Over Betting Polarized Ranges in the Short Run: How Big to Bet?

by Brian Space

Two Plus Two Magazine, Vol. 12, No. 2

Poker is a game driven by mathematics clouded in a statistical haze. Luckily for experts, this austere math reality is well hidden by a large human element and a lot of immediate unpredictability. However, after a sufficiently large number of hands, the unpredictable becomes the inescapable. This inevitability is played out in the online poker world, where a reasonable sample size of hands is played in a relatively short time period; amateurs lose quickly to players with increasingly superior strategies.

To design a mathematically sound strategy it is reasonable to consider "toy games" and/or reduced complexity situations that are nonetheless representative of actual game conditions to varying degrees. This article was inspired by a Ben Sulsky Run it Once video called *Toy Gaming*, itself a descendent of earlier foundational poker writings like the seminal the *Mathematics of Poker.* Simple games make certain plays straightforward to mathematically derive and offer obvious principles to apply in a variety of more complex situations. One common play is over betting a polarized range versus a capped range in no-limit games; this particular tactic is very difficult to defend and is often a license to print money used judiciously. It also is a fairly common scenario that arises in a variety of situations. For example, it is very often the case in no-limit hold 'em that the river aggressor is betting only very strong hands in their range, along with bluffs. The person faced with a decision has a medium strength hand that beats the bettor's bluffs but has no chance versus the value part of the bettor's range. Betting large with this polarized range is a winning strategy element for all superior players. The question to be considered here is how large to bet as a practical matter.

This is an example of game theory optimal (GTO) concepts applied to poker. The GTO solution for an individual, collusion aside, is going to exist in every spot. This is related to but different from equilibrium solutions that can exist for any strategy - any play can be balanced. Multiway, the GTO solution would depend on the exact stack sizes of every player and other participants also playing optimally. Responses throughout the hand depend on what action is made by each player, further complicating things. Thus, the GTO solution is not generally relevant but the concepts that can be extracted from it are useful. Even a heads-up river solution, popular now given the existence of GTO solvers, provide strategies with low fractional frequencies that are only sensible if one plays an extremely large number of hands, even given that they consider limited input turn ranges. Note, there are only so many responses a player can make: check, fold or bet some size. So almost any vacuum play may be part of a GTO strategy, albeit perhaps with a very low frequency. This doesn't mean they need to be part of a good winning strategy. GTO play is a set of frequencies, not a required play in a particular spot - it need not be the best play for a given situation but rather for the range on average. For example, deep stacked very low frequency plays may be included to threaten unlikely run outs given the possibility of large later returns. The point here is that to take advantage of optimal solution strategies, frequencies of actions need to be realistic for the scope of the game play.

Start by considering the statistical convergence of the over betting play. The mathematics, outlined below, cannot be denied. However, given that strategies offer only a statistical edge versus our opponents it is worth considering both how our advantage is achieved in practice and how liberally the concept of over betting can practically be applied. The idea here can be summarized that, while any size over bet represents a mathematically sound strategy, the benefit is only achieved after a large statistical sample. It works in the long run, but as Keynes pointed out, in the long run we are all dead.

Consider a situation where Villain is the aggressor on the river and is betting a polarized range of AA and 22 and Hero has KK. This is an idealized, best case scenario, in that Hero has perfect information to defend. First, consider the situation asymptotically to demonstrate how iterative improvement from the extreme cases will converge to an equilibrium strategy where no one can improve their play. In this case, the bettor of the polarized range will be shown to capture an ever-increasing pot share with increasing bet size relative to the pot. To make things concrete, consider first $100 stacks for both Hero and Villain with $100 in the middle - a pot size bet remaining on the river. Villain has AA (6 combinations) and 22 (6 combinations) and Hero has KK (6 combinations).

First, suppose Villain is a nit and bets only AA, then a clairvoyant Hero simply folds. This gives Villain's betting range an expected value (EV) of $50, where Villain wins the pot half the time betting (AA>KK) and Hero wins the other times (KK>22) when Villain checks. Now if the nitty Villain is feeling frisky and starts bluffing one of the six 22 combinations things get worse for Hero. Hero still has no better option than to always fold, because the nitty Villain still isn't bluffing enough to alter our strategy, but the expectation of Villain's range is now EV = (7*$100) / 12 = $ 58.33. This is a clear improvement and Villain has easily captured an extra 1/12 or 8.3% of the pot, demonstrating the power of bluffing with a polarized range.

Note, if Hero overreacts to our tilting nit and always calls, Villain's EV increases to EV = (6*$200 - 1*100) / 12 = $ 91.67 and it is easy to show no strategy by the clairvoyant Hero is better than always folding. Imagine our Villain is now more dangerous, adding bluffs with Hero always folding: considering two-six 22 combinations bluffed, Villain's EV becomes respectively: $66.67, $75, $83.33, $91.67 and finally, with Villain always bluffing and the demoralized Hero always folding, $100, or the whole pot. Clearly at some point the aspiring Hero needs to start calling.

Absent a guessing game, called exploitive play, the solution to this conundrum is well known and it is easy to show that Villain can simply make Hero indifferent to calling or folding. Before designing a practical strategy, first consider the detailed mathematics of the situation. If Hero is getting x to y on a call, Villain bluffs with a frequency y/(y+x) adding a proportion of (y/x) bluffs to each value hand. This, like all GTO solutions, can be thought of as the end result of an iterative improvement of strategies by both parties until neither can improve.

Here, Villain is betting the pot so Hero gets x=2 to y=1 on a call, making Villain bluffing frequency (1/3) and letting Villain bluff (1/2) as many combinations as there are value bets; (1/2) *6 = 3 combinations of 22 may be bluffed optimally for this bet size. The optimal, equilibrium EV for Villain is $75. For example, in optimal play if Hero always calls, the EV = (6*$200 - 3*$100) /12 = $ 75, and always folding yields, (9*$100) /12 = $75. Any other Hero strategy is a weighted average of these strategies and gives the same EV. Note, optimal here refers to the highest EV equilibrium play for the fixed Villain bet size - the choice of this sizing is the subject of the current discussion. A larger, all-in sizing would always be optimal in the common math use of the word.

Note, all is not lost for Hero as an optimal defense frequency of 1-(y/x) = .50 here will insure Hero is unexploitable for any Villain strategy and punishes Villain for deviating from playing optimally, but is not maximally exploitive in that case. For example, if Villain over-bluffs 4 combinations for a pot size bet, Hero's optimal strategy is to call every time and then Villain's EV falls to EV = (6*$200 - 4*$100) / 12 = $66.67. If Hero uses the optimal defense frequency of 50%, Villain's EV = 0.5*(6*$200-4*$100) /12 + 0.5*$100*10/12 = $75. This is a general feature of optimal play, when Villain's strategy deviates, Hero's best approach changes from an optimal frequency to a defined exploitative strategy, here always calling.

There is no defense for the betting of a polarized range under these conditions; once Hero has capped their range versus a polarized range much EV is unrecoverably lost against a skilled opponent. Put another way, *a priori* Hero has the best hand, KK, half the time but can't capture all the equity. But the situation is even worse when stacks are deeper. If both Hero and Villain have $200 with $100 in the pot, Villain maximizes the EV by betting two times the pot and now Hero gets x=3 to y=2 on a call, making Villain's bluffing frequency (2/5) and letting Villain bluff (2/3) as many combinations as there are value bets; (2/3) *6 = 4 combinations of 22 may be bluffed optimally for this bet size. Under this scenario Villain's EV = (10*$100) / 12 = $ 83.33; because the EV is independent of Hero's action we can consider the simplest case of Hero always folding in disgust - knowing 22 is raking the pot quite often. As stack sizes increase the mathematics insist that betting all-in for increasing multiples of the extant pot capture more and more of the equity for Villain, asymptotically approaching 100% of the pot. All this given that Hero has the best hand half the time.

Consider an extreme case betting 99 times the $100 pot, $9900. Hero is getting x=100 to y=99 pot odds. Using the formulas presented demonstrates that optimally Hero has ninety-nine 22 bluffs for every one hundred AA combinations in a betting range and one 22 hand in a checking range. The EV of this strategy is $99.50. For example, Hero is indifferent and can fold every time with Villain winning (199/200) of the pot.; Hero only wins when Villain checks 22 one in two hundred times. Alternatively, if Hero calls every time Villain wins (100/200)*$10,000 - (99/200)($9900) = $99.50. They are the same as they must be and any other defense frequency, including the optimal 1% frequency is a weighted average of these numbers giving the identical EV of $99.50. This result suggests taking care with sizing in betting a polarized range. When Villain bets large multiples of the pot over a finite sample, Hero often regains almost 50% equity by always calling in the (AA,22):(KK) bluff catcher toy game. Villain is betting nearly 50% each AA and 22 as the math permits almost an equal fraction of value bets and bluffs to offer indifference to Villain. A clairvoyant Hero, being indifferent, can now adopt any strategy and isn't limited to an equilibrium or optimal choice. Thus it will be infrequent, if Hero chooses a calling strategy, for Villain to gain any equity by betting a polarized range for large multiples of the pot over a realistic number of representative hands.

The rigorous math still holds up because Villain checks 22 once in a blue moon and wins often enough to make Villain's expectation the original pot size. Though indifferent over an infinite sample, Hero almost always does better by always calling versus always folding over a finite sample, and rarely much worse. Thus, practical considerations and actual sample sizes start to matter. Statistical convergence is faster for pot size bets and the poker is also more sensible as fewer "human" considerations are involved when the bet size resembles the pot size; unusually large bet sizes might have an unpredictable effect on Hero's behavior making an exploitative strategy more relevant; variance is also much less of an issue when bet sizes are similar to pot sizes as well.

The issue is that over a finite sample, when Villain bets large multiples of the pot, it involves betting nearly equal numbers of bluffs and value hands. Thus, a strategy of always calling will very often capture nearly the full original range equity of half the pot. This is a realistic issue that needs to be accounted for in a strategy. For example, in tournament play where chip accumulation and survival are critical, betting relatively smaller may be preferable to more frequently realize an advantage and to reduce variance. This is also true for live play with limited opportunity to repeat the same scenario. This suggests the more common approach of building strategies around small multiples of pot size bets is generally more sensible but also that mathematically exact solutions and asymptotic conclusions should be treated with care.

None of this changes the exact math of over betting. In a sense, it exposes an asymptotic property of the over betting that is itself exploitable. In the large limit of over betting, the pot itself is small compared to the bet. At equilibrium, nearly half of the bets are bluffs, and if Hero folds at an unexploitable frequency, Villain gets almost the whole pot. But the pot is not negligible, it is the money in the pot Villain seeks to win and this changes the frequency of bluffing in just the right way to capture an increasing fraction of the pot. And even if done exactly correctly, Hero can nonetheless call over a small finite sample and capture half the pot most of the time in the large over bet limit. Further, if Villain is betting to make Hero indifferent, Hero can call as frequently as they like and not be exploited. This suggests Hero can simply call and, given the rarity of such situations *e.g.* in live play Villain will have a hard time making exploitative betting range adjustments to *e.g.* bluff less versus a call happy Hero.

Exploitative considerations may be most important. Consider if Hero knows Villain is bluffing half the time, as a practical matter why give away most of the pot? Skilled players are aware of the result that, as bet sizes get large, they can bluff more and in the limit of large over bets, have a range of (nearly but not) half bluff combinations. The point is Villain cannot bluff equal numbers of combinations or a Hero strategy of calling recaptures half the pot – which is the equity Hero had versus the nit. Care is needed to keep careful track of the number of allowed bluff combinations in a given spot and get the frequencies just right for larger over bets – I suggest this is unlikely in all but the most common spots for online players. All of this suggests over betting in a live setting is going to be dominated by exploitive considerations if not constructed with extreme care.

Brian Space is a scientist, but can't find people to play quantum statistical mechanics for money.