Finite and Infinite Dialogues

"But how does he know where and how
he is to look up the word 'red' and
what he is to do with the word 'five'?"
Well, I assume he *acts* as I have described.
Explanations come to an end somewhere.
`Ludwig Wittgenstein`

Philosophical Investigations I.1

Two players Ann and Bob are told that the following will happen. Some positive integer n will be chosen and *one* of n, $n+1$ will be written on Ann's forehead, the other on Bob's. Each will be able to see the other's forehead, but not his/her own. After this is done, they will be asked repeatedly, beginning with Ann, if they know what their own number is.

Theorem 1: In those cases where Ann has the even number, the reponse at the nth stage will be, "my number is $n+1$", and in the other cases, the response at the $(n+1)$st stage will be "my number is $n+1$". In either case, it will be the person who *sees* the smaller number, who will respond first.

Definition 1: A *Kripke model* M for a (two person) knowledge situation consists of a state space W and two equivalence relations $\equiv_1$ and $\equiv_2$. Intuitively $s \equiv_1 t$ means that states s and t are indistinguishable to player 1 (Ann) and $s \equiv_2 t$ means that they are indistinguishable to player 2 (Bob). We shall assume in this paper that W is finite or countable.

In the example we are looking at, $W = \{(m,n)|m, n \epsilon N^+$ and $|m-n| = 1\}$. If $s, t \epsilon W$ and $i \epsilon \{1,2\}$, then $s \equiv_i t$ iff $(s)_j = (t)_j$, where $j = 3 - i$, and $(s)_j$ is the j-the component of s. Intuitively, $s \equiv_i t$ means that when the dialogue begins, player i cannot distinguish between s and t, where Ann is player 1 and Bob is player 2.

Definition 2: A subset X of W is *i-closed* if $s\epsilon X$ and $s \equiv_i t$ imply that $t\epsilon X$. X is *closed* if it is both 1-closed and 2-closed.

Definition 3: Given Kripke model M, $X \subseteq W$, and $s\epsilon X$, then i *knows* X at s iff for all t, $s \equiv_i t$ implies that $t\epsilon X$. X is *common knowledge* at s iff there is a closed set Y such that $s\epsilon Y \subseteq X$.

Observation: If an announcement of a formula ϕ is made, then the new Kripke structure is obtained by deleting all states $s \in W$ where ϕ is false.

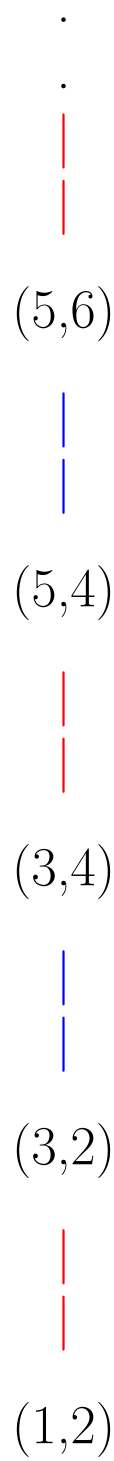
(5,6)
(5,4)
(3,4)
(3,2)
(1,2)

.

.

.

(5,6)

(5,4)

(3,4)

(3,2)

~~(1,2)~~

.

.

.

(5,6)

(5,4)

(3,4)

~~(3,2)~~

~~(1,2)~~

.

.

.

(5,6)

(5,4)

~~(3,4)~~

~~(3,2)~~

~~(1,2)~~

.

.

.

(5,6)

~~(5,4)~~

~~(3,4)~~

~~(3,2)~~

~~(1,2)~~

However, there is a serious defect in the argument in that both Ann and Bob's reasoning depends heavily on what the other one is thinking, including a consideration of what the other does not know. Ann's reasoning is justified if Bob thinks as she believes he does, and Bob's reasoning is justified if she thinks as he believes she does. But there is no guarantee that they do indeed think this way. How do we justify what each thinks and what each does and does not know?

Definition 4: An IDS (interactive discovery system) for M is a map
$f : W \times N^+ \longrightarrow \{\text{``}no\text{''}\} \cup W$ such that for each odd n, $f(s, n)$ (Ann's response at stage n) depends only on the $\equiv_1$ equivalence class of s and on $f(s, m)$ for $m < n$. For each even n, $f(s, n)$ depends only on the $\equiv_2$ equivalence class of s and on $f(s, m)$ for $m < n$.

Definition 5: The IDS f is *sound* if for all s, if $f(s,n) \neq$ "*no*", then $f(s,n) = s$. We define $i_f(s) = \mu_n(f(s,n) \neq$ "*no*") and $p(s) = 1$ if $i_f(s)$ is odd and 2 if $i_f(s)$ is even. (Here μ stands for "least". $i_f(s) = \infty$ if $f(s,n)$ is always "no". We may drop the subscript f from i_f if it is clear from the context.)

Lemma 1: Let f be a sound IDS. Let $s \equiv_i t$, $i(s) = k < \infty$ and $p(s) = i$. Then $i(t) < k$ and $p(t) \neq i$.

Proof: At stage $i(s)$, i has evidence distinguishing between s and t. Since all previous utterances associated with s were "no", some previous utterance associated with t must have been nontrivial. Formally, $f(s,i(s)) = s \neq f(t,i(s))$. But $s \equiv_i t$. Hence $(\exists m < i(s))(f(s,m) \neq f(t,m))$. Since $m < i(s)$, $f(s,m) =$ "*no*" and so $f(t,m) \neq$ "*no*". Thus $i(t) \leq m < i(s)$. Now, if $p(t) = i$, then, by a symmetric argument, we could prove also that $i(t) < i(s)$. But this is absurd. Hence $p(t) \neq i$. $\square$

Corollary: Suppose that $p(s) = i$ and there is a chain $s = s_1 \equiv_1 s_2 \equiv_2 s_3 \equiv_1 ... s_m$. Then $i(s) \geq m$.

Corollary: Suppose that there is a chain $s_1 \equiv_1 s_2 \equiv_2 s_3 \equiv_1 ... s_m \equiv_2 s_1$, with $m > 1$. Then $i(s_i) = \infty$ for all i.

Proof: If, say, $i(s_1) = k < \infty$, we would get $i(s_1) > i(s_2) > ... > i(s_m) > i(s_1)$, a contradiction. $\square$

Remark 1: Theorem 1 is really a proof that the IDS f is sound where f is defined by:
Ann's strategy: If you see 2n+1, then say n "no"'s and then, if Bob has not said his number, say "2n+2". If you see 2n, then say n "no"'s and if Bob has not said his number, say "2n+1".
Bob's strategy: If you see 2n+1, then say n "no"'s and then, if Ann has not said her number, say "2n+2". If you see 2n, then say n "no"'s and if Ann has not said her number, say "2n+1".

These strategies yield: $i(2n+2, 2n+1) = 2n+1$, $i(2n, 2n+1) = 2n$, $i(2n+1, 2n+2) = 2n+2$ and $i(2n+1, 2n) = 2n+1$. In other words, the smaller number if Ann's number is even, and the bigger number if it is odd. These strategies are *optimal*. E.g. we have

$$(6,5) \equiv_1 (4,5) \equiv_2 (4,3) \equiv_1 (2,3) \equiv_2 (2,1)$$

and hence $i(6,5)$ has a minimum value of 5, the value achieved by the strategy above.

Theorem 2: The strategies implicit in theorem 1 and described in remark 1 are optimal. I.e. if h is any other sound IDS, then $i_f(s) \leq i_h(s)$ for all s.

Proof: By cases. Suppose, for example, that Ann has an even number and $s = (2n, 2n-1)$. $i_f(s) = 2n-1$. Suppose Bob is the one who first notices the state. Then we have $(2n, 2n-1) \equiv_2 (2n, 2n+1) \equiv_1 (2n+2, 2n+3)$..., and by lemma 1, $i_h(s)$ could not be finite. So Ann *does* first discover s. But then we have $(2n, 2n-1) \equiv_1 (2n-2, 2n-1) \equiv_2 (2n-2, 2n-3) ... \equiv_2 (2, 1)$ and so, by lemma 1, $i_h(s) \geq 2n-1$. $\square$

Infinite Dialogues

Instead of using the function $f(n) = n + 1$ we use a somewhat more interesting function g defined as follows:

$g(n) = 1$ if $n = 2^k$ for some $k > 0$
$g(n) = n + 2$ if n is odd
$g(n) = n - 2$ otherwise, i.e. if n is even, not a power of 2.

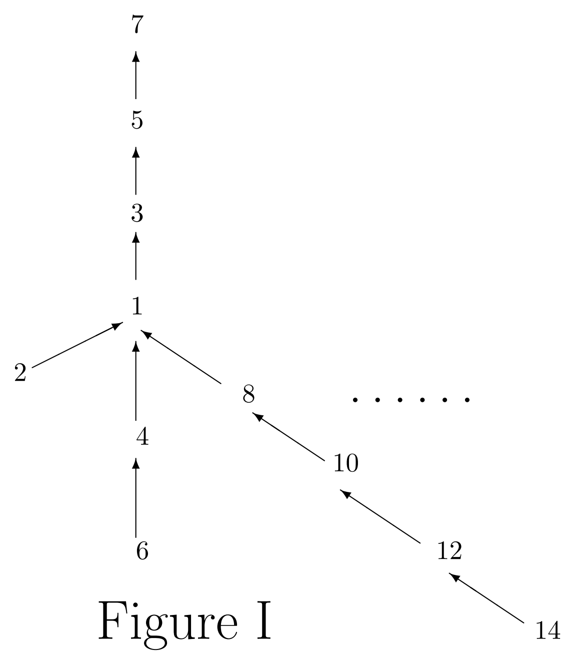

Figure I

(the dots represent numbers not shown, like 16,18,...,32,... etc.)

Again the game proceeds by picking a positive integer n, and writing one of $n, g(n)$ on Ann's forehead, the other on Bob's. Figure II shows states (a, b), where a is written on Ann's forehead and b on Bob's and either $g(a) = b$ or $g(b) = a$.

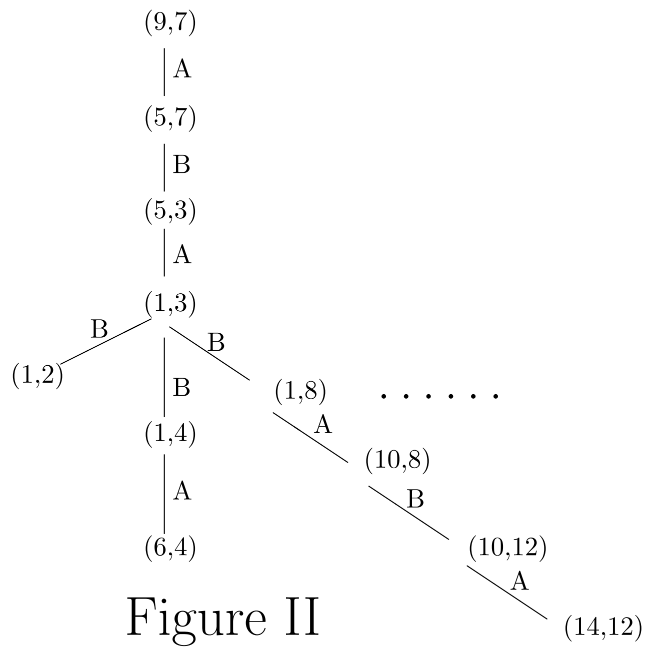

Figure II

Consider now what happens if the state is (1,3). Bob realises after Ann's first "I don't know", that his number is not 2, for otherwise Ann would have known that her number is 1. After her second "I don't know", he realises that his own number is not 4, for otherwise she would have guessed her own number. More generally, after $2^{k-1}+1$ stages, he realises that his number is not 2^k.

Thus when ω stages pass, and Ann has *still* not guessed her own number, Bob will realise that his number is not any power of 2, and hence it must be 3. Thus, in the case of the state (1,3), it is at stage $\omega+1$ that one of the two players realises his number. We can easily see now that if the state is $(5,3)$, then Ann will realise her own number at stage $\omega+2$, and so on through all ordinals of the form $\omega+n$.

This construction is quite similar to that in the Cantor-Bendixson theorem, [Mo], where a closed set is gradually diminished by removing isolated points until, at some countable ordinal, either nothing is left or else a perfect set is left. We now show that the parallel is exact except that we are dealing simultaneously with two topologies on the same space.

The Cantor-Bendixson Theorem

Let X be a subset of the Euclidean space E^n and $p \in X$. Then p is *isolated* if there is a neighbourhood U of p which contains no points of X except p.

Theorem: Let X be a closed subset of E^n and X' be the subset of X (its derivative) obtained by removing all isolated points. X' may have new isolated points if all their neighbours have been removed. Let X'' be the derivative of X' and let X^ω be the limit for all finite stages. Continue this process, then after a countable number of steps, there are no more isolated points. The limit X^∞ may be either empty, or else a perfect set (a closed set which is dense in itself).

Fact: Every perfect set has cardinality that of the continuum.

Corollary: Every closed subset of E^n is either countable (or finite) or has cardinality that of the continuum.

In other words, the continuum hypothesis holds for closed sets.

Definition 6: Let $\mathcal{O}$ be the set of countable ordinals, M a Kripke structure. A $TIDS$ (transfinite interactive discovery system) for M is a pair of maps
$p : \mathcal{O} \to \{1, 2\}$ and $f : W \times \mathcal{O} \to \{\text{``}no\text{''}\} \cup W$ such that for each s, α, If $j = p(\alpha)$, then $f(s, \alpha)$ depends only on the $\equiv_j$ equivalence class of s and on $f(s, \beta)$ for $\beta < \alpha$. Intuitively, $p(\alpha)$ is the person who responds at stage α and $f(s, \alpha)$ is his response at stage α. Again, "no" stands for "I don't know".

Definition 7: The TIDS f, p is *sound* if for all s, α, if $f(s, \alpha) \neq \text{``}no\text{''}$, then $f(s, \alpha) = s$.

We again define $i_f(s) = \mu_\alpha(s(\alpha) \neq \text{``}no\text{''})$. Again, $i_f(s) = \infty$ if $f(s, \alpha)$ is always "no". We think of ∞ as larger than *all* the ordinals, even the infinite ones. By abuse of language, we will write $p(s)$ for $p(i(s))$. This makes our usage consistent with that of the previous section.

First define:
$W_0 = W$, $\mathcal{T}_{i,0} = \mathcal{T}_i$, where the topologies $\mathcal{T}_i$ were defined in definition 2.
$W_{\alpha+1} = W_\alpha -$ the i-isolated points of W_α, where $i = p(\alpha)$.
$\mathcal{T}_{i,\alpha+1} = \mathcal{T}_{i,\alpha}$
$\mathcal{T}_{j,\alpha+1} = \mathcal{T}_{j,\alpha} \oplus W_{\alpha+1} = \{X \cap W_{\alpha+1} | X \epsilon \mathcal{T}_{j,\alpha}\}$ for $j \neq i$
If λ is a limit ordinal, then
$W_\lambda = \bigcap_{\alpha<\lambda} W_\alpha$
$\mathcal{T}_{i,\lambda} = \{X \cap W_\lambda | \exists \alpha < \lambda, X \epsilon \mathcal{T}_{i,\alpha}\}$.

Note that the i-isolated points are not j-isolated for $j \neq i$. Thus, in general, $W_{\alpha+1}$ has to be *added* to j's topology. E.g. in figure II, the point (6,4) is an isolated point for Bob but not for Ann. When that point is removed, Ann gets more sets in her topology.

Now define the functions p, f by: $p(\alpha) = 1$ if α is even and 2 if α is odd. (We think of Ann as beginning with the first ordinal, 0, and re-starting the dialogue at each limit ordinal. Thus for instance, she responds at ω, an even ordinal.) Let the function f be given by: at stage α, if s is an i-isolated point of W_α and i$=p(\alpha)$ then answer s. If the answer s has ever been given, then answer s. Otherwise answer "no". We show now that this is a sound and optimal strategy for all Kripke structures M_g arising from *some* function g from N^+ to N^+.

Theorem 3: f is an optimal (among all strategies which question Ann at all even ordinals and Bob at all odd ordinals.) sound strategy and yields,
$i(s) = i_f(s) = \mu_\alpha(s \epsilon W_\alpha - W_{\alpha+1})$.

Proof: f is evidently sound if it is a strategy. To see that it *is* a strategy, suppose, if possible, that there exist s, t, α such that $s \equiv_i t$ where $i = p(\alpha)$ and $f(s, \beta) = f(t, \beta)$ for all $\beta < \alpha$, but $f(s, \alpha) \neq f(t, \alpha)$. We may assume that α is the smallest ordinal for which this happens, so that $f(s, \beta) = f(t, \beta) =$ "*no*" for all $\beta < \alpha$. Obviously, one (and by soundness exactly one) of $f(s, \alpha), f(t, \alpha)$, say the first, is different from "no". Now $s, t \epsilon W_\alpha$ (since all previous answers were "no") but s is an i-isolated point of W_α. This contradicts the fact that $s \equiv_i t$.

Suppose now that h is another sound strategy and, there is some s such that $i_h(s) = \alpha < i_f(s)$. I.e., h yields knowledge earlier in some case. Assume α is the smallest ordinal for which h is faster than f. Let $i = p(\alpha)$. Now we have $h(s, \beta) = f(s, \beta) =$ "*no*" for all $\beta < \alpha$ and $f(s, \alpha) =$ "*no*", but $h(s, \alpha) = s$. Since $f(s, \alpha) =$ "*no*", s is not an i-isolated point of W_α. Pick $t \neq s$ such that $s \equiv_i t$ and $t \epsilon W_\alpha$. Then t is not an i-isolated point of W_α, and hence of W_β for any $\beta < \alpha$. Thus we have $f(t, \beta) =$ "*no*" for all $\beta < \alpha$ and by minimality of α, $h(t, \beta) =$ "*no*" for all $\beta < \alpha$. Since h is a strategy, this yields $h(t, \alpha) = f(s, \alpha) = s$. Thus h is not sound. $\Box$

Let us consider the problem now over a general Kripke structure with a countable W. Let $W_\infty = \bigcap W_\alpha : \alpha \epsilon \mathcal{O}$.

Definition 8: $< W, \mathcal{T}_1, \mathcal{T}_2 >$ is *scattered* if $W_\infty = \emptyset$.

Theorem 4: $< W, \mathcal{T}_1, \mathcal{T}_2 >$ is scattered iff there is a sound strategy for M which *always* yields a non-trivial answer.

Proof: If $< W, \mathcal{T}_1, \mathcal{T}_2 >$ is scattered, then the CB strategy always yields an answer. If it is not scattered, then clearly the CB strategy cannot always yield an answer. For there is a perfect core (W_∞) which is never removed. However, the CB strategy is optimal. Hence no sound strategy can yield an answer in all cases. $\Box$.

Definition 9: g is *well founded* if there is no infinite chain $x_1, x_2, \ldots$ such that $g(x_{n+1}) = x_n$ for all n. g is *finite-one* iff for all n the set $g^{-1}(n) = \{m|g(m) = n\}$ is finite.

Some of the folowing results will depend on the assumption that $g(n) = n$ or $g(g(n)) = n$ never holds and we make this a **blanket assumption** from now on. The reason this condition is relevant is that if $g(g(n)) = n$ or $g(n) = n$, then the point $(n, g(n))$ might be isolated *even though* g is not well founded.

Theorem 5: (a) The space $< W, \mathcal{T}_1, \mathcal{T}_2 >$ arising from g is scattered iff g is well founded.
(b) If g is well founded and finite-one, then $W_\omega = \emptyset$, i.e. every state is learned at some finite stage.
Proof: The first part has been proved already. To see the second part, notice that König's lemma applies to the tree of g so that every state has only finitely many states under it. $\Box$

Corollary: g is well founded iff the dialogue between Ann and Bob is guaranteed to terminate (with the CB strategy).

We remark that for computable well founded functions g, all ordinals less than Church-Kleene ω_1 can arise as ordinals of the corresponding trees.

The Probabilistic Case

We now show that if we are dealing with *justified risk* rather than knowledge, then the situation of the last section, which required infinite dialogues, improves dramatically.

Suppose that the number n is chosen in accordance with some probability distribution, say $\mu_1(n) = \frac{1}{n(n+1)}$. Thus $\mu_1(1) = 1/2$, $\mu_1(2) = 1/6$, $\mu_1(3) = 1/12$ etc. This μ_1 induces a probability measure μ on W if we assume that the states (a, b) and (b, a) are equally likely.

Now the game is played as follows: each person risks \$1,000 by saying "I know my number, it is ...". If (s)he is right, (s)he receives one dollar. If (s)he is wrong, (s)he loses \$1,000. It is assumed that the parties are rational and that rationality is common knowledge. Thus, for example, if Ann did not guess her number yet, Bob can assume that it was not yet profitable for her, and conversely.

Then it will always make sense to take the risk after a *finite* number of steps. I.e. after a finite number of stages, the expected payoff will be positive for some person.

Theorem 6: If some function g is well founded, μ is a probability distribution such that $\mu(s)$ is positive for all s, B is some bet with positive payoff for a correct guess, and negative payoff for an incorrect guess, and it is common knowledge that the parties are rational, then after a finite number of rounds, someone will take the risk (and will be justified in taking the risk).

Proof: If not, then there is some x of lowest rank in the tree of g such that the bet is never profitable for either side. The person who sees x knows that his number is either $g(x)$ or else in $X = \{y|g(y) = x\}$. However, since x has the lowest possible rank as above, all these y, being of lower rank, are finitely bettable, i.e. it is justified to bet on them at some finite stage. Hence, as time passes, as elements of X which *should* have been guessed are *not* guessed, the set X steadily approaches the empty set and its probability approaches 0. Hence after some finite stage, its probability will be as small as needed. At this point it *will* make sense for Bob to take the risk. This contradiction proves the theorem. $\square$

Definition 10: Let M be a Kripke structure, μ be a probability measure on W and ϵ be a real number > 0. An interactive discovery system f for M, μ is ϵ-*good* if for all s, there is an n such that $f(s, n) = s$, and if n is the least such, then $\mu(\{s\})/\mu(\{t|f(t, n) = s\}) > 1 - \epsilon$.

Theorem 7: Let M be a Kripke structure arising from a well founded computable g. Suppose that μ_1 is a computable probability measure on N^+ and $\delta > 0$. Then there is a δ-good, computable strategy f for M, μ.

Proof: Let d be an integer such that $1/d < \delta$. Define strategies $h_A(s)$, $h_B(s)$ as follows:
$h_A(s)$: Let $n = (s)_2$. Let k be the least integer greater than $\frac{2d}{\mu_1(n)}$.
Let $X = \{m|m < r(k) \text{ and } g(m) = n\}$.
Then $h_A(s) = 1 + max(h_B(m) : m\epsilon X)$; $h_A(s) = 1$ if X is empty.
$h_B(s)$: Let $n = (s)_1$. Let k be the least integer greater than $\frac{2d}{\mu_1(n))}$.
Let $Y = \{m|m < r(k) \text{ and } g(m) = n\}$.
Then $h_B(s) = 1 + max(h_A(m) : m\epsilon Y)$; $h_B(s) = 2$ if Y is empty.

We claim first that this gives us computable functions h_A, h_B. The claim follows from the fact that $h_A(s)$ depends only on $(s)_2$ and on $h_B(m)$ for m such that $g(m) = (s)_2$. Similarly for h_B. Since g is well founded, this is a legitimate definition by reursion.

We now combine h_A, h_B into a strategy f. If n is odd,

$n \geq h_A(s)$ and all previous values $f(s,p)$ have been trivial, then $f(s,n) = (g((s)_2),(s)_2)$. If some previous value has been t then $f(s,n) = t$. Otherwise $f(s,n) =$ "no". Similarly with n even, using h_B instead of h_A.

It is easiy seen that h_A depends only on information that Ann has, and h_B depnds only on information that Bob has. Hence f is a trategy.

We now show that this strategy is $(1/d)$-good, this will imply that it is δ-good. Given s, let n be the least integer such that $g(s,n) \neq$ "no". Assume without loss of generality that n is odd.

If X is empty, then the set $\{m|g(m) = (s)_2\}$ is contained in the set $\{m|m > r(k)\}$ and hence has measure less than $\mu_1(g((s)_2))/d$. Thus the probability that $(s)_1 = g((s)_2)$ is larger than $1 - 1/d$.

If X is not empty, then $n = h_A(s)$. Suppose $(s)_1$ were such that $g((s)_1) = (s)_2$, then if $(s)_1 \epsilon X$, we would already have a non-trivial value earlier. Hence, the probability that $g((s)_1) = (s)_2$, given that there have been only trivial answers so far, is less than $\mu_1(g((s)_2) \times (1/d)$. Hence the probability that the state is $((g(s)_2), (s)_2)$ exceeds $1 - (1/d)$. $\Box$

Theorem 8: g is well founded iff for all μ, δ, there exist δ-good strategies.