MECHANICAL CHESS PLAYER

—W. ROSS ASHBY

THE question I want to discuss is whether a mechanical chess player can outplay its designer (1). I don't say "beat" its designer; I say "outplay." I want to set aside all mechanical brains that beat their designer by sheer brute power of analysis. If the designer is a mediocre player, who can see only three moves ahead, let the machine be restricted until it, too, can see only three moves ahead. I want to consider the machine that wins by developing a deeper strategy than its designer can provide. Let us assume that the machine cannot analyze the position right out and that it must make judgments. The problem, then, becomes that the machine must form its own criteria for judgment, and, if it is to beat its designer, it must form better judgments than the designer can put into it. Is this possible? Can we build such a machine? The problem that faces the designer is the same as that of the father who is not a good chess player and who wants his son to become world champion. Obviously, he must be very careful about what he teaches the boy. If he teaches him rules like: "Always get your queen in the middle of the board as quickly as you can," he may do permanent injury to the child's chess-playing powers.

The problem, then, is how is the machine to develop better criteria of judgment than the designer himself can produce. We can get a line of argument by considering a chess position—I haven't one available, but probably you all can supply one out of your memories—that looks fairly ordinary and yet, in fact, has a powerful move possible. How are we to get the machine to play that move? Suppose the move is so subtle that not even our best players can see it; how are we to get that move played at all? I say that there are only two classes of players that are capable of making that move. One is the beginner, who is so bad that he can make any silly move, and the other is a random player that just draws its moves out of a hat. The one player who can never make that move is the mediocre player: he has his rules of thumb for playing and they are not good enough.

Wiesner: I don't think that necessarily follows. He may be so mediocre that he is playing a random game, too.

Ashby: In that case, he is in the other class.

Wiesner: The fact that a man is a mediocre player doesn't necessarily mean that his rules exclude his making a superhuman move.

Ashby: Not rigorously, but the tendency for the mediocre player is to reject the very good moves and go for the mediocre moves, not necessarily, of course, but usually our first deduction, then, is that if the designer wants to get moves that are better than he can provide himself, he must go for them to a random source of moves.

The next principle is that the machine, as it faces a position, must form a great variety of transformations from the position, must form a valuation from each transformation, must follow the line of action that these valuations suggest, and then, when the game is over, must go back to modify the transformations and valuations simply according to whether the game was won or lost. The transformations and valuations can be formed, at first, entirely at random. A Geiger counter suitably worked up will provide variety and may lead to such rules as: always keep rooks as far apart as possible, or, always keep one bishop on a white square and one on a black.

Pitts: You can't do otherwise.

Bateson: They stay that way anyway.

Ashby: That was a good random one, then. The result is, of course, that the first games will be silly. But if there is corrective feedback that is operated by results, such a machine, breaking up the transformations when the game is lost and holding them when it is won, will inevitably move its population of transformations from the completely random toward those transformations that are the right ones for winning the game.

Bateson: The noisegenerator on your machine is now in a different relationship to the organism and the environment from what it was when you were using the machine to illustrate homeostasis.

Ashby: Yes. It is wanted simply to provide variety. It could be a Geiger counter, providing a stream of irregularly varying numbers. If this machine is a determinate machine, it cannot make randomness out of nothing; it will demand a specific instruction. If you say "Watch the Geiger counter, take the last three numbers on it, and form your transformation in that way," you are giving it specific instructions that it can follow; but because it is getting its instructions off a Geiger counter, you are getting possibilities which are not limited by the limitations of the designer. If the designer said, "I have provided you with a great number of good transformations; select the best," the machine is restricted by the very best that he can produce. But if the transformations take Brownian movements as their source, theoretically, they have no limitations at all. However complicated and subtle the really good transformations are, Brownian movement can provide them. I suggest that something like Brownian movement is the only place where they can be found. They can't be provided by hypothesis.

Bigelow: Why not?

Ashby: That is my basic hypothesis, that our intelligence goes so far and then stops, and that we cannot provide these transformations. We are trying to find something better than what we can provide from our own skill.

Wiesner: You can provide all the elements of the transformation.

Ashby: Yes, but then we have still got to provide some instructions for their combination. If we try to construct it in detail as a determinate machine, it will be limited by our ideas. With Brownian movement, we can just let it go.

Bigelow: It is not at all clear that the addition of the Brownian movement adds one iota of information to the system.

Wiesner: Suppose you tell the machine to try all the possible combinations. One main difficulty is that it would come out no better, as we said previously, than the Brownian movement.

Bigelow: Exactly.

Wiesner: If you have a stack of cards and you shuffle through to find something, without knowing anything about the order, it doesn't matter much whether you do it in a systematic way, if there are a fair number of operations to perform, or do it randomly, provided you examine each thing only once. If you inject the Brownian motion, you run the possibility of sometimes taking longer because you do certain operations more often. Perhaps occasionally you will come out better, but on the average, I think you will come out exactly the same.

Bigelow: Exactly so. Furthermore, I see no possible way to distinguish between the analysis of the situation and the formulation of a strategy in a game like this. I think if you put any limitations on the ability of either machine or human operator to analyze, you put an exact equivalent limitation on the ability to form strategy. To that extent, the problem is closed as soon as you state it. It has little further interest. If you limit the ability of the person or the machine to analyzing three moves ahead, then you put an absolute limit on the variety of strategies they can choose.

Ashby: I am suggesting just the opposite, that this random method can get past the limit.

Bigelow: I thought you started out by saying that you chose to discuss those games on which this limitation is imposed.

Hutchinson: Your idea is that some games would be played in which the play could not be known more than three moves ahead to be good and therefore could not be chosen on any reasonable criterion. Such good moves could, however, come in from the noise-producing mechanism and so will happen to be played. There is then a selector device so that when this happens, memory comes in, and the process is of some use in the future?

Ashby: Yes.

Hutchinson: So the real criterion is what the opponent can learn.

Ashby: Yes.

Bigelow: It is impossible to separate the function of analysis—and by analysis, I mean computing the possible things that can happen—from the operation of formulating the best strategy for a move at any one point. They are one and the same.

Ashby: I deny that.

Bigelow: Well, let's take it up.

Ashby: I went through Capablanca's *Chess Fundamentals* (2) the other day, and found many sentences each of which gave clear advice in a general way without making any specific analysis on specific squares. One example I will quote because it did in fact go beyond my standard of play. When the game has reached a point at which there are only pawns and bishops left, the beginner always takes his pawns off the squares that the enemy bishop covers. Capablanca, however, advises that the pawns should be moved on to the squares the enemy bishop covers to restrict the enemy bishop's movements, regarding that as more important than the mere safety of the pawns. His statement has nothing to do with analysis in the sense of following out in detail the exact position on the board.

Bigelow: But this is a statistical observation about what beginners do. This is an analysis.

Ashby: Let me define what I mean by "analysis." I mean the actual working out on 64 squares that if this bishop moves to that square, it will attack that knight, let this piece in, and so on. By "analysis," I understand specific reference of the actual pieces to the actual squares on the board.

Bigelow: But precisely the statement you made is the summation of experience with such restricted types of analysis as that. It is a statistic of how people behave in playing this game. The question as to whether or not a computing machine can outplay its maker really contains the question as to whether or not it can gather statistics on the probabilities of the situation at each move more rapidly than its maker can, so as to ever exceed the amount of information about the probabilities that its maker has. This is the question, in my language, that you are trying to discuss.

Ashby: Yes, I agree.

Hutchinson: May I make my few remarks before I have to leave, Mr. Chairman?

McCulloch: Please do.