Bonus Maths 3: A New Toy Poker Game: NLHE Without The Inconvenience of Cards? (JB)

A few years ago I was playing SnGs and started to have some doubts about whether ICM really is an accurate guide to the shove or fold game. I wanted to try some different models but, being old and lazy, couldn't face either writing code to compute preflop equities in NLHE or learning to use somebody else's code. I wanted a model that had all the important features of preflop NLHE without any of the computational pain – something that I could write down in one equation. Once you've stopped laughing, read on and I'll show you the game that I invented, demonstrate that it does approximate preflop NLHE very well, and speculate on some extensions that might make this sort of model applicable to other poker variants.[1]

Hand Strength, Not Hand Rank

In the book, I discussed the [0,1] game, where a hand is replaced by a number between zero and one. If there's a showdown, the lowest number wins, e.g. if I'm dealt 0.265 and you're dealt 0.891, I would win at showdown because 0.265<0.891. Although hands are dealt randomly, there's no subsequent random element in the game because the ranking of the hands is static. This sort of toy game is very useful for learning how to think about situations on the river in NLHE, where hand values are indeed static. However, before the final round of betting, the cards that appear on flop, turn and river change the relative strength of each player's hand. In preflop shove or fold NLHE, the players have to decide whether to go with their hand by shoving in their stack, taking into account the likely effect of the five community cards on their final hand strength.

In *The Mathematics of Poker*, Chen and Ankenman discuss a HU [0,1] shove or fold game where, if one player shoves and the other calls, the better hand wins two thirds of the time. They found that this is a much better approximation to real NLHE than the game where the better hand wins every time. Let's see whether we can do any better.[2]

Consider a game where each hand, h, with $0 \leq h \leq 1$, has a strength given by a *hand strength function*, $H(h)$. If two hands h_1 and h_2 go to showdown, h_1 wins with probability $H(h_1)/\big(H(h_1)+H(h_2)\big)$, and h_2 wins with probability $H(h_2)/\big(H(h_1)+H(h_2)\big)$. As you can see, these probabilities add nicely to one, and there's a natural extension when more than two hands go to showdown.

The simplest hand strength function that I can think of is $H(h) = 1$. In this game, all hands have equal strength, and $H(h_1)/\big(H(h_1)+H(h_2)\big) = 1/2$, so all showdowns are coin flips. This isn't a very interesting game. The next simplest hand strength function is $H(h) = a - h$, where $a \geq 1$ is a constant. Now we're starting to get somewhere. The strongest hand, $h = 0$, has strength a, and the weakest hand, $h = 1$, has strength $a - 1$. If we take $a = 4/3$, then

[1] I'm not going to talk about alternatives to ICM, but if you want to see some ramblings on this subject, you can look at http://www.pokerstrategy.com/forum/thread.php?threadid=94918. To cut a long story short, I found that ICM does pretty well, but had lots of fun playing off various models against each other in three-handed bubble situations.

[2] I've tried to shovel as much of the maths as possible into an appendix after my co-authors claimed that the equations made their heads hurt, but you have to see at least this little bit.

$H(1)/\left(H(1)+H(0)\right) = a/(2a-1) = 0.8$, so that the strongest hand has 80% equity against the weakest. That's not a million miles away from preflop Holdem.

HU Shove or Fold

We'll use the standard blind structure, and assume that John posts the small blind and Tom the big blind. Each player starts with a stack of S big blinds. John must decide whether to shove or fold, and Tom must decide whether or not to call a shove. If John has a hand j, he will shove if this hand is strong enough, i.e. if $0 \le j \le j_s$. If Tom has a hand t, he will call if this hand is strong enough, i.e. if $0 \le t \le t_c$, and there will be a showdown. A little bit of calculus is all that's required to find the optimal values of these shoving and calling ranges, $j_s$and t_c, and the details can be found in the *Appendix*.

I've plotted some typical solutions in Figure 1 as a function of stack size, S, and also showed you the solution for NLHE, as given by the HoldemResources Nash calculator. As you can see, the linear hand strength game and NLHE have qualitatively similar shoving and folding ranges, provided that a is small enough. However, John's equity in NHLE is significantly lower. Is there any rational way in which we can choose a slightly different hand strength function and capture a little more of the way that preflop equity is distributed in NLHE?

Inspired by the plot of hand strengths in Andrew Seidman's book *Easy Game*, we can note that a few hands (AA, KK, AK, QQ, etc.) have significantly more equity than most other hands in NLHE. We might therefore expect that a hand strength function that is peaked close to the strongest hand, $h = 0$, might be a better fit to NLHE. The simplest function like this that I can think of is $H(h) = a + be^{-kh} - 1$. Armed with three parameters (a, b, k) with which to fit this game to NLHE, I was able to reproduce NLHE's shove or fold ranges and equities quite well, as you can see in Figure 2. The hand strength function that produced this fit is shown in Figure 3, and is biased towards the top 15% of hands. The best hand has about 80% equity against the worst hand. As promised, we have a proxy for NLHE using one simple equation.

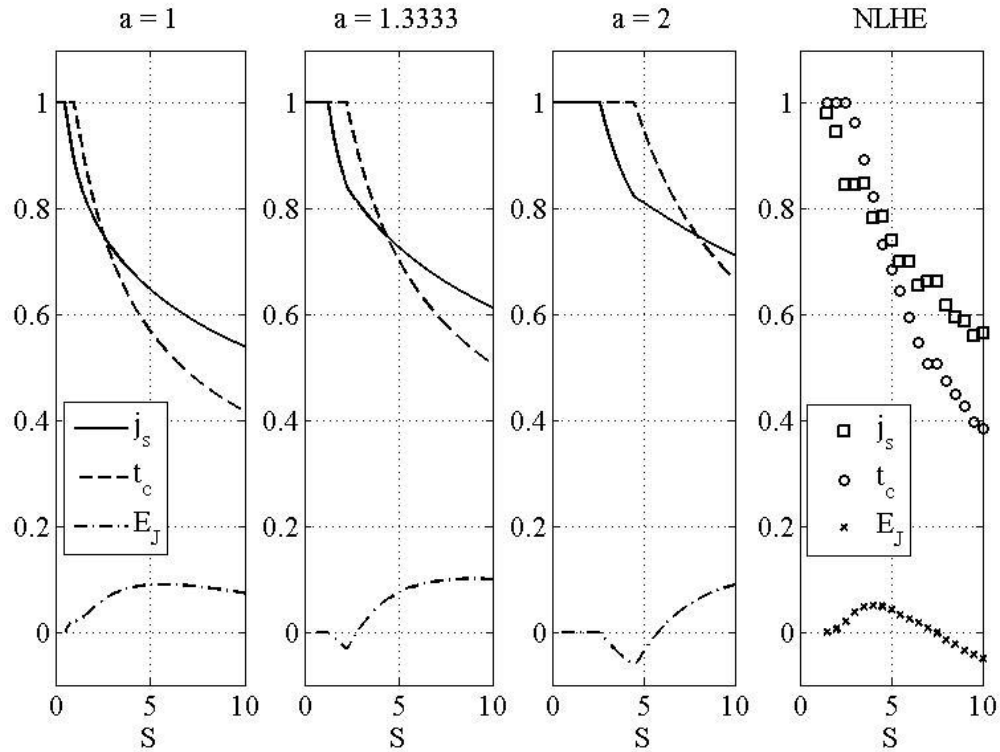

Figure 1: Shoving ranges, calling ranges and John's equity for H(h) = a-h and NLHE.

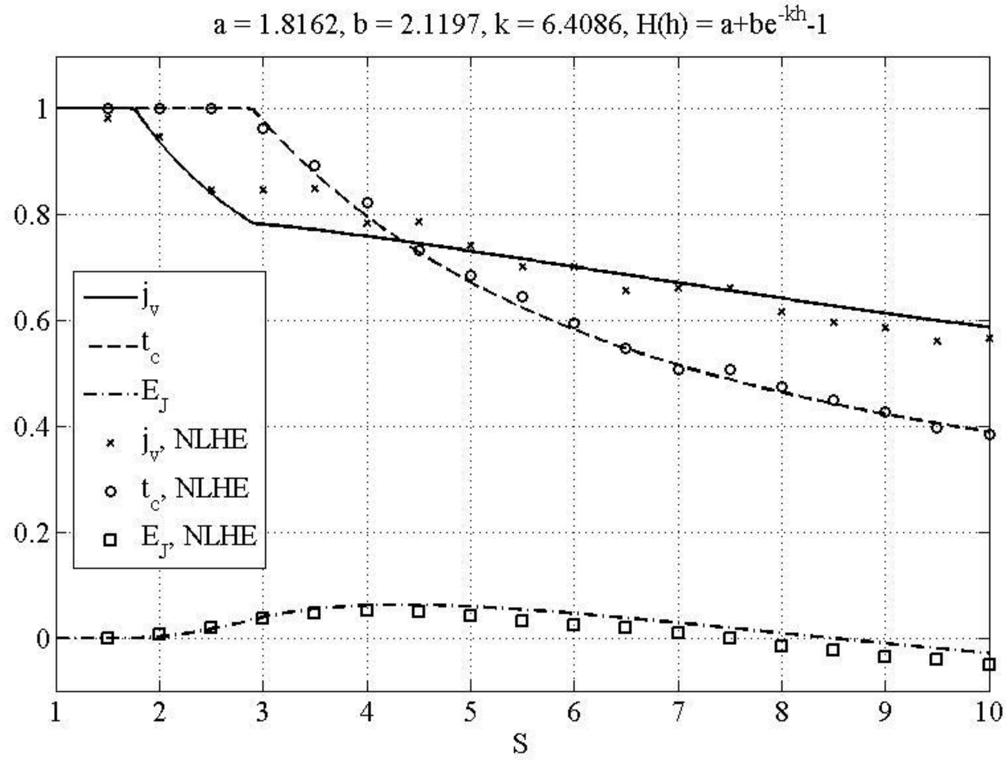

Figure 2: The best fit of a simple, nonlinear hand strength function to shove or fold NLHE.

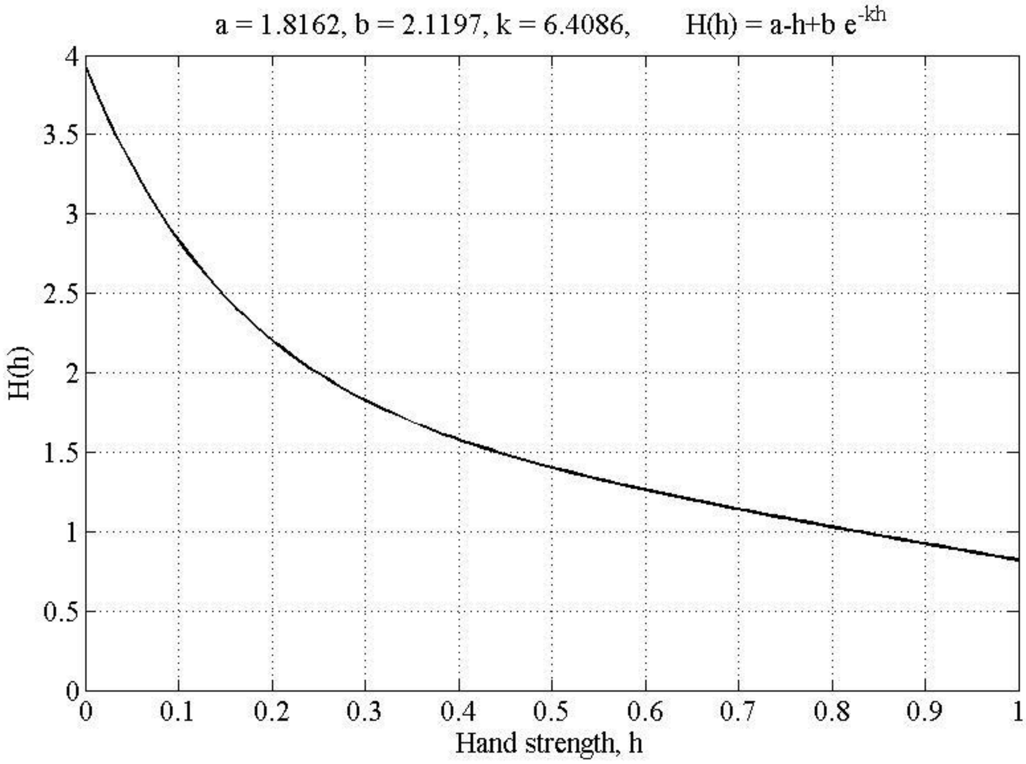

Figure 3: The hand strength function that gives the best fit to shove or fold NLHE.

Fixed Pot, Fixed Bet Size

Another version of this game that we can take a look at is the equivalent of the half street [0,1] game, which we studied in the book. There is a pot of size $$P$, John has to check and then Tom can decide either to check behind and show his hand down or bet $1. Remember that in our new game 'showdown' means the hands are revealed and Tom wins with probability $H(t)/(H(t)+H(j))$. If Tom bets, John has to decide whether to call. For the standard [0,1] game, we found that Tom value bets his best hands, bluffs with his worst hands and checks behind to show down his medium strength hands. John must call with his best hands using a range wider than Tom's value range so that he has some bluff catchers. What do you think that the effect of using the hand strength function shown in Figure 3, instead of simple [0,1] hand ranking, would be on John and Tom's optimal strategies? It turns out (see the appendix) that Tom should never bluff. His worst hands have too much equity to waste them as bluffs against an opponent whose hands also have a lot of equity and will be calling with a lot of hands. This is in complete and, at least initially to me, surprising contrast to the standard [0,1] game.

I'd like to use this example as an excuse to introduce a method for solving games numerically that proper Game Theorists often use, called *fictitious play*. Just because I'm not a proper Game Theorist doesn't mean I can't use it too. The idea is that each computerised, fictitious player takes it in turn to choose the parameters of his strategy to maximise his expectation. Fictitious Tom value bets his strongest hands, $0 \leq t \leq t_v$, bluffs with his weakest hands, $t_b \leq t \leq 1$, and shows down his midstrength hands. Fictitious John calls with his best hands, $0 \leq j \leq j_c$, and folds the rest. In this game, Fictitious John chooses j_c to minimise Fictitious Tom's expectation, then Fictitious Tom chooses t_b and t_v to maximise his own expectation, and this turn taking carries on until their strategies stop changing. In practice, you need to introduce some damping into the procedure to stop it just cycling around forever, but it seems to work well (I'm no expert).

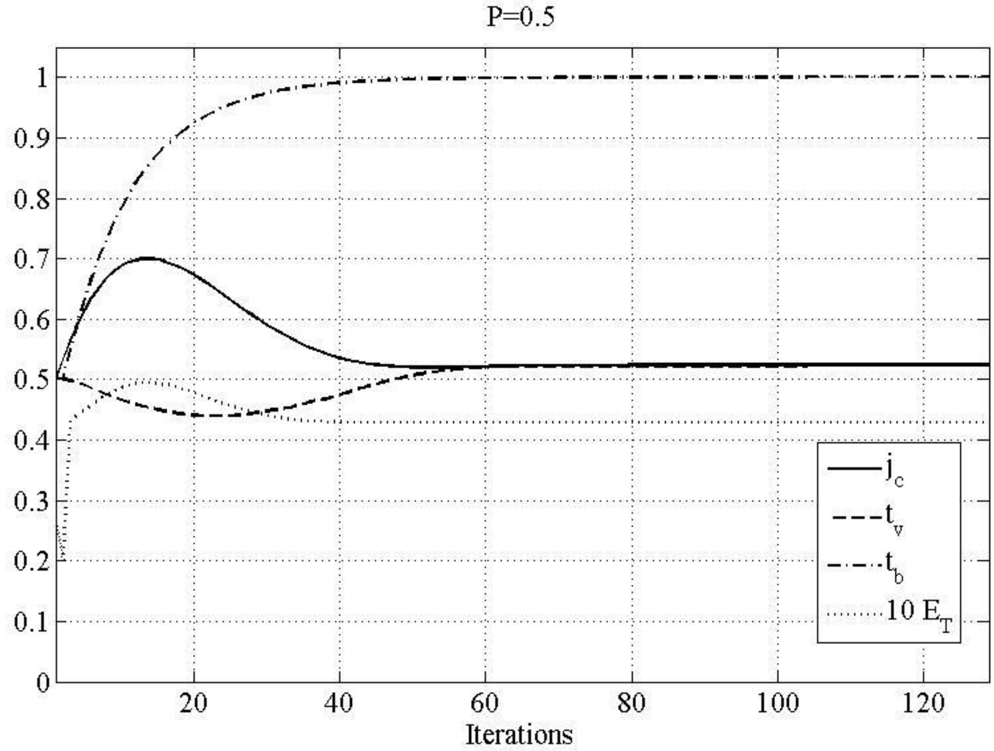

Figure 4: An example of fictitious play in operation. Tom's expectation, E_T, has been multiplied by 10 so that you can see what's going on.

Figure 4 shows an example where initially $j_c = t_b = t_v = 0.5$. Fictitious John starts by calling wider, which causes fictitious Tom to bluff less and value bet tighter, until eventually, as predicted, fictitious Tom stops bluffing and their value betting and calling ranges become close to each other.

It may, in principle, be possible to construct a hand strength function $H(h)$ for which the optimal solution involves bluffing, but I haven't been able to after a few (rather desultory) attempts. Even if I had been successful, the resulting hand strength function would probably not have had much relevance to any real games.

Extensions?

I hope I've convinced you that I've constructed a simple proxy for preflop NLHE, which is also an interesting counterpoint to the half street [0,1] game in that bluffing is not required in the optimal solution. Of course shove or fold NLHE – basically very short-stacked NLHE - is much simpler than even shallow-stacked NLHE. Can this simple game be extended to make it more like deeper stacked NLHE? Once the flop appears, hand strengths change. In order to capture this in a toy game, there needs to be a round of betting, and then something random has to happen to the hand strength function, with the strength of some previously weak hands improving. I'm not sure how to do this yet, but I have it in mind to try. A variant that might be more amenable to this sort of modelling is draw poker, with No Limit 2-7 Single Draw the simplest example. A modification of the hand strength function would correspond to the draw, but postdraw the strength of each hand would be fixed, i.e.

the best hand has 100% equity postdraw. I'm not quite sure where to go with this, but if I come up with anything interesting I'll let you know. I also have the option of passing this on to undergraduate project students to see what they make of it. I'm sure they'll be pleased.

Appendix: The Maths

HU Shove or Fold

We want to determine the equilibrium values of j_s and t_c for a given hand strength function, $H(h)$. In order to do this, we first need to work out John's expectation for this game, which is

$$E_J = \frac{3}{2} j_s(1 - t_c) + \int_0^{j_s} dj \int_0^{t_c} dt \, \frac{\left(S + \frac{1}{2}\right) H(j) - \left(S - \frac{1}{2}\right) H(t)}{H(j) + H(t)} - \frac{1}{2}.$$

If John folds, his expectation is zero. If John shoves and Tom folds, John wins 3/2 big blinds, which gives the first term. If John shoves and Tom calls, there's a showdown and John wins $\left(S + \frac{1}{2}\right)$ big blinds with probability $H(j)/\big(H(j) + H(t)\big)$ and loses $\left(S - \frac{1}{2}\right)$ big blinds with probability $H(t)/\big(H(j) + H(t)\big)$. We need to add this up over all of John's possible shoving hands and Tom's possible calling hands, which gives us the integral in the equation. Finally, we subtract half a big blind to account for the fact that the small blind came from John's stack. This means that if John always shoves and Tom always calls, $E_J = 0$, so this is John's ex-showdown expectation. A small manipulation of the integral gives us the form

$$E_J = \frac{3}{2} j_s - j_s t_c - \frac{1}{2} + S \int_0^{j_s} dj \int_0^{t_c} dt \, \frac{H(j) - H(t)}{H(j) + H(t)},$$

which is somewhat easier to work with.

At the equilibrium, this expectation is stationary with respect to the parameters j_s and t_c, so we have to find the solution of the simultaneous, nonlinear equations

$$\frac{\partial E_J}{\partial j_s} \equiv \frac{3}{2} - (1 + S) t_c + 2S \, H(j_s) \int_0^{t_c} \frac{dt}{H(j_s) + H(t)} = 0,$$

$$\frac{\partial E_J}{\partial t_c} \equiv (S - 1) j_s - 2S \, H(t_c) \int_0^{j_s} \frac{dj}{H(j) + H(t_c)} = 0.$$

With $H(h) = a - h$ you can work out these integrals, but even then you get some horrid, nonlinear equations, so I just slapped the whole lot, in the form given above, onto my computer and told it to find j_s and t_c for me. The results are given in Figure 1

Fixed Pot, Fixed Bet Size

Let's assume that the optimal strategy is the same as for the comparable [0,1] game. Tom value bets his strongest hands, $0 \le t \le t_v$, bluffs with his weakest hands, $t_b \le t \le 1$, and shows down his midstrength hands. John calls with his best hands, $0 \le j \le j_c$, and folds the rest. One thing that we can note immediately is that, since John needs to call \$1 to win \$$(P + 1)$, he can certainly call if he has more equity than $1/(P + 2)$ against Tom's strongest hand, i.e. if $H(j)/(H(j) + H(0)) \ge$

$1/(P+2)$. In particular, we must have $H(j_c)/(H(j_c)+H(0)) \geq 1/(P+2)$. We can also deduce that if $H(1)/(H(1)+H(0)) \geq 1/(P+2)$, John's weakest hand has enough equity to call a bet from Tom's strongest hand, so he will call every time. For the hand strength function shown in Figure 3, this is true if $P \geq 3.80$. For a roughly ¼ pot bet, 20% equity is enough to call.

Tom's expectation using this strategy is

$$E_T = \int_0^{j_c}\int_{t_b}^{1}\frac{(P+1)H(t)-H(j)}{H(t)+H(j)}dt\,dj + \int_0^{j_c}\int_0^{t_v}\frac{(P+1)H(t)-H(j)}{H(t)+H(j)}dt\,dj$$
$$+\int_0^{1}\int_{t_v}^{t_b}\frac{PH(t)}{H(t)+H(j)}dt\,dj\ + P(1-j_c)(1-t_b+t_v)-\frac{1}{2}P.$$

If Tom either bluffs or value bets and John calls, Tom wins the $\$(P+1)$ in the pot with probability $H(t)/\big(H(j)+H(t)\big)$ and loses his \$1 bet with probability $H(j)/\big(H(j)+H(t)\big)$ – this is captured by the first two integrals. If Tom checks behind, he wins the pot with probability $H(t)/\big(H(j)+H(t)\big)$ – this is the third integral. If John folds to a bet, there's no showdown and Tom wins the pot, which is given by the fourth term in the equation. Averaged over all possible hands, each player wins an equal share of the pot, so the final term is there to make this Tom's ex-showdown expectation.

Now, to find the optimal solution, we need to solve

$$\frac{\partial E_T}{\partial j_c} \equiv 1-t_b+t_v-\ (P+2)H(j_c)\left(\int_{t_b}^{1}+\int_0^{t_v}\ \right)\frac{dt}{H(t)+H(j_c)}=0,$$

$$\frac{\partial E_T}{\partial t_v} \equiv P-(P+1)j_c+\ H(t_v)\left((P+2)\int_0^{j_c}\quad -P\int_0^{1}\ \right)\frac{dj}{H(t_v)+H(j)}=0,$$

$$\frac{\partial E_T}{\partial t_b} \equiv -P+(P+1)j_c-\ H(t_b)\left((P+2)\int_0^{j_c}\quad -P\int_0^{1}\ \right)\frac{dj}{H(t_b)+H(j)}=0.$$

If we subtract the last two equations from each other, we get

$$H(t_v)\left((P+2)\int_0^{j_c}\quad -P\int_0^{1}\ \right)\frac{dj}{H(t_v)+H(j)}$$
$$=H(t_b)\left((P+2)\int_0^{j_c}\quad -P\int_0^{1}\ \right)\frac{dj}{H(t_b)+H(j)}.$$

One solution is therefore $t_b = t_v$. If the function that appears on each side of this equation is not monotone, there may be solutions with $t_b \neq t_v$, but a numerical evaluation of this function for the NLHE-like values of a, b and k suggests that it is indeed monotone, and therefore that $t_b = t_v$. This means that Tom bets with all his hands.

Can this possibly be right? Well, actually no, it can't. When $t_b = t_v$, these equations show that $\frac{\partial E_T}{\partial t_v} = -\frac{\partial E_T}{\partial t_b}$, and therefore that, although Tom maximises E_T by choosing t_v according to this approach, he *minimises* E_T by choosing t_b like this. In fact, he needs to choose $t_b = 1$, i.e. never bluff, in order to maximise his expectation.