The one-up game

Albert Julius Liu

January 29, 2018

An extension to [1].

In this appendix we analyze the "one-up" game. This demonstrates the technique of analyzing a continuous strategy set by discretizing it and running it through our algorithm. We also analyze the effects of restricting one player's strategy set.

0.1 Introduction

The game is as follows:

- The strategies for both players are numbers in the (continuous) interval $[0, 1)$.
- If the maximizer picks a larger number than the minimizer, they gain a base payoff of b. Otherwise their base payoff is 1.
- A uniform Nash equilibrium is to be induced by applying a multiplicative handicap.

This type of mechanic occasionally appears explicitly; for example, as in *Hearts of Iron IV*, it could represent the penetration capability of a cannon versus the thickness of armor, with a flat damage multiplier being applied if the cannon can penetrate the armor. When the base payoff is plotted as a function of the difference between the maximizer and minimizer strategies, it is a step function. As such, it can also be seen as the limiting case of any sigmoid-shaped function as the x-axis becomes compressed; for example, "clipped-difference" payoff functions, where the probability of a successful attack is the difference of an attacker and a defender statistic, but with a minimum and maximum chance, as seen in Figure 1.

0.2 Balancing for uniform Nash equilibrium

To find the multiplicative handicaps that produce a uniform Nash equilibrium, we can use the following technique:

1. Discretize the strategy space.
2. Solve the discretized problem.
3. Graph the result.
4. Guess what function represents the result.
5. Prove that this function is indeed the solution in the continuous case.

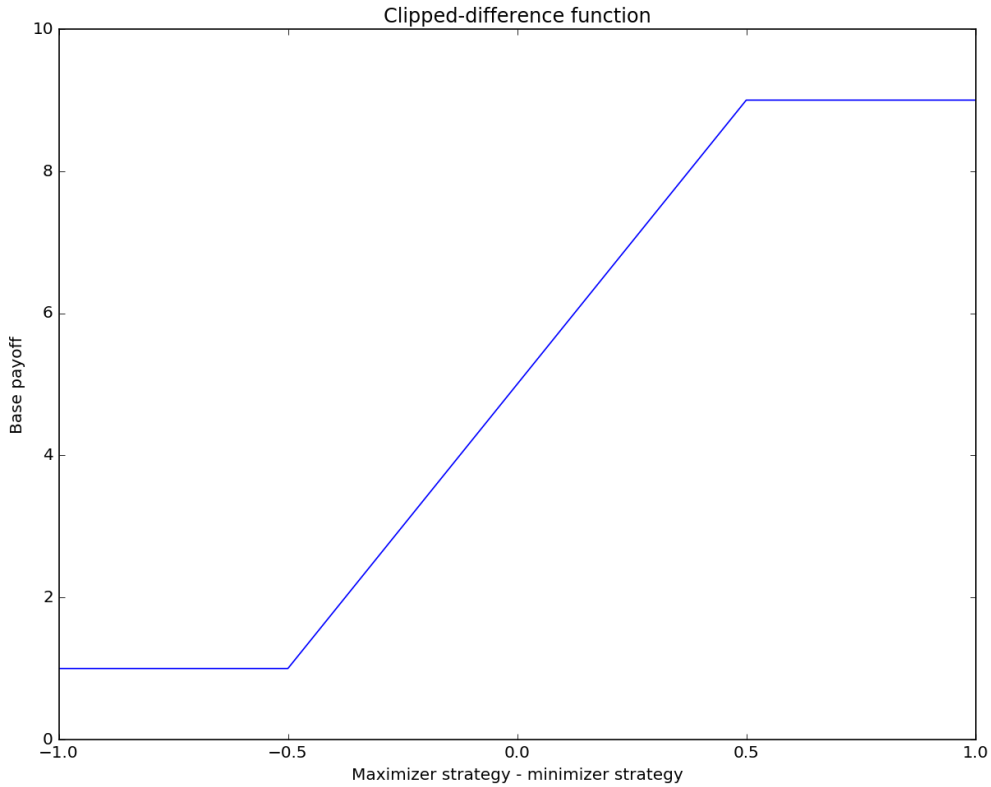

Figure 1: Clipped difference payoff function. As the x-axis is compressed it increasingly resembles the step function of the one-up game.

In this case the desired multiplicative handicap for a strategy x is simply proportional to

$$h(x) = b^x \tag{1}$$

for both players. The payoff if the maximizer plays x and the minimizer plays y is then equal to

$$p(x, y) = b^{y-x} \cdot \begin{cases} b & \text{if } x > y \\ 1 & \text{otherwise} \end{cases} \tag{2}$$

$$= b^{(y-x) \bmod 1} \tag{3}$$

This makes every strategy symmetric, so the uniform distribution must be a Nash equilibrium.

Non-uniform Nash equilibria. For this particular game, producing a desired non-uniform Nash equilibrium is trivial: since the step function only depends on the ordering of the strategies

and not their particular value, to determine the appropriate handicap for a strategy $\tilde{x}$, we can map it through the inverse cumulative distribution function of the desired Nash equilibrium, and then assign it the corresponding handicap from the uniform case:

$$x = \mathrm{CDF}^{-1}\tilde{x} \tag{4}$$

$$\tilde{h}\left(\tilde{x}\right) = b^{\mathrm{CDF}^{-1}\tilde{x}} \tag{5}$$

Essentially, we start with the desired non-uniform distribution and "stretch" the domain until the distribution becomes uniform. This does require that the desired Nash equilibrium is atomless so that the CDF is invertible. An atomless distribution also ensures that the result does not depend on the tiebreaking behavior, the same as with the uniform distribution.

Determining Nash equilibrium from handicaps. We can use a similar approach as an alternative method of solving the inverse problem, i.e. that of determining the Nash equilibrium given a set of handicaps. Specifically, let $h\left(x\right)$ be a continuous function with a positive minimum value normalized to 1. Make the following deletions from the domain:

- Remove all x such that $h\left(x\right) >= b$. These will never be played because their high handicap means that they cannot be lowest-handicap strategy instead.
- Remove all x such that there exists $x' > x$ such that $h\left(x'\right) \leq h\left(x\right)$. These strategies will never be played since there is a higher strategy available for equal or lesser cost.

The remaining function is strictly monotonically increasing and has a range $[0, 1)$. The CDF of the Nash equilibrium is then simply

$$\mathrm{CDF}\left(x\right) = \log_b h\left(x\right) = \frac{\log h\left(x\right)}{\log b} \tag{6}$$

$$\mathrm{PDF}\left(x\right) = \frac{1}{\log b}\frac{\frac{d}{dx}h\left(x\right)}{h\left(x\right)} \tag{7}$$

In other words, the probability distribution is proportional to the proportional rate of change in the handicap (over the trimmed domain). An example is shown in Figure 2.

0.3 Restricted strategy space

Suppose we balance the game as above, but instead of being able to play any strategy in the interval $[0, 1)$, one player has the disadvantage of only being able to play strategies in the interval $[0, a)$—perhaps they are limited by technology or economics and do not have access to higher strategies. What does this do to the Nash equilibrium and expected payoff?

Clearly the other, advantaged player has no incentive to play higher than a. Playing a already achieves the more favorable base payoff (e.g. penetrates the armor, or blocks the cannon) against all available strategies of the other player, and playing a higher strategy merely incurs a more severe handicap for no benefit. To find the exact Nash equilibrium we again use the discretization technique. Based on the result we conjecture that each player uses an extremal

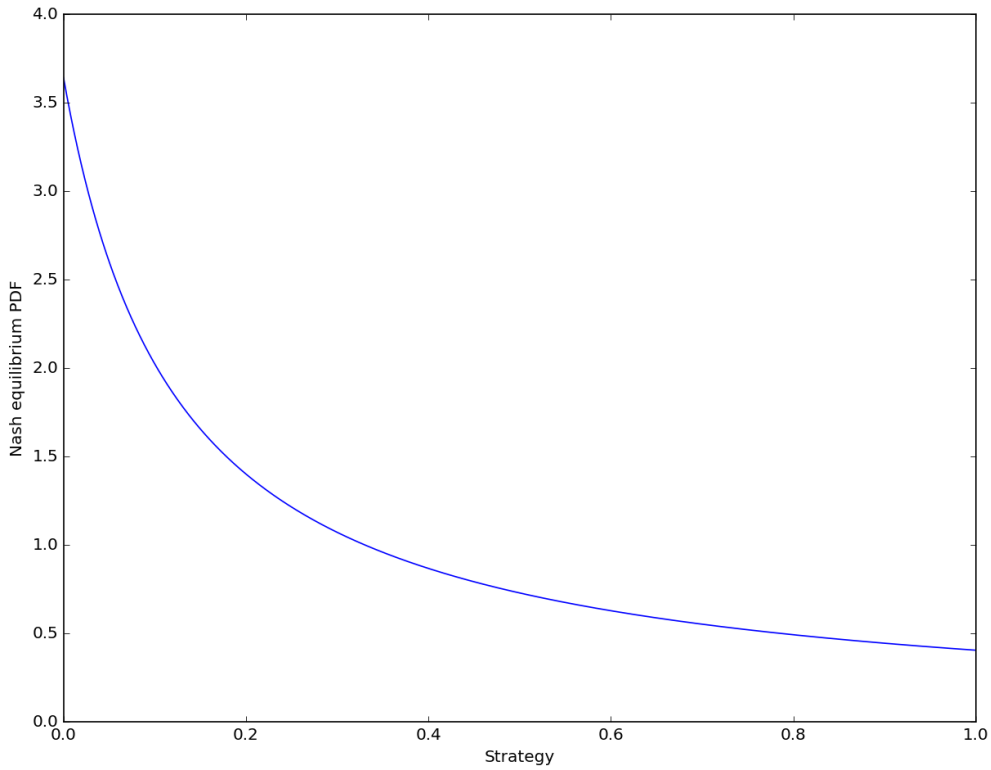

Figure 2: Nash equilibrium for the one-up game with $b = 9$ and handicaps that grow linearly with respect to the strategy $h(x) = 1 + (b-1)x$.

strategy with probability p: the player with advantage plays the "trump card" a, and the player with disadvantage "folds" and plays 0. Otherwise, with probability $1-p$, each player plays uniformly at random in $[0, a)$. We now proceed to find p and prove that this is indeed a Nash equilibrium.

Minimizer advantage. The expected payoff if the maximizer plays x against such a strategy is

$$\underbrace{\frac{1-p}{a}\left(\underbrace{b\int_0^x b^{y-x}dy}_{\text{maximizer greater}} + \underbrace{\int_x^a b^{y-x}dy}_{\text{minimizer greater}}\right)}_{\text{minimizer plays uniformly}} + \underbrace{pb^{a-x}}_{\text{minimizer plays } a} \tag{8}$$

$$\tag{9}$$

At Nash equilibrium this should be constant across its support. Differentiating with respect to x and setting to 0 yields

$$0 = \frac{1-p}{a}\left(b^{1-x} - b^{a-x}\right) - pb^{a-x}\ln b \tag{10}$$

Multiplying through by ab^x removes the dependence on x, as desired for a Nash equilibrium:

$$0 = (1-p)\left(b - b^a\right) - pab^a \ln b \tag{11}$$

$$p = \frac{b - b^a}{b - b^a + ab^a \ln b} \tag{12}$$

The expected (maximizer) payoff if the minimizer plays y against this is

$$\frac{1-p}{a}\left(\int_0^y b^{y-x}dx + b\int_y^a b^{y-x}dx\right) + pb^y \tag{13}$$

$$\tag{14}$$

Differentiating with respect to y and setting to 0 yields

$$0 = \frac{1-p}{a}\left(b^y - b^{1+y-a}\right) + pb^y \ln b \tag{15}$$

Multiplying through by $-ab^{a-y}$ gives

$$0 = (1-p)\left(-b^a + b\right) - pab^a \ln b \tag{16}$$

exactly the same.
Substituting this back in, we have expected payoff

$$\frac{b^a\left(b-1\right)}{b - b^a + ab^a \ln b} \tag{17}$$

Maximizer advantage. This is the same except the final term changes to pbb^{-x} and pbb^{y-a} for the maximizer and minimizer respectively, resulting in

$$p = \frac{b - b^a}{b - b^a + ab\ln b} \tag{18}$$

and expected payoff

$$\frac{b\left(b-1\right)}{b - b^a + ab\ln b} \tag{19}$$

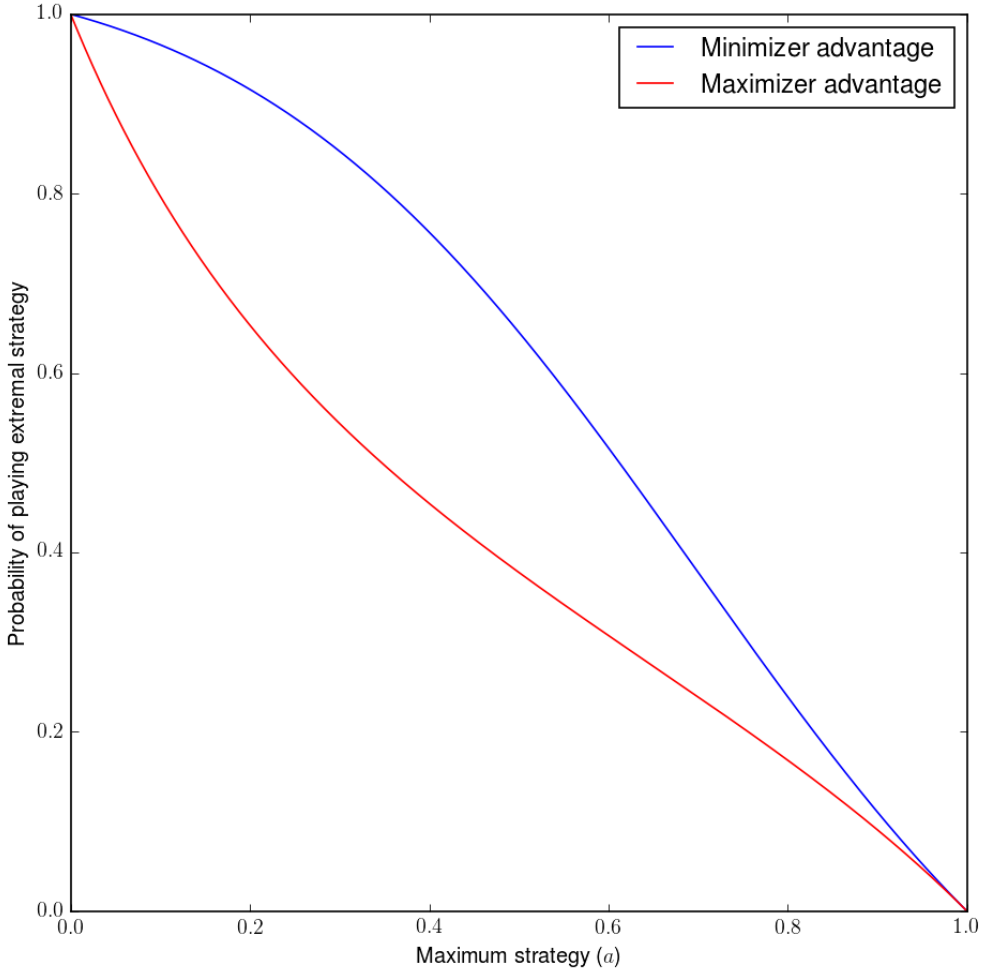

Figure 3: Probability at Nash equilibrium of playing an extremal strategy, i.e. the maximum strategy a for the player with advantage, or the minimum strategy 0 for the player with disadvantage. This plot is for maximum payoff $b = 9$.
For $a \neq 0, 1$ the probability of extremal strategy is not symmetric with respect to which player has the advantage. Namely, for a given a, extremal strategies are played more often when the minimizer has the advantage than when the maximizer has the advantage. This gap increases with b. In either case the Nash equilibrium becomes more uniform (i.e. p decreases) as a increases.

Commentary. These are plotted for $b = 9$ in Figures 3 and 4, with further commentary in the captions. (Imagine e.g. a jump from a 10% hit chance to a 90% hit chance.)

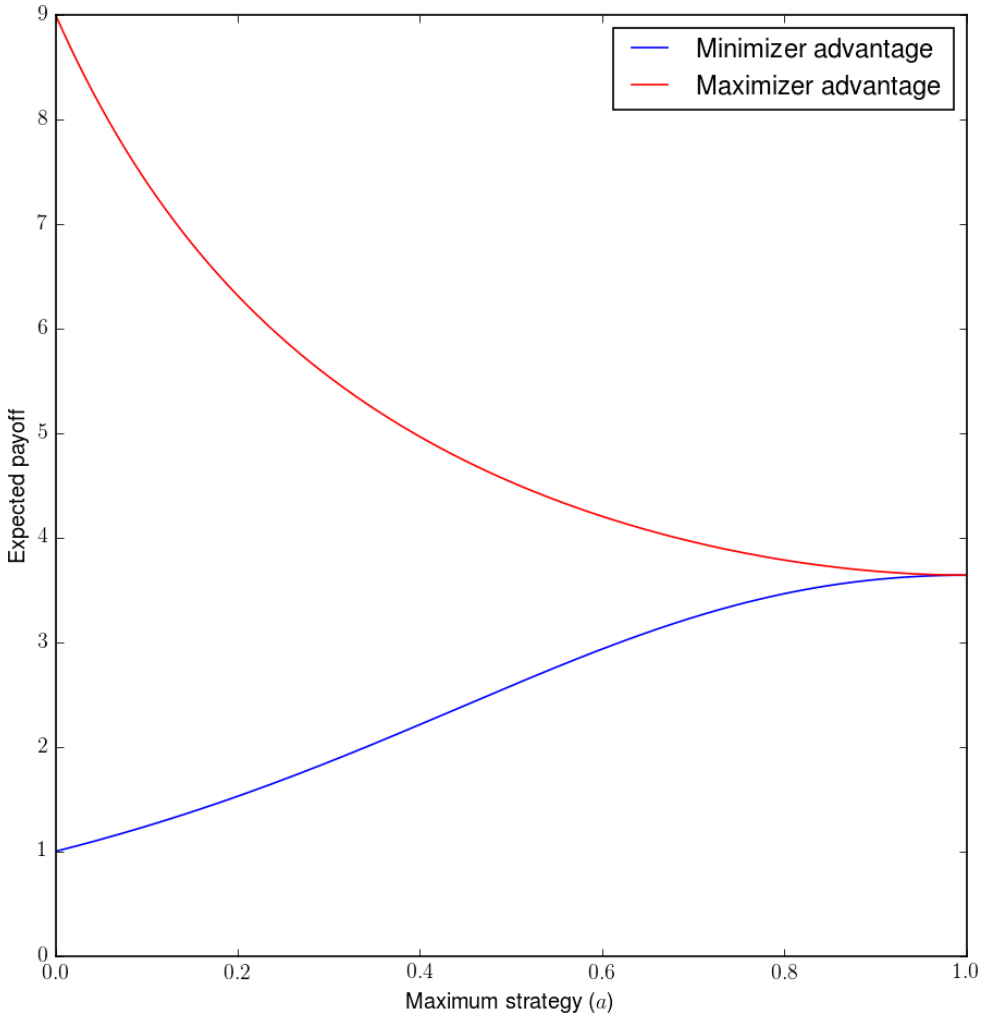

Figure 4: Expected payoff at Nash equilibrium in the one-up game. This plot is again for maximum payoff $b = 9$.
Having the advantage matters less as a increases, as shown by the difference in expected payoffs for a given a depending on whether the minimizer or the maximizer has the advantage.

References

[1] Albert Liu and Steve Marschner. Balancing zero-sum games with one variable per strategy. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, October 2017.