Poker as War: Reducible & Quantum Games, Randomness and Free Will?

by Brian Space
Two Plus Two Magazine, Vol. 15, No. 9

Start with War. War, the card game where half the shuffled deck is dealt to two players. Next, each player exposes a card; the high card wins and collects both cards. Ties are broken by some iteration of the initial process. After the first set of trials, the now possibly uneven player decks are reconstituted in the order the tricks were won; the next round proceeds identically until someone wins the entire deck. The competitive simplicity and randomness of the game make it a common children's card game. My young son Quincy enjoys playing and revels in the delight as incomplete information is revealed in victory or defeat. War is a simple example of intermittent reinforcement hidden in fluctuations that casinos use to great reward.

For our purposes, focus on War as a strategically reducible game. War maps exactly onto the process of flipping a fair coin, where the probability of heads and tails is 50% for a given toss. The starting deck is shuffled to access one of the 52! ≈ 8.x1067 possible arrangements of a standard deck of cards. Incidentally, that number of arrangements means that for a well shuffled deck, one will never see even a small subset of possible starting card arrangements. The combinatorics of having 52 choices for the first card, 51 for the second and so on, leads to the factorial mathematics giving 52! = 52x51x50… 5x4x3x2 distinct starting decks. Digressing further, it has been shown that this vast space of deck configurations can be accessed with only a dozen standard dovetail shuffles and the cards are well randomized after only 4-5 shuffles. (Ann. Appl. Probab., Volume 2, Number 2 (1992), 294-313) Here, random means only that each starting configuration of cards is equally likely to appear in a string of decks produced by the shuffling procedure. This manifestation of randomness can be referred to as equal a priori probabilities, with each shuffled deck appearing with the same likelihood.

Consider, in War the starting arrangement of each player's deck determines a priori who will win. There are no strategic decisions in the game; it is a pure exploration of randomness and variance. The game reduces to flipping a coin, but the card dynamics draw out the determination of your ultimate fate. A perfectly played game of chess may be the same, depending on who moves first. But chess is different in that the dynamics are complex. There are so many tactics that we are not certain of which side is initially advantaged, only that there is an inherent asymmetry. In chess, an enormous number of imperfect games can reasonably be played, and many strategic choices arise.

In contrast, the nature of War can be understood by considering the simpler model game of tossing a coin. For example, the probability of each player winning is 50% and the game is symmetric to choosing heads or tails. In War, the homologous choice is which initial half deck of cards a player chooses. Both games terminate with a deterministic winner and the player has no choice in the matter. A player's fate is sealed the moment the initial choice of deck is made or when the coin is tossed in the air. The games are isomorphic, and conclusions can be drawn from the simpler model coin game without the confounding details of the War card dynamics.

However, to avoid what he perceives as a doomed fate in War, my clever son Quincy will take to shuffling his cards to alter his luck after a round that doesn't go his way. Does this change anything? For the purposes of this article, let's pretend he is using a quantum process to generate random numbers to "perfectly" randomize his cards. (This will make things seem weird for the next bit of the article!) Whether Quincy himself is such a quantum decision maker is an open question. Yet, consider that colloquially randomness is a manifestation of incomplete information. For example, if one measured a flipped coin trajectory with sufficient accuracy, the outcome of the toss could be predicted to an arbitrary precision. Indeed, people have used such approaches to try and gain an edge at games like roulette by predicting a range of outcomes for a spin of the ball. In War, if one knew the order and content of a player's initial deck, the winner could be reliably predicted. The information contained in the hidden cards is obtainable but it is simply not known to us. Thus, in War or flipping a coin, the outcome reflects our ignorance rather than a truly random trial -- the probabilities are 50%/50% from our completely ignorant point of view.

Further, after tossing the fair coin a large number of times, there will be a Gaussian distribution of outcomes with the most likely being half heads and half tails. The variance in the distribution of outcomes falls like the number of trials - after a very large number of flips the outcomes become highly clustered around half heads and tails and large deviations become rare. This captures our intuition of randomness via a completely deterministic process. Note the statistics of winning and losing at War are identical to the coin flip results as the games map onto each other perfectly.

Conversely, quantum processes have been demonstrated to be "truly" random in the sense there is no more information available no matter how carefully measured. The outcome of an experiment can only be known probabilistically. Indeed, random numbers that drive computations might be generated by quantum processes for particularly sensitive applications like encryption.

Suppose Quincy has 20 cards left after losing in the first round of War. His quantum re-shuffling produces 20! = 2432902008177640000 ≈ 2.4x1018 different possible outcomes -- there are 20! possible outcomes of the partial deck shuffling. Consider, many physicists have reasonably inferred that the full set of possible random quantum events each actually occurs, not just the one outcome that an individual perceives. The idea is that the universe splits into different branches realizing each possible outcome with its associated probability. That would be fully 2,432,902,008,177,640,000 branched, equally likely, Quincy universes in this case. This interpretation is a logical consequence of the dynamics of the universe being driven by the time dependent Schrödinger equation, its mathematical behavior taken literally. Each quantum shuffle generates an ensemble of Quincys each with a different fate in War. Quincy the creator -- truth is stranger than fiction.

So, after Quincy quantum-shuffles / randomizes the cards in the middle of the War game, how does this change things back on our branch of the universe? Clearly, the outcome is no longer determined by the initial two half deck of cards. Indeed, because the outcome of the shuffling procedure is a random process, the winner can no longer be predicted from the initial conditions. This does not change the strategic framework of the game, but rather alters the deterministic nature of the dynamics. The ultimate outcome of the game is unknowable ahead of time. It also creates new branches of the universe with my son and I each winning half of the time, averaged over this multiverse. This interpretation of quantum reality is known as Many Worlds or Everettian Quantum Mechanics.

Thus, in randomized War the entire game must be played out to determine a winner. We are now flipping a quantum coin. In the game of War Quincy and I actually play, the deterministic shuffling of the cards mid-game serves an equivalent purpose as we have perfectly incomplete information about the mixing of the cards. The winner can no longer be predicted from the starting decks, but the outcomes still follow coin flipping statistics.

This distinction between incomplete information, i.e. probabilistic randomness, and quantum or "true" randomness has implications for assessing free will. Decision making maps onto the coin tossing game in a sense. Free will is a term without meaning in a deterministic universe - one is destined to their fate by the initial conditions of the universe. In such a universe, our decisions and past behaviors reveal our predetermined yet evolving nature to ourselves and others.

However, it has been suggested there are quantum random elements to human decision-making due to, e.g. certain biochemical processes in the brain. Some have suggested this permits the possibility of some type of free will. Consider, this situation is analogous to the quantum game of war - it makes the outcomes unpredictable but no less deterministic. Alas, there is no room for free will via adding quantum processes to human decision making. Humans are incomplete information processing machines. Adding quantum random elements to a decision process makes the outcomes ultimately unpredictable but not free in any sense. The universe as it is currently understood has perhaps many (infinite?) outcomes that each unfold as prescribed by some initial condition of unknown origin. If interested, such ideas are presented thoughtfully by Sean Carroll in his writings including books for the public and a terrific podcast.

Back to the idea of reducible games -- this all relates to how I hear professionals talk muddily about poker in the following sense. People are motivated to ascribe a reason to the plays they make. These ultimately rely on the value of the play in some model of the full complexity poker game. For example, considering short stacks, one can construct an unexploitable shoving range as the game simplifies when the stack to pot ratio becomes sufficiently small. No other bet can maximize value for our range better than moving all-in -- simple mathematical models can be evaluated to decide how to play each particular holding. In this limit, characterizing a bet for value or as a bluff has limited utility. One does best by betting our entire stack with a set of hands and folding the rest. With deeper stacks, players can often meaningfully characterize their bets as a bluff or for value. This is based on implicit, intuitive assumptions - essentially a reduced model of the full poker game.

Generally, complex games cannot exactly be reduced to simpler models, yet we rely on approximate model games to guide us. An experienced poker player will find spots where a useful bet, one that optimizes a player's expected value, cannot be easily characterized. It is simply the best of all the available tactics. Sometimes, different possible actions are of equivalent value and one can choose between the plays without cost. For example, there are game states in no-limit hold 'em where the value of raising, calling and folding are all equal. In this case if we choose, perhaps randomly, to raise in the spot what is the purpose? Is it for value or a bluff? If it is for value, is the call or fold in that situation also for value? Actions have distinct costs versus different parts of an opponent's range that are averaged to get an overall expected value. Simple characterizations are only clear in reducible, well-defined situations. Ultimately, some complex game states may be irreducible -- the game is just the game and the best play is the best play. It may not be possible to ascribe simple logic to the merit of the optimal play. This may also be true of our Universe. Some of the properties of reality may not follow from reduced, physical models of the cosmos.

The essence of the matter is that complex games will involve optimal plays that can't easily be characterized by some reduced model of the fully complex game. The bet may not be for value, a bluff or even a merge. Merge characterizes situations where a bet might fold out better hands and get hands that are behind, typically with a draw, to call. The coining of such words reflects the difficulty in characterizing strategic situations that arise in no-limit hold 'em. A given play may simply be the best strategic option. For example, this kind of thing is seen when one simulates wide range situations using solver software. Consider heads up play or button vs. blind confrontations where the "solution" to the situation, given the assumptions required in the input to the solver, can involve doing what seem to be crazy plays that occur with low frequencies. Imagine seeing someone raising large with no pair or draw, it would be hard to characterize the play logically. Nonetheless, the unusual play can be part of the actual solution equilibrium.

Expert poker players search for simplified models, like betting large with a polarized range, where the game dynamics are well understood. One can attempt to force our opponents onto parts of the game tree where the chosen plays are known to be highly profitable. Poker is a beautifully complex game, where ideas of GTO, optimal, exploitive and / or counter strategic play all have their place. We seek to understand and simplify the game dynamics to get an edge on our opponents.

Much of the confusion and argumentation between poker pundits and acolytes is a result of not understanding the assumptions inherent in the model of the game a person is implicitly assuming in their argument. Ultimately, the only observable of merit in poker is the expected value of the play. To teach an artificial intelligence to play poker, it requires only the rules and the value of the outcome as the objective function to optimize. For humans, the expected value is a quantity that cannot generally be evaluated without a large number of assumptions with regard to the game dynamics. AI's can run enormous numbers of trials and learn from vast experience beyond human capability. Put another way, in most of the arguments I hear about the value of poker strategies, people literarily are unsure of what they are talking about but appear to enjoy a good argument nonetheless.

As an aside, I have an interest in poker theory. I believe that the nature of the optimal solutions is not yet completely understood. It seems to me from an information theoretic perspective that bet sizing should draw from distributions, providing information hiding. Solver work suggests this might be right. I also see an analogy between (statistical) mechanical energy and expected value in that the solution space of poker is a surface of constant expected value. Lastly, I have noted that the game theoretical optimal strategy is the one that requires no information of the opponents play suggesting a Shannon entropy tie in. All of this suggests to me a statistical mechanical approach to poker solutions that I have not formulated. If anyone is interested pursuing this feel free to get in touch.

Brian Space is a scientist and professor seeking people to play Quantum Statistical Mechanics for money. He plays poker in the Tampa Bay Florida area. His poker articles are available on his web site: http://drbrian.space/poker.html