $A(s)$ is the loser. Note that the function A serves a dual purpose. On finite sequences, it indicates who is to move next; on infinite sequences, it indicates who has lost the play. It will be convenient to have a single function playing both roles. Let us also agree that, if A is a game on X, if s is a sequence some of whose terms are not in X, and if $s(k)$ is the first such term, then $A(s)$ means $A(\bar{s}(k))$. This convention may be interpreted as saying: If one of the players selects something outside X, then it remains his move forever, and he loses.

Let A be a game on X. A *strategy* for player i in A is a function σ from $\mathrm{Fin}(X) \cap A^{-1}(i)$ into X. It is a *winning strategy* if

(1) $$\forall s \in {}^{\omega}X[\forall n \in \omega\big(A(\bar{s}(n)) = i \to s(n) = \sigma(\bar{s}(n))\big) \to A(s) \neq i].$$

A strategy σ for i should be thought of as a set of instructions for player i, telling him, if the sequence t has already been played and $A(t) = i$, to play $\sigma(t)$ next. The condition (1) for a winning strategy says that if i follows these instructions then he wins. If one of the players has a winning strategy in a game, then that game is said to be *determined* and to be a win for that player. Clearly, no game can be a win for both players. Gale and Stewart [2] have given an example of an undetermined game on $\{0, 1\}$. (Although we shall always assume the axiom of choice, upon which the Gale-Stewart construction depends, one can prove without this axiom that there is an undetermined game on $\aleph_1$; see [3], page 217).

Our definitions of game and strategy make sense even when X is empty, although such games cannot be played since neither player has a move. As $\mathrm{Seq}(0)$ consists of just the empty sequence $\langle\ \rangle$, there are exactly two games, $\underline{O}$ and $\underline{I}$ on 0, defined by

$$\underline{O}(\langle\ \rangle) = 1, \quad \underline{I}(\langle\ \rangle) = 0.$$

According to our definitions, $\underline{O}$ is a win for 0, while $\underline{I}$ is a win for 1, and, in either case, the empty function is a winning strategy.

We shall define several operations on games. The first of these is just to interchange the roles of the two players. For any game A, we define the game $\sim A$ on the same set by

$$(\sim A)(s) = 1 - A(s).$$

Clearly, $\sim\sim A = A$, and $\sim A$ is a win for one player iff A is a win for the other.

Our second operation, which will give the greatest lower bound in the orderings to be defined later, acts on an indexed family $\{A_i |\ i \in I\}$ of games and yields a new game $\bigwedge_{i \in I} A_i$ played as follows. Player 0 begins by choosing an $i \in I$, and, from then on, the players play A_i. Formally, $\bigwedge_{i \in I} A_i$ is the game on $I \cup \bigcup_{i \in I} X_i$ defined by

$$(\bigwedge_{i \in I} A_i)(s) = \begin{cases} 0 & \text{if } s = \langle\ \rangle, \\ 0 & \text{if } \quad s(0) \notin I, \\ A_i(t) & \text{if } \quad i \in I \text{ and } s = \langle i \rangle * t \end{cases}$$

(where $*$ means concatenation). In this definition, the first clause says 0 moves first, and the second says, in effect, that he must choose an $i \in I$, for otherwise he loses. Once 0 has chosen $i \in I$, both players must choose moves in X_i, for the first player to move outside X_i loses. Finally, according to the third clause, the rest of the game proceeds just like A_i. It is easy to check that $\bigwedge_{i \in I} A_i$ is a win for 0 (resp. 1) iff one (resp. each) of the A_i's is.

There is an operation $\bigvee$, dual to $\bigwedge$, which will give the least upper bound in our orderings. $\bigvee_{i \in I} A_i$ is just like $\bigwedge_{i \in I} A_i$ except that the first move, the choice of which A_i to play, belongs to 1 rather than 0. Formally,

$$\bigvee_{i \in I} A_i = \sim\bigwedge_{i \in I} \sim A_i.$$

Then $\bigvee_{i \in I} A_i$ is a win for 0 (resp. 1) iff each (resp. one) of the A_i's is.

We shall define one more operation in this section, the tensor product. It will be used in the definitions of the orderings of games. The tensor product $\bigotimes_{i \in I} A_i$ of the indexed family $\{A_i |\ i \in I\}$ is played as follows. Player 0 begins by choosing an $i \in I$, and the players start to play A_i.

However, at any of his moves, 0 may abandon A_i and choose another index $i' \in I$. Then both players start to play $A_{i'}$ until 0 again changes games. He may change games as often as he likes, and may also return to a previously abandoned game; that game is then resumed at the point where it was abandoned (not at the beginning). Player 0 wins if at least one of the A_i's is finished (i.e. infinitely many moves are made in A_i) and won by 0; otherwise 1 wins. Turning to a formal definition of $\bigotimes_{i\in I} A_i$, we assume that A_i is on X_i and that I is disjoint from $\bigcup_{i\in I} X_i$. (Otherwise, replace I by another index set.) Then $\bigotimes_{i\in I} A_i$ is a game on $I \cup \bigcup_{i\in I} X_i$. Given a sequence s of moves in $\bigotimes_{i\in I} A_i$, it will be necessary to extract the subsequence of moves in a particular A_i. In general, if $s \in \mathrm{Seq}(I \cup J)$ where I and J are disjoint, and if $i \in I$, then $s_i \in \mathrm{Seq}(J)$ is the subsequence of s consisting of those terms of s in J which are preceded by a term in I and whose last predecessor in I is i. (Strictly speaking, s_i depends not only on s and i but also on the decomposition of $I \cup J$ into I and J.) Note that s_i may be finite even if s is infinite. Now we can define

$$(\bigotimes_{i\in I} A_i)(s) = \begin{cases} 1 & \text{if } (\exists k < \mathrm{length}(s))\ (s(k) \in I \text{ and } (\bigotimes_{i\in I} A_i)(\bar{s}(k)) = 1), \\ 1 & \text{if } (\exists i \in I \cap \mathrm{Range}(s))\ A_i(s_i) = 1, \\ 0 & \text{otherwise}. \end{cases}$$

(This is a definition by induction on the length of s.) The first clause, in effect, prevents 1 from choosing elements of I, for once he makes such a choice he loses. Let us restrict our attention to sequence s to which this clause does not apply. As long as no element of I appears in s, it is 0's move (in particular, 0 moves first); if 0 never gets around to choosing an element of I, he loses. If s is finite and $i \in I$ appears in s but is not the last element of I in s (game A_i was started, but, after s, the players are in another A_j), then $A_i(s_i) = 0$. For, if $s(k)$ is the first element of I to appear in s after the last occurrence of i, then, since clause 1 doesn't apply, $(\bigotimes_{i\in I} A_i)(\bar{s}(k)) = 0$. By clause 2, $A_i(\bar{s}(k)_i) = 0$, but $\bar{s}(k)_i = s_i$, so $A_i(s_i) = 0$. Hence, for finite s, clause 2 will apply iff $A_j(s_j) = 1$ where j is the last element of I in s. Clauses 2 and 3 say that, after s is played, the next move in $\bigotimes_{i\in I} A_i$ belongs to whoever has the next move in A_j after s_j is played. If s is infinite, $i \in I$ appears in s, and s_i is finite, then, by an argument like the one above, $A_i(s_i) = 0$. Thus, the only i's which are relevant in clause 2 are those for which s_i is infinite. By clauses 2 and 3, then, 0 wins the play s of $\bigotimes_{i\in I} A_i$ iff, for some $i \in I$, s_i is infinite and 0 wins the play s_i of A_i.

It is not hard to see that $\bigotimes_{i\in I} A_i$ is a win for 1 iff all the A_i's are. In this respect, $\otimes$ resembles $\bigwedge$. However, it is possible for $\bigotimes_{i\in I} A_i$ to be a win for 0 even if none of the A_i are. We conclude this section by constructing an example of this phenomenon.

Let D be a non-principal ultrafilter on ω. For any strictly increasing sequence $s \in {}^{\omega}\omega$, let

$$E(s) = \{x \in \omega \mid (\exists n \in \omega)\ s(2n-1) \leqslant x < s(2n)\}$$

(where $s(-1)$ means 0). If we think of ω as partitioned into the segments with endpoints $s(n)$, then $E(s)$ is the union of the even-numbered segments. Define a game A on ω by

$$A(s) = \begin{cases} A(\bar{s}(k)), & \text{if } k \text{ is the least number with } s(k) \leqslant s(k-1), \\ \mathrm{length}(s) \text{ modulo } 2, & \text{if there is no such } k \text{ and } s \text{ is finite}, \\ 1, & \text{if there is no such } k,\ s \text{ is infinite, and } E(s) \in D, \\ 0, & \text{otherwise}. \end{cases}$$

The first clause forces the players to produce a strictly increasing sequence s; let us restrict our attention to such plays. By the second clause, the players move alternately, with 0 moving first. If we imagine that, by playing $s(n)$, a player "takes" the integers from $s(n-1)$ to $s(n)-1$ (inclusive), then the last two clauses say that whoever takes almost all (with respect to D) integers wins.

Let i and j be two indices not in ω, and let A_i and A_j both be A. Then $\bigotimes_{k\in\{i,j\}} A_k$ $(= A_i \otimes A_j = A \otimes A)$ is a win for 0 by means of the following strategy. Begin by choosing i and choosing 1 in A_i. From now on, whatever move your opponent makes in A_i or A_j, immediately switch to the other game and make the same move there. If 0 follows this strategy, then $E(s_i)$ and $E(s_j)$ are complements of each other (except that 0 is in both), so one of them is in D, and 0 wins.

An analogous strategy shows that $(\sim A) \otimes (\sim A)$ is also a win for 0. If either A or $\sim A$ were a win for 0, the other would be a win for 1, and its tensor product with itself would also be a win for 1, contrary to what we have shown. So A is undetermined, even though $A \otimes A$ is a win for 0.

2. The strong ordering. The intended meaning of the strong ordering relation $A \leqslant B$ was indicated in the introduction. In terms of the operations introduced in § 1, we have the following reformulation of the informal description. There is a strategy σ for 0 in $A \otimes \sim B$ such that, if 0 uses this strategy, and if the moves of 1 in A are made by an oracle for A, then 0 wins the play of $\sim B$. (The player 0 in this description corresponds to player 1 of B in the introduction.) To eliminate the reference to an oracle, we observe that the only objective criterion for whether the moves of 1 in A were made by an oracle is that 1 wins the play of A if it is finished. The description thus becomes: If 0 uses σ and the play of A is unfinished or won by 1, then 0 wins the play of $\sim B$. Equivalently, σ is a winning strategy for 0 in $A \otimes \sim B$. Therefore, we define

$$A \leqslant B \quad \text{iff} \quad A \otimes \sim B \text{ is a win for } 0 .$$

THEOREM 1. *The relation $\leqslant$ is a pre-ordering (i.e. reflexive and transitive).*

Proof. Reflexivity: The game $A \otimes \sim A$ consists of two games of A. Player 0 acts as 0 in the first but as 1 in the second, and he is free to change games at any of his moves. Suppose he uses the following "mimicking" strategy. Start by choosing the game where your opponent moves first. Whenever it's your move thereafter, switch games, and do in the new game what your opponent just did in the other one. The result of this strategy is that the two plays of A are identical. As 0 plays opposite roles in the two games, he wins one (and loses the other), thereby winning $A \otimes \sim A$.

Transitivity: Let σ and τ be winning strategies for 0 in $A \otimes \sim B$ and $B \otimes \sim C$, respectively, and let 0 play $A \otimes \sim C$ according to the following strategy. Imagine, in addition to the games of A and C actually being played, a fictitious game of B. Begin by playing σ as long as it dictates moves (of 0) in A; your opponent must, of course, reply in A. If σ ever dictates a switch to $\sim B$, make any move dictated by σ (for 1) in the fictitious B, and begin playing τ as long as it dictates moves (of 1) in C. When τ dictates a switch to B, make whatever moves τ dictates (for 0) in the fictitious B. Continue playing B, using σ (resp. τ) to determine 1's (resp. 0's) moves until σ or τ dictates a switch to A or $\sim C$ (respectively); then make the indicated move in the actual game. (Thus, you use σ and τ against your opponent in A and C, and against each other in B.)

If you do not finish and win the play of A (as 0), then, since σ wins $A \otimes \sim B$, the fictitious play of B is finished and won by 1. But then, since τ wins $B \otimes \sim C$, you must finish and win the play of C (as 1). Therefore, this strategy is a winning strategy for 0 in $A \otimes \sim C$. ■

We associate an equivalence relation $\equiv$ with the pre-ordering $\leqslant$ in the usual way:

$$A \equiv B \quad \text{iff} \quad A \leqslant B \text{ and } B \leqslant A .$$

The equivalence classes are partially ordered by (a relation induced by) $\leqslant$.

If f is a bijection from X to Y, then any game A on Y is equivalent to the game B on X defined by

$$B(s) = A(f \circ s) .$$

Player 0 wins $A \otimes \sim B$ and $B \otimes \sim A$ by means of mimicking strategies as in the proof of Theorem 1; the only difference is that, before copying his opponent's move in the other game, he must "translate" it via f. Therefore, when studying properties of games that are invariant under equivalence, we may confine our attention to games on cardinals.

In fact, we may usually confine our attention to games on infinite cardinals. Notice first that games on 0 or 1 are determined, hence usually uninteresting. And any game on ω is equivalent to a game on 2, for the choice of an $n \in \omega$ can be coded as n choices of 1 followed by a 0. We leave to the reader the task of explicitly writing out the game B on 2 corresponding to a game A on ω and proving that $A \equiv B$.

In the next theorem we collect several elementary properties of the strong ordering and the operations defined in § 1.

THEOREM 2.

(a) *$\sim B \leqslant \sim A$ iff $A \leqslant B$.*

(b) *$\bigwedge_{i \in I} A_i$ is a greatest lower bound, and $\bigvee_{i \in I} A_i$ is a least upper bound, of $\{A_i \mid i \in I\}$.*

(c) *If $\{J_i \mid i \in I\}$ is a disjoint partition of K, then $\bigotimes_{i \in I} (\bigotimes_{j \in J_i} A_j) \equiv \bigotimes_{k \in K} A_k$. The same is true with $\bigwedge$ or $\bigvee$ in place of $\otimes$, even if the J_i are not disjoint.*

(d) *If $I \subseteq J$, then $\bigotimes_{j \in J} A_j \leqslant \bigotimes_{i \in I} A_i$. The same is true with $\bigwedge$ in place of $\otimes$, but with $\bigvee$ the inequality is reversed.*

(e) *If $I = \{0\}$, then $\bigwedge_{i \in I} A_i \equiv \bigvee_{i \in I} A_i \equiv \bigotimes_{i \in I} A_i \equiv A_0$.*

(f) *If $A_i \leqslant B_i$ for all $i \in I$, then $\bigotimes_{i \in I} A_i \leqslant \bigotimes_{i \in I} B_i$, and the same is true with $\bigwedge$ or $\bigvee$ in place of $\otimes$, and also with $\equiv$ in place of $\leqslant$.*

(g) *$\underline{O} \leqslant A \leqslant \underline{I}$ for all games A.*

(h) *$A \equiv \underline{O}$ (resp. $A \equiv \underline{I}$) iff A is a win for 0 (resp. 1).*

(i) *$\bigotimes_{i \in I} A_i \leqslant \bigwedge_{i \in I} A_i$.*

The proofs of all these assertions are straightforward and left to the reader. (Note that the assertions about $\bigwedge$ and $\bigvee$ in (c) through (f) all follow from (b) and that (i) follows from (b), (d), and (e).)

We could deduce from Theorem 2 that the equivalence classes form a complete lattice if were not for two set-theoretical difficulties. The lesser of these is that the equivalence classes are not sets. This problem can be overcome by standard devices, such as choosing representatives, and we shall say no more about it. The more serious difficulty is that there are too many equivalence classes. All we can say is that we have a complete "lattice" $\mathcal{S}$, i.e. a partially ordered class in which every subset has a greatest lower bound and a least upper bound.

To illustrate the problems caused by the fact that $\mathcal{S}$ is not a set, we recall that $\bigwedge_{i\in I} A_i$ is a win for 0 iff one of the A_i's is. In view of Theorem 2, this means that the greatest lower bound of a set of non-zero elements of $\mathcal{S}$ is non-zero. In a complete lattice, this property would imply the existence of a unique atom, namely the greatest lower bound of all the non-zero elements. Because of the excessive size of $\mathcal{S}$, this argument is not applicable, and in fact $\mathcal{S}$ has no atoms.

THEOREM 3. *For every undetermined game A, there is a game B incomparable with A.*

Proof. Let A be given, and assume, without loss of generality, that A is on a cardinal $\varkappa$ such that $\varkappa^\omega = 2^\varkappa$. (There are arbitrarily large such $\varkappa$'s, and if A is on $\lambda < \varkappa$ then A is equivalent to a game on $\varkappa$ in which the first player to move outside λ loses.) B will also be on $\varkappa$, and, for $s \in \mathrm{Fin}(\varkappa)$,

$$B(s) = \mathrm{length}(s) \,\mathrm{modulo}\, 2\,,$$

so the players move alternately in B. Thus, $(A\otimes \sim B)(s)$ and $(B\otimes\sim A)(s)$ are already defined for all finite s, and the set of strategies for 0 in $A\otimes\sim B$ and $B\otimes\sim A$ is thereby defined. As both of these games are on sets of cardinality $\varkappa$, there are $2^\varkappa$ such strategies. Enumerate them in a sequence $\{\sigma_\alpha |\ \alpha < 2^\varkappa\}$. To complete the definition of B, we shall define an increasing sequence $\{B_\alpha |\ \alpha < 2^\varkappa\}$ of partial functions from ${}^\omega\varkappa$ into 2, and then let $B\restriction {}^\omega\varkappa$ be any total extension of $\bigcup_{\alpha<2^k} B_\alpha$. The B_α's will be defined by induction on α in such a way that B_α has cardinality at most α. Begin by setting $B_0 = 0$. For limit ordinals α, set $B_\alpha = \bigcup_{\beta<\alpha} B_\beta$. Suppose $\alpha = \beta+1$ and B_β has been defined. Suppose also that σ_β is a strategy for 0 in $B\otimes\sim A$. (The case that σ_β is a strategy in $A\otimes\sim B$ is handled similarly.) We shall define B_α so as to guarantee that σ_β is not winning.

For each $s\in {}^\omega\varkappa$, let $s'(n) = s(2n+1)$, so s' is the sequence of moves of 1 in the play s of B. As B_β has cardinality $\leqslant \alpha < 2^\varkappa$ and ${}^\omega\varkappa$ has cardinality $\varkappa^\omega = 2^\varkappa$, we can find an $f\in {}^\omega\varkappa$ such that $f\neq s'$ for all $s \in \mathrm{Domain}(B_\beta)$. Consider the plays t of $B\otimes\sim A$ that result when 0 uses strategy σ_β, 1 plays $f(n)$ at his nth move in B, but 1 plays arbitrarily in A. For any such t, let t_B and $t_{\sim A}$ be the subsequences of t consisting of the moves in B and in $\sim A$, respectively.

Case 1. For every such t, the play $t_{\sim A}$ of $\sim A$ is finished and won by 0. Then 0 wins $\sim A$ by means of the following strategy. Imagine, in addition to the actual game of $\sim A$, a fictitious game of B; use strategy σ_β, and imagine that your opponent is playing f in B. Thus, $\sim A$ is a win for 0, contrary to the hypothesis that A is undetermined.

Case 2. For one such t, the play $t_{\sim A}$ is unfinished or won by 1. Notice that t_B is not in the domain of B_β, for $t_B' = f \neq s'$ for all s in that domain. Therefore, we may define B_α to be the extension of B_β which also maps t_B to 0. Then the play t of $B\otimes\sim A$ is won by 1 even though 0 used σ_β, so σ_β is not winning.

Since every possible strategy for 0 in $A\otimes\sim B$ and $B\otimes\sim A$ is σ_β for some β, we conclude that B is incomparable with A. ∎

COROLLARY 3a. *Every undetermined game is < another undetermined game. Every undetermined game is the first element of arbitrarily long increasing well-ordered sequences in $\mathcal{S}$.*

Proof. For the first claim, take the least upper bound of the given game and one incomparable with it. (Recall that $\bigvee_{i\in I} A_i$ is undetermined if all the A_i are.) The second claim follows by transfinite iteration of the first, using the completeness of $\mathcal{S}$ at limit stages. ∎

COROLLARY 3b. *Given $2^\varkappa$ or fewer undetermined games on $\varkappa$, there is a game on $\varkappa$ incomparable with them all, provided $\varkappa^\omega = 2^\varkappa$. In particular, for every set of undetermined games, there is a game incomparable with them all.*

The proof is a trivial modification of the proof of Theorem 3.

We define $\mathcal{S}(\varkappa)$ to be the sublattice of $\mathcal{S}$ consisting of the equivalence classes that contain games on $\varkappa$. Unlike $\mathcal{S}$, $\mathcal{S}(\varkappa)$ is a set; its cardinality is $\leqslant 2^{(k^\omega)}$. $\mathcal{S}(\varkappa)$ is not complete, but every subset of power $\leqslant\varkappa$ has a greatest lower bound and a least upper bound.

COROLLARY 3c. *If $2^\varkappa = \varkappa^\omega$, then $\mathcal{S}(\varkappa)$ includes chains of order type $\varkappa^+$ and antichains of cardinality $(2^\varkappa)^+$.*

The first assertion is proved like Corollary 3a, and the second follows from Corollary 3b.

Our next theorem will imply that every partially ordered set can be order-isomorphically embedded in $\mathcal{S}$ and every free lattice can be lattice-isomorphically embedded in $\mathcal{S}$. Let $(P, \leqslant)$ be any partially ordered set. By a *term*, we will mean an expression built up from elements of P by means of the (formal) lattice operations $\wedge$ and $\vee$. Whenever P is mapped into a lattice L, every term denotes an element of L, and every inequality, $S\leqslant T$, between such terms becomes either true or false. We define inductively the notion of a *necessary* inequality between terms.

(a) If $p\leqslant q$ in P, then $p\leqslant q$ is necessary.

(b) If $S\leqslant T$ and $S\leqslant U$ are necessary, so is $S\leqslant T\wedge U$.

(c) If $S\leqslant U$ and $T\leqslant U$ are necessary, so is $S\vee T\leqslant U$.

(d) If $S \leqslant U$ is necessary, so are $S \leqslant T \vee U$ and $S \leqslant U \vee T$.

(e) If $S \leqslant U$ is necessary, so are $S \wedge T \leqslant U$ and $T \wedge S \leqslant U$.

An inequality is necessary only if its being so follows from (a)–(e). If P is mapped into a lattice in an order-preserving way, then all necessary inequalities clearly become true.

THEOREM 4. *Assume P has cardinality $\leqslant 2^{\varkappa} = \varkappa^{\omega}$. Then P can be order-isomorphically embedded in $\mathcal{S}(\varkappa)$ in such a way that no unnecessary inequalities become true.*

Before proving this theorem we point out a few consequences. If P is embedded in $\mathcal{S}(\varkappa)$ in the prescribed way, then the sublattice $\bar{P}$ of $\mathcal{S}(\varkappa)$ generated by P is completely free on P. This means that any order-preserving map of P to any lattice extends uniquely to $\bar{P}$; it follows immediately from the theorem. If we take P to be an antichain, then $\bar{P}$ is a free lattice. Thus, every partially ordered set (resp. free lattice) can be order-(resp. lattice-) isomorphically embedded in $\mathcal{S}$. As another corollary, we obtain the theorem of Whitman [4] that the only inequalities between terms, that hold in the completely free lattice on P, are the necessary ones.

Proof of Theorem 4. We shall define, for each $p \in P$, a game, also called p, on $\varkappa$. In each of these games,

$$p(s) = \text{length}(s) \,\text{modulo}\, 2, \quad \text{if} \quad s \in \text{Fin}(\varkappa)\,.$$

As in the proof of Theorem 3, $p \upharpoonright {}^{\omega}\varkappa$ will be defined by an induction of length $2^{\varkappa}$. At each step, we consider a strategy that threatens to make an unnecessary inequality true, and we make sure that it doesn't work. Also, if $p \leqslant q$ in P, then whenever we define $p(s) = 0$ for some s, we also define $q(s) = 0$ for the same s (at the same stage of the induction), and whenever we define $q(s) = 1$, we also define $p(s) = 1$. This ensures that $p \otimes \sim q$ is a win for 0 (by means of a mimicking strategy) so $p \leqslant q$ in $\mathcal{S}(\varkappa)$. At each stage of the induction, only one or two new sequences s will be added to $\bigcup_{p \in P} \text{Domain}(p)$.

Clearly, there are only $2^{\varkappa}$ terms and therefore only $2^{\varkappa}$ unnecessary inequalities $S \leqslant T$. For each of these, there are only $2^{\varkappa}$ strategies for 0 in the corresponding game $S \otimes \sim T$, so there are enough steps in the induction to make sure that each of these strategies can be defeated.

Suppose we are at a particular stage of the induction, at which the strategy σ for 0 in $S \otimes \sim T$ is under consideration. ($S \leqslant T$ is unnecessary). There are fewer than $2^{\varkappa} = \varkappa^{\omega}$ sequences $s \in \bigcup_{p \in P} \text{Domain}(p)$, so we can choose an $f \in {}^{\omega}\varkappa$ which doesn't occur as the sequence of moves of either 0 or 1 in any such s.

A play of $S \otimes \sim T$ consists of two phases. In phase 1, the players are deciding which of the subterms of S and T to play. For example, if S is $X \wedge Y$, then 0's choice of X or Y belongs to phase 1. These moves continue until S and T have been reduced to atomic terms $p, q \in P$. Then the players are essentially playing $p \otimes \sim q$; this is phase 2. (In fact the phases may overlap. Once S is reduced to p, the players may start to play p before finishing the reduction of T.)

Consider the following play of $S \otimes \sim T$. Player 0 uses strategy σ. At a phase 1 move of player 1, if S and T have been reduced to X and Y with $X \leqslant Y$ unnecessary, he moves so that the resulting reduction still corresponds to an unnecessary inequality. He can do this because of clauses (b) and (c) in the definition of necessary. Furthermore, 0's phase 1 moves cannot produce necessary inequalities from unnecessary ones, by clauses (d) and (e). Thus, the reduced games $X \otimes \sim Y$ correspond to unnecessary inequalities $X \leqslant Y$. At his phase 2 moves, player 1 play the sequence f in each of the components p, q of $p \otimes \sim q$.

Case 1. In this play s, T is not finished. Then S is finished. It is ultimately reduced to some $p \in P$, and the play t of p (subsequence of s) is not in Domain(r) for any $r \in P$, by choice of f. For $p' \geqslant p$, extend p' by setting $p'(t) = 0$; leave the other games unchanged. Then 0 loses the play s of $S \otimes \sim T$ although he used σ, so σ is not winning.

Case 2. S is not finished. This is entirely analogous to Case 1. We take the play t of the $q \in P$ to which T eventually reduced, and set $q'(t) = 1$ for all $q' \leqslant q$.

Case 3. Both S and T are finished in s. They are reduced to games p and q such that $p \leqslant q$ is unnecessary, which implies $p \nleqslant q$ in P by clause (a). The plays t and u of p and q (subsequences of s) are not in Domain(r) for any $r \in P$, by choice of f. If $r \geqslant p$, extend it by setting $r(t) = 0$; if $r \leqslant q$, set $r(u) = 1$; otherwise, do nothing to r. Even if $t = u$, these definitions do not conflict with each other because $p \nleqslant q$.

This completes the inductive definition of the games p. The construction assures that P is mapped into $\mathcal{S}(\varkappa)$ in an order-preserving way, and no unnecessary inequalities hold. In particular, if $p \nleqslant q$ in P, then $p \leqslant q$ is unnecessary (by inspection of the definition of necessary), so $p \nleqslant q$ in $\mathcal{S}(\varkappa)$. The map of P into $\mathcal{S}(\varkappa)$ is therefore order-isomorphic. ∎

COROLLARY. *$\mathcal{S}(\varkappa)$ is not modular.*

Proof. It suffices to consider $\mathcal{S}(\omega)$, as sublattices of modular lattices are modular. Let P be $\{0, 1, 2\}$ ordered so that $0 < 2$ and 1 is incomparable with 0 and 2. Then the modular inequality

$$(0 \vee 1) \wedge 2 \leqslant 0 \vee (1 \wedge 2)$$

is unnecessary, hence false for a certain embedding of P into $\mathcal{S}(\varkappa)$. ∎

With P and $\varkappa$ as in Theorem 4, let us extend the notion of term by allowing formal g.l.b.'s and l.u.b.'s of $\varkappa$ or fewer, rather than only two, terms at a time. We also extend the definition of necessary to include the new infinitary terms. The clauses corresponding to (b) and (e) are

(b′) If $S \leqslant T_i$ is necessary for all $i \in I$, so is $S \leqslant \bigwedge_{i\in I} T_i$,

(e′) If $S_i \leqslant T$ is necessary for some $i \in I$, then so $\bigwedge_{i\in I} S_i \leqslant T$,

and (c) and (d) are modified similarly. (The index set I is assumed to have cardinality $\leqslant \varkappa$.) Theorem 4 remains true with these extended definitions of term and necessary. It follows that $\mathcal{S}$ contains an isomorphic copy of the completely free λ-complete lattice on P, where λ is any cardinal and P is any partially ordered set. It also follows that completely free λ-complete lattices satisfy only necessary inequalities.

It should be pointed out that not every lattice can be embedded in $\mathcal{S}$.

THEOREM 5. *If $\bigwedge_{i\in I} A_i \leqslant \bigvee_{j\in J} B_j$, then either $A_i \leqslant \bigvee_{j\in J} B_j$ for some $i \in I$, or $\bigwedge_{i\in I} A_i \leqslant B_j$ for some $j \in J$.*

Proof. Let σ be a winning strategy for 0 in

$$\bigwedge_{i\in I} A_i \otimes \sim\bigvee_{j\in J} B_j = \bigwedge_{i\in I} A_i \otimes \bigwedge_{j\in J} \sim B_j .$$

Since 0 moves first in a tensor-product game, σ must begin by specifying a choice of $\bigwedge_{i\in I} A_i$ or $\bigwedge_{j\in J} \sim B_j$; suppose it chooses the former. (The argument in the other case is analogous.) By definition of $\bigwedge$, it is still 0's move, and σ must choose an $i \in I$. From here on, the players are, in effect, playing $A_i \otimes \sim\bigvee_{j\in J} B_j$, and σ provides a winning strategy for 0 in this game. Hence, $A_i \leqslant \bigvee_{j\in J} B_j$. ■

It is easy to give examples of lattices where $A_1 \wedge A_2 \leqslant B_1 \vee B_2$ does not imply $A_i \leqslant B_1 \vee B_2$ or $A_1 \wedge A_2 \leqslant B_i$ for either i. According to Theorem 5, such a lattice cannot be lattice-isomorphically embedded in L.

3. The weak ordering. The weak ordering relation, $A \lesssim B$, was informally described in the introduction. It is similar to the strong ordering, but, instead of playing A against the oracle, we play a more complicated game which we call $R(A)$. This game resembles the tensor product of ω copies of A in that player 0 may abandon plays of A to start new ones and may later resume previously abandoned plays. Also, the oracle (as 1) is expected to win every finished play of A, so the rules for winning in $R(A)$ are the same as for $\bigotimes_{i\in\omega} A$. However, $R(A)$ differs from this tensor product in that 1 (the oracle) is required to answer the same moves the same way in all plays.

Turning to a formal definition of $R(A)$, we assume that A is on a set X disjoint from ω. (Otherwise, replace X by another set.) Then $R(A)$ is the game on $\omega \cup X$ such that

$$R(A)(s) = \begin{cases} 1 & \text{if } (\exists k < \text{length}(s))\ (s(k) \in \omega \text{ and } R(A)(\bar{s}(k)) = 1), \\ 1 & \text{if } (\exists i, j, k \in \omega)\ (\overline{s_i}(k) = \overline{s_j}(k) \text{ and } s_i(k) \neq s_j(k) \\ & \qquad \text{and } A(\overline{s_i}(k)) = 1), \\ 1 & \text{if } (\exists i \in \omega \cap \text{Range}(s)) A(s_i) = 1, \\ 0 & \text{otherwise}. \end{cases}$$

(The notation s_i was defined in § 1, before the definition of $\otimes$.) The second clause of this definition says that 1 loses if the ith and jth plays of A first differ at the kth move and this was a move of 1. The remaining clauses are as in the definition of $\bigotimes_{i\in\omega} A$, and the remarks following the definition of $\otimes$ in § 1 also apply here.

THEOREM 6.

(a) $R(A) \leqslant \bigotimes_{i\in\omega} A \leqslant A$.

(b) $R(A) \equiv RR(A)$.

(c) *If $A \leqslant B$, then $R(A) \leqslant R(B)$.*

(d) $R \bigvee_{i\in I} A_i \equiv \bigvee_{i\in I} RA_i$.

(e) $R(A) \otimes R(B) \equiv R(A \otimes B) \equiv R(A \wedge B)$.

Proof. The first inequality in (a) is proved by a mimicking strategy and the second follows from Theorem 2. So $RR(A) \leqslant R(A)$. For (b), we need only show that $R(A) \otimes \sim RR(A)$ is a win for 0. If $f\colon \omega \times \omega \to \omega$ is any bijection, then 0 wins this game by a mimicking strategy that ensures the same play in the (i, j)th copy of A in $RR(A)$ and the $f(i, j)$th copy of A in $R(A)$. For (c), suppose σ is a winning strategy for 0 in $A \otimes \sim B$. Then 0 wins $R(A) \otimes \sim R(B)$ by the following strategy. When player 1 chooses the ith copy of B in $R(B)$, reply, using σ, in that copy of B and the ith copy of A in $R(A)$. For each i, player 0 wins the ith copy of A or of B. Either he wins all the B's or at least one of the A's, so he wins $R(A) \otimes \sim R(B)$. (Note that 0 will reply differently to the same moves in two copies of B, thereby losing $R(B)$, only if 1 has already committed the same error in the corresponding copies of A.) Half of (d) follows from (c). For the other half, 0 wins $R(\bigvee_{i\in I} A_i) \otimes \sim\bigvee_{i\in I} R(A_i)$ by the following strategy. Begin by choosing $R(\bigvee_{i\in I} A_i)$ and index 0. Then 1 must reply, in the 0th copy of $\bigvee_{i\in I} A_i$, by choosing an index $j \in I$. By definition of R, he must choose the same index in all other copies that are ever played in. At your (0's) next move, choose $\sim\bigvee_{i\in I} R(A'_i)$, and there choose the same index j. From this point on, you are, in effect, playing $R(A_j)\otimes\sim R(A_j)$ which you win by mimicking. Finally, part (e) is proved by mimicking; we leave the details to the reader. ■

We define the weak ordering by

$$A \lesssim B \quad \text{iff} \quad R(A) \leqslant B.$$

In view of parts (a) and (b) of Theorem 6, we also have

$$A \lesssim B \quad \text{iff} \quad R(A) \leqslant R(B),$$

from which it is clear $\lesssim$ is a pre-ordering. The associated equivalence relation is

$$A \simeq B \quad \text{iff} \quad A \lesssim B \lesssim A.$$

Clearly, $A \leqslant B$ implies $A \lesssim B$, so each weak equivalence class is a union of (strong) equivalence classes. We have a canonical projection of $\mathcal{S}$ onto the partially ordered class $\mathcal{W}$ of weak equivalence classes of games.

For any set $\{A_i \mid i \in I\}$ of games and any game C, we have

$$\begin{aligned}(\forall i \in I)\, C \lesssim A_i &\Leftrightarrow (\forall i \in I)\, R(C) \leqslant A_i\\ &\Leftrightarrow R(C) \leqslant \bigwedge_{i\in I} A_i\\ &\Leftrightarrow C \lesssim \bigwedge_{i\in I} A_i,\end{aligned}$$

and, using Theorem 6(d),

$$\begin{aligned}(\forall i \in I)\, A_i \lesssim C &\Leftrightarrow (\forall i \in I)\, R(A_i) \leqslant C\\ &\Leftrightarrow R(\bigvee_{i\in I} A_i) \equiv \bigvee_{i\in I} R(A_i) \leqslant C\\ &\Leftrightarrow \bigvee_{i\in I} A_i \lesssim C.\end{aligned}$$

Therefore, the operators $\bigwedge$ and $\bigvee$ give g.l.b.'s and l.u.b.'s in $\mathcal{W}$ as well as in $\mathcal{S}$. $\mathcal{W}$ is a complete "lattice" and the canonical projection from $\mathcal{S}$ to $\mathcal{W}$ is a complete "homomorphism".

THEOREM 7. *$\mathcal{W}$ is Brouwerian; the relative pseudo-complement $B{:}A$ is $\sim(R(A)\otimes\sim B)$.*

Proof. For any games A, B, C, we have

$$\begin{aligned}C\wedge A \lesssim B &\Leftrightarrow R(C)\otimes R(A) = R(C\wedge A) \leqslant B\\ &\Leftrightarrow R(C)\otimes\sim(\sim(R(A)\otimes\sim B)) = R(C)\otimes R(A)\otimes\sim B\\ &\quad\ \text{is a win for } 0\\ &\Leftrightarrow R(C) \leqslant \sim(R(A)\otimes\sim B)\\ &\Leftrightarrow C \lesssim \sim(R(A)\otimes\sim B).\ ■\end{aligned}$$

COROLLARY. *$\mathcal{W}$ is distributive; in fact,*

$$A\wedge\bigvee_{i\in I} B_i \simeq \bigvee_{i\in I}(A\wedge B_i).$$

In particular, $\mathcal{W}$ is modular.

Proof. See [1], pages 147 and 65.

Thus, $\mathcal{W}$ has pleasant properties which $\mathcal{S}$ lacks. On the other hand, it lacks the symmetry property of $\mathcal{S}$ expressed by Theorem 2(a). To see this, consider the game A constructed at the end of § 1. As $A\otimes A$ is a win for 0, $R(A) \equiv \underline{O}$, so $A \simeq \underline{O}$. But $\sim A \not\simeq \underline{I}$; in fact $\sim A \simeq 0$ because $\sim A\otimes\sim A$ is a win for 0.

This example also shows that a game can be $\simeq \underline{O}$ without being $\equiv \underline{O}$. On the other hand, if $A \simeq \underline{I}$ then $A \equiv \underline{I}$, for $R(\underline{I})$ is a win for 1.

$\mathcal{W}$ shares with $\mathcal{S}$ the property that, given any set of elements other than $\underline{O}$ and $\underline{I}$, there is an element incomparable with all of them. We omit the proof because it is essentially the same as for $\mathcal{S}$. We cannot embed free lattices in $\mathcal{W}$, for $\mathcal{W}$ is distributive; the best we can hope for is to embed free distributive lattices. We shall show that this is indeed possible.

Let P be a partially ordered set. We define terms as in § 2 (before Theorem 4), but we consider inequalities $\Gamma \leqslant T$ where T is a term and Γ is a finite set of terms. The definition of necessary is modified as follows.

(a) If $p \leqslant q$ in P and $p \in \Gamma$, then $\Gamma \leqslant q$ is necessary.

(b) If $\Gamma \leqslant T$ and $\Gamma \leqslant U$ are necessary, so is $\Gamma \leqslant T\wedge U$.

(c) If $\{S\}\cup\Gamma \leqslant U$ and $\{T\}\cup\Gamma \leqslant U$ are necessary, so is $\{S\vee T\}\cup\Gamma \leqslant U$.

(d) If $\Gamma \leqslant T$ is necessary, so are $\Gamma \leqslant T\vee U$ and $\Gamma \leqslant U\vee T$.

(e) If $\{S\}\cup\Gamma \leqslant U$ is necessary, so are $\{S\wedge T\}\cup\Gamma \leqslant U$ and $\{T\wedge S\}\cup\Gamma \leqslant U$.

THEOREM 8. *P can be order-isomorphically embedded into $\mathcal{W}$ in such a way that if $S \lesssim T$ in $\mathcal{W}$ then $\{S\} \leqslant T$ is necessary.*

Sketch of proof. As in the proof of Theorem 4, we inductively defeat all strategies that threaten to make $S \lesssim T$ when $\{S\} \leqslant T$ is unnecessary. Such a strategy would be in the game $R(S)\otimes\sim T$ (rather than $S\otimes\sim T$ as in § 2). At phase 1 moves, T and the various copies of S in $R(S)$ are reduced to subgames. At any stage of the game, only finitely many copies of S have been started. We let Γ be the finite set consisting of S together with the subgames to which it has been reduced, and we let Y be the subgame to which T has been reduced. Let player 1 strive to keep Γ and Y such that $\Gamma \leqslant Y$ is unnecessary. He can do this successfully, because of the definition of necessary. In particular, if $p \in P$ eventually appears in Γ and T is eventually reduced to q, then $p \nleqslant q$ in P because of (a). The rest of the proof is just like the proof of Theorem 4. ■

For any embedding of P into a lattice, we interpret the inequality

$\Gamma \leqslant T$ as $\bigwedge_{S \in \Gamma} S \leqslant T$. If the embedding preserves order and if the lattice is distributive, then all necessary inequalities become true. (Distributivity is needed because of clause (c).) Hence, the sublattice $\bar{P}$ of $\mathfrak{W}$ generated by P is the completely free distributive lattice on P. We have thus shown that all free distributive lattices are, up to isomorphism, sublattices of $\mathfrak{W}$. We have also shown that an inequality $S \leqslant T$ holds in the completely free distributive lattice on P iff $\{S\} \leqslant T$ is necessary.

As in § 2, we can generalize Theorem 8 by permitting infinitary, say $\varkappa$-ary, lattice operations. The required extension of the definition of necessary is obvious; for example, (c) is replaced by

(c′) If $\{S_i\} \cup \Gamma \leqslant U$ is necessary for every $i \in I$, then so is $\{\bigvee_{i \in I} S_i\} \cup \Gamma \leqslant U$.

We then have order-isomorphic embeddings of P into $\mathfrak{W}$ such that, if S and T are terms built up from P by applying the lattice operations to $\varkappa$ or fewer terms at a time, then $S \lesssim T$ in $\mathfrak{W}$ only if $\{S\} \leqslant T$ is necessary.

Interpreting inequalities as above, we find that the necessary ones are true for any embedding of P into a $\varkappa^+$-complete lattice satisfying the generalized distributive law

$$(\bigvee_{i \in I} S_i) \wedge T \leqslant \bigvee_{i \in I} (S_i \wedge T).$$

(The converse inequality always holds.) Since $\mathfrak{W}$ is Brouwerian, it satisfies this generalized distributive law, and so do all its $\varkappa^+$-complete sublattices. It follows that, for any P and $\varkappa$, the completely free, $\varkappa^+$-complete, generalized-distributive lattice on P can be $\varkappa^+$-completely embedded into $\mathfrak{W}$, and satisfies $S \leqslant T$ iff $\{S\} \leqslant T$ is necessary.

References

[1] G. Birkhoff, *Lattice theory*, AMS Colloquium Publications 25 (1948).

[2] D. Gale and F. M. Stewart, *Infinite games with perfect information*, Contributions to the Theory of Games, Vol. 2, Annals of Math. Studies 28 (1953).

[3] J. Mycielski, *On the axiom of determinateness*, Fund. Math. 53 (1963), pp. 205–224.

[4] P. Whitman, *Free lattices*, Annals of Math. 42 (1941), pp. 325–330.

Reçu par la Rédaction le 6. 9. 1971

A geometric form of the axiom of choice

by

J. L. Bell (London) and **D. H. Fremlin** (Cambridge)

Consider the following well-known result from the theory of normed linear spaces ([2], p. 80, 4(b)):

(∗) the unit ball of the (continuous) dual of a normed linear space over the reals has an extreme point.

The standard proof of (∗) uses the axiom of choice (AC); thus the implication AC→(∗) can be proved in set theory. In this paper we show that this implication can be reversed, so that (∗) is actually *equivalent* to the axiom of choice. From this we derive various corollaries, for example: the conjunction of the Boolean prime ideal theorem and the Krein-Milman theorem implies the axiom of choice, and the Krein-Milman theorem is not derivable from the Boolean prime ideal theorem.

1. Preliminaries. Throughout this paper we shall assume that all linear spaces we consider have the real number field, $\boldsymbol{R}$, as their underlying field of scalars.

DEFINITION. Let L be a linear topological space. A subset A of L is said to be *quasicompact* if whenever $\mathcal{F}$ is a family of closed convex subsets of L such that $\{F \cap A\colon F \in \mathcal{F}\}$ has the finite intersection property, then $\bigcap \{F \cap A\colon F \in \mathcal{F}\} \neq \emptyset$. An element $a \in A$ is called an *extreme point* of A if $x, y \in A$ and $a = \frac{1}{2}(x+y)$ imply $a = x = y$.

Now consider the following propositions:

(BPI) Every Boolean algebra contains a prime ideal.

(HB) Let M be a linear subspace of a linear space L and let p be a sublinear functional on L (that is, $p(x+y) \leqslant p(x)+p(y)$ for all $x, y \in L$ and $p(ax) = ap(x)$ for all $0 \leqslant a \in R$ and all $x \in L$). If f is a linear functional on M such that $f(x) \leqslant p(x)$ for all $x \in M$, then f can be extended to a linear functional g on L such that $g(x) \leqslant p(x)$ for all $x \in L$.