Ultimate X Bonus Streak Analysis

Gary J. Koehler

John B. Higdon Eminent Scholar, Emeritus
Department of Information Systems and Operations Management, 351 BUS, The Warrington College of Business, University of Florida, Gainesville, FL 32611, (koehler@ufl.edu).

This paper extends an analysis of Ultimate X Video Poker to a new variation on its theme. Instead of an outcome generating an immediate return plus establishing a multiplier of the next round's return, in Bonus Streak a set of multipliers is established for subsequent hands. This paper analyzes this new type of game.

Key words: Gambling, non-discounted Markov Decision Problem, Video Poker, Ultimate X.

January, 2017
Revised March, 2017
Revised September, 2018 (added vulturing section)
Revised January, 2019 (more added to vulturing section)

1. Introduction

We refer the reader to our earlier paper analyzing Ultimate X[1] Poker [2] for basic concepts. Ultimate X Bonus Steak alters the basic idea of Ultimate X Poker by offering a stream of multipliers (a streak) for different outcomes to be applied to subsequent hands of play, not just a single multiplier for the next hand as in the original Ultimate X games. Like Ultimate X, this game costs twice the normal underlying game's maximum bet amount to activate the Bonus Streak (e.g., the normal maximal bet amount is 5 coins per line in Jacks or Better). That is, it cost 10 coins per line in Ultimate X. As is usual for multi-line games, each game starts with the same hand dealt to all lines of play and the held cards apply to each line. The outcomes come from independent draws from decks with the cards of the initial hand removed.

Table 1 shows per coin payouts (based on the initial 5 coins) and multiplier streaks for each possible outcome for a Deuces Wild game. For example, if on a line of play the current multiplier is 1 and one gets a Straight Flush then he will be paid 65 coins (5 times the outcome payout of 13). The "5" is because we are showing payouts on a per-coin bet basis and 5 coins were bet (the additional 5 coins wagered were to enable the bonus streak feature). This win sets up a streak so the next hand's multiplier will be 2, the subsequent 4 and so forth. However, if when in the midst of using a streak's multipliers, the player gets an outcome with another non-unit streak, then the current streak's remaining multipliers are changed to multipliers of 12.

Outcome	Per Coin Payout	Streak
Royal Straight Flush	**800**	**2,4,7,10,12**
Four Deuces	**200**	**2,4,7,10,12**
Wild Royal Straight Flush	**25**	**2,4,7,10,12**
Five of a Kind	**16**	**2,4,7,10,12**
Straight Flush	**13**	**2,4,7,10,12**
Four of a Kind (4K)	**4**	**2,2,4**
Full House (FH)	**3**	**2,2,4**
Flush	**2**	**2,2,4**
Straight	**2**	**1**
Three of a Kind (3K)	**1**	**1**
Nothing	**0**	**1**

Table 1: Ultimate X Bonus Steak Multiplies, Deuces Wild

[1] Both Ultimate X and Ultimate X Bonus Streak were created by IGT (https://www.igt.com/) and are offered in their video poker machines.

For example, suppose there is just one multiplier in place in the current streak for a line of play. Then let's track what happens with the following sequence of hands and outcomes shown in Table 2. The first hand results in a Three of a Kind and the payout is multiplied by the Outcome Multiplier of 1. The new streak is just "1". The Straight Flush with Hand 2 sets up a streak of future multipliers (2,4,7,10,12). We see these successively applied in the next two hands. However, the Full House outcome at Hand 4 would normally establish a streak of 2,2,4 but since we already have a streak longer than one element, the current remaining streak (7,10,12) is changed to all 12 multipliers (i.e., to 12,12,12).

Hand	Starting Streak	Outcome	Outcome Multiplier	New Streak
1	1	Three of Kind	1	1
2	1	Straight Flush	1	2,4,7,10,12
3	2,4,7,10,12	Nothing	2	4,7,10,12
4	4,7,10,12	Full House	4	12,12,12
5	12,12,12	Three of Kind	12	12,12
6	12,12	Nothing	12	12
7	12	Nothing	12	1
8	1	Nothing	1	1

Table 2: Example of Multiplier Evolution

Table 3 shows the possible streaks one might see at the start of a hand.

Streak	Streak Values
1	1
2	4
3	12
4	2,4
5	10,12
6	12,12
7	2,2,4
8	7,10,12
9	12,12,12
10	4,7,10,12
11	12,12,12,12
12	2,4,7,10,12

Table 3: Possible Observable Multiplier Streaks

2. Expected Value Analysis

Let M be the set of possible starting multiplier streaks. For example, for the streaks in Table 3 we have

$$M = \left\{ \begin{array}{l} (1),(2,2,4),(2,4),(4),(12,12),(12),(2,4,7,10,12), \\ (4,7,10,12),(7,10,12),(10,12),(12,12,12,12),(12,12,12) \end{array} \right\}.$$

Likewise, let Ω be the set of permutations of the elements of M taken L (the number of lines) at a time with repetition. So for a 3-Line game, each $\pi \in \Omega$ looks like $\pi = (\pi_1, \pi_2, \pi_3)$ where $\pi_i \in M$ and the j$^{\text{th}}$ multiplier of π_i is $\pi_i(j)$. Ω gives all of the possible streak states a player might see for the L lines before starting a hand of play.

Technically, the starting state of each round of play is (π, H) where $\pi \in \Omega$ results from the previous hands' outcomes and $H \in \mathbb{H}$ is a randomly generated next hand and $\mathbb{H}$ is the set of all possible starting hands. Since the outcome of any action depends on just (π, H) and what a decision maker chooses to hold in H, and not the history leading one to this state, the Markov property holds and the resulting problem is a Markov Decision problem[2]. This is not to say that all states can be reached in one step as was the case with the Ultimate X game in [2]. For example, for a one-line game, if the starting state is $((2,2,4), H)$, the only states that could be reached are $((2,4),*)$ and $((12,12),*)$. That is, the only realizable ending streaks are $(2,4)$ and $(12,12)$.

As in [2], we choose to study the non-discounted stream of returns and, for practical matters, assume the horizon is infinite. Thus we focus on solving the infinite horizon, non-discounted, Markov Decision problem (ndMDP) which is represented by

$$v_\pi + g = \sum_{H \in \mathbb{H}} P_H \max_i \left(\sum_{l=1}^{L} \pi_l(1) R_{H_i} + \sum_{\gamma \in \Omega} P_{\pi,\gamma}(H_i) v_\gamma \right) \qquad \pi \in \Omega$$

$$\sum_{\pi} P_\pi v_\pi = 0$$

[2] The associated Markov chains are readily shown to be ergodic.

Here g is the maximal gain per round of play, v_π is the relative bias for state $\pi \in \Omega$, P_π is the steady-state probability of being in state π (before a hand is dealt) under optimal decisions, and P_H is the probability of being dealt hand H. Note that $g/(2L)$ is the optimal expected return per bet unit for the game, the value we wish to compute. The "2" comes from the game costing twice the normal amount on which the payouts are based. For each hand, one must decide which of the possible $i = 1, \ldots, 32$ ways to hold subsets of H, designated by H_i. Each possible decision results in an expected outcome for the hand, R_{H_i}, and a probability of transitioning to state γ of $P_{\pi,\gamma}(H_i)$. Note that in the formulation above, we have reduced the starting state from (π, H) to π by averaging out the impact of the random starting hand (hence the $\sum_{H \in \mathbb{H}}$).

Since R_{H_i} is independent of the multipliers, let $m(\pi) = \sum_{l=1}^{L} \pi_l(1)$ and we can rewrite the problem as

$$\begin{aligned} &v_\pi + g = \sum_{H \in \mathbb{H}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma \in \Omega} P_{\pi,\gamma}(H_i) v_\gamma \right) \qquad \pi \in \Omega \\ &\sum_\pi P_\pi v_\pi = 0 \end{aligned} \tag{1}$$

Consider $P_{\pi,\gamma}(H_i)$. This is the probability of starting in state π and transitioning to state γ. This depends on which cards in H are held (designated by decision i leading to holding H_i) and the various possible outcomes (Straight, Flush, etc.) afterwards. Let $\mathbb{O}$ be the set of possible outcomes and $P_o(H_i \mid H)$ be the probability of outcome $o \in \mathbb{O}$ when cards H_i are held from hand H. For each outcome there is a payout and a streak (see Table 1 for example). The resulting streak is a function of the starting streak and the outcome represented by $s(o, \pi_l)$. Note, for regular Ultimate X, $s(o, \pi_l)$ is independent of π, it depends only on the hand's outcome. States in Bonus Streak having only single-length streaks also exhibit this property.

For example in a 2-Line game, if the starting state has $\pi = \left((2,2,4),(1)\right)$ the possible resulting streaks are

$$(2,2,4) \to \begin{cases} (2,4) & o \in \{Straight, 3K, Nothing\} \\ (12,12) & otherwise \end{cases}$$

$$(1) \to \begin{cases} (1) & o \in \{Straight, 3K, Nothing\} \\ (2,2,4) & o \in \{4K, FH, Flush\} \\ (2,4,7,10,12) & otherwise \end{cases}$$

Thus

$$P_{\pi,\gamma_l}(H_i) \equiv \sum_{\substack{o \in \mathbb{O} \\ \gamma_l = s(o,\pi_l)}} P(o \mid H_i)$$

$$P_{\pi,\gamma}(H_i) = \prod_{l=1}^{L} P_{\pi,\gamma_l}(H_i) = \prod_{l=1}^{L} \sum_{\substack{o \in \mathbb{O} \\ \gamma_l = s(o,\pi_l)}} P(o \mid H_i)$$

and $P_{\pi,\gamma_l}(H_i)$ is the probability of outcomes having an associated multiplier of γ_l given one starts in state (π, H) and chooses to hold H_i. As in [2], we can iteratively solve (1) by

$$v_\pi^{n+1} + g^{n+1} = e_\pi^{n+1} = \sum_{H \in \mathbb{H}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma \in \Omega} P_{\pi,\gamma}(H_i) v_\gamma^n \right) \qquad \pi \in \Omega \tag{2}$$

$$g^{n+1} = \sum_{\pi} P_\pi^{n+1} e_\pi^{n+1}$$

$$P_\pi^{n+1} = \sum_{\gamma \in \Omega} P_\gamma^n P_{\gamma,\pi}(H_i^*) \qquad \pi \in \Omega$$

The term $P_{\gamma,\pi}(H_i^*)$ stands for the value of $P_{\pi,\gamma}(H_i)$ with an optimal decision i.

As discussed in [2], the number of permutations (with repetition) of $|M|$ things L at a time is $|M|^L$, so a 10-Line version of Ultimate X Bonus Streak with the multipliers shown in Table1 has $12^{10} = 61{,}917{,}364{,}224$ multiplier patterns a player may see. So the true number of states is $\binom{n}{5}|M|^L$ where n is the size of the deck of cards used (assuming order of the cards is not important). For example, for decks of 52 cards and a 10-line game, the number of states is on the order of 10^{17}, over 100 quadrillion.

Fortunately, some of the problem size reductions discussed in [2] can be used in the Bonus Streak game. In particular, the reductions are:

1. Use equivalent suite permutations of hands to reduce $H \in \mathbb{H}$ to unique hands $H \in \bar{\bar{\mathbb{H}}}$. This is easily implemented by letting P_H reflect the number of different suite permutations for a given hand. For games with 52 cards, this reduces the size of $\mathbb{H}$ from 2,598,960 hands to 134,459 in $\bar{\bar{\mathbb{H}}}$.
2. Use state permutations to reduce the state space. For example, in a 3-Line game, state $\{(1),(2,4),(12,12)\}$ will give the same expected payouts as state $\{(2,4),(1),(12,12)\}$ and state $\{(1),(12,12),(2,4)\}$ since the order of the multipliers across the lines of play is not important. As in [2] we let $C \subseteq \Omega$ contain just the unique combinations (say those in sorted order) and denote equivalent states $\gamma \approx \pi$ in Ω for each $\gamma \in C$.

Unfortunately, a third reduction in [2] first suggested by Michael Shackelford [4] is not valid here. That reduction stated that all states having the same value of $m(\pi)$ are equivalent. The proof given in [2] relied on the fact that $P_{\pi,\gamma_1}(H_i)$ was independent of π which is not the case with Bonus Streak unless the states are composed of single-length streaks.

Let $C \subseteq \Omega$ contain just the unique combinations (say those in sorted order). So

$$|C| = \binom{|M|+L-1}{|M|-1}.$$

With the reductions, we wish to solve

$$v_\pi^{n+1} + g^{n+1} = e_\pi^{n+1} = \sum_{H \in \bar{\bar{\mathbb{H}}}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma \in C} P_{\pi,\gamma}(H_i) v_\gamma^n \right) \qquad \pi \in C \tag{3}$$

$$g^{n+1} = \sum_\pi P_\pi^{n+1} e_\pi^{n+1}$$

$$P_\pi^{n+1} = \sum_{\gamma \in \Omega} P_\gamma^n P_{\gamma,\pi}\left(H_{i \in S_{\gamma,H}}\right) \qquad \pi \in \Omega$$

With the reductions, we need to adjust our definition of $P_{\pi,\gamma}(H_i)$. Let

$$P_{\pi,\gamma}\left(H_i\right)=\prod_{l=1}^{L}\sum_{\substack{\eta\in\Omega\\ \eta\approx\gamma}}P_{\pi,\eta_l}\left(H_i\right)=\prod_{l=1}^{L}\sum_{\substack{\eta\in\Omega\\ \eta\approx\gamma}}\sum_{\substack{o\in\mathbb{O}\\ \eta_l=s(o,\pi_l)}}P\left(o\mid H_i\right)\qquad \pi,\gamma\in\Omega$$

Note, the original values are $v_{\gamma}^{n+1}=v_{\pi}^{n+1}$ for $\gamma\in\Omega/C,\gamma\approx\pi$. As in [2], we stop (3) when

$$\left|g^{n+1}-g^{n}\right|+\sum_{\pi\in C}\left|v_{\pi}^{n+1}-v_{\pi}^{n}\right|+\sum_{\pi\in C}\left|P_{\pi}^{n+1}-P_{\pi}^{n}\right|<10^{-10}\left|C\right|. \tag{4}$$

We solved a hypothetical[3] 1-Line version of Deuces Wild in Table 1 to get a gain (g) of 1.94665 and steady state values shown in Table 4. The Expected Value (EV) is 1.94665/2 = 0.973325.

Deuces Wild – 1 Line		
π	v_{π}	P_{π}
1	-2.506	0.680794
4	0.363605	0.069155
12	8.08679	0.031268
2,4	1.30152	0.079446
10,12	15.8166	0.006048
12,12	17.7519	0.014929
2,2,4	3.38067	0.091454
7,10,12	20.8671	0.006858
12,12,12	27.4171	0.002111
4,7,10,12	23.5686	0.007795
12,12,12,12	37.0822	0.001174
2,4,7,10,12	25.2629	0.008969

Table 4: Solution to one line version of the game with multiples in Table 2

Table 5 gives the outcomes for the 1-3 Line versions of this Deuces Wild game. Actual machines in casinos currently only offer 3, 5 and 10-Line versions, so the 1-Line and 2-Line versions are hypothetical.

Deuces Wild Video Poker	**g**	**EV**
1-Line	**1.94665**	**0.973325**
2-Lines	**3.88404**	**0.971010**
3-Lines	**5.81832**	**0.969721**

Table 5: Optimal expected returns for Deuces Wild Ultimate X Bonus Streak.

[3] Although we have not seen a 1-Line version of the game, we anticipate their introduction just as 1-Line games of Ultimate X were eventually released by IGT.

Interestingly, the Bonus Streak game appears to exhibit the same phenomenon that the Ultimate X games showed (Page 16, [2]):

"the impact on expected return as the number of lines increases is negative"

Note the EVs reduce as the number of lines increase in Table 5.

As another example, Table 6 gives the payouts and streaks for 7-5 Bonus Poker Deluxe.

Outcome	Payout	Streak
Royal Straight Flush	**800**	**2,5,8,10,12**
Straight Flush	**50**	**2,5,8,10,12**
Four of a Kind (4K)	**80**	**2,5,8,10,12**
Full House (FH)	**7**	**2,5,8,10,12**
Flush	**5**	**2,5,8**
Straight	**4**	**2,5**
Three of a Kind (3K)	**3**	**2,5**
Two Pair	**1**	**1**
Jacks or Better Pair	**1**	**1**
Nothing	**0**	**1**

Table 6: Ultimate X Bonus Steak Multiplies, Bonus Poker Deluxe

Table 7 gives the outcomes for the 1-3 Line versions of Bonus Poker Deluxe and Table 8 its steady-state values for 1-Line.

Bonus Deluxe	g	EV
1-Line	**1.93818**	**0.969092**
2-Lines	**3.86879**	**0.967198**
3-Lines	**5.79847**	**0.966412**

Table 7: Optimal expected returns for Bonus Poker Deluxe Ultimate X Bonus Streak.

Bonus Deluxe – 1 Line		
π	v_π	P_π
1	-2.43554	0.750595
5	1.3749	0.064623
8	4.2542	0.011538
12	8.10293	0.023618
2,5	2.14009	0.073308
5,8	7.56512	0.013025
10,12	15.79	0.008086
12,12	17.7151	0.00536
2,5,8	8.7366	0.014784
8,10,12	21.7695	0.009117
12,12,12	27.3272	0.002571
5,8,10,12	25.2758	0.010296
12,12,12,12	36.9393	0.001391
2,5,8,10,12	26.6274	0.011688

Table 8: Optimal relative biases and steady state probabilities for Bonus Deluxe.

The challenge with analyzing games beyond 3-Lines is easily seen in Table 9 where we show the sizes of the states for the Deuces Wild game of Table 1.

	1-Line	3-Lines	5-Lines	10-Lines
$\lvert\Omega\rvert = \lvert M\rvert^L$	12	1,728	248,832	61,917,364,224
$\lvert C\rvert = \binom{\lvert M\rvert + L - 1}{L}$	12	364	4,368	352,716

Table 9: Size of Sets for Ultimate X Bonus Streak Deuces Wild

For example, using the state reduction to C for a 10-Line game gives 352,715 states. For each state we need to find the optimal hold of 134,459 hands, each requiring 32 probability vectors and expected value calculations. That is, over 1.5 trillion calculations for each are needed at each iteration in (3). With Ultimate X, the third state size reduction (which is not generally applicable here) to set D (in [2]) reduced the state space size dramatically. For the Deuces Wild game examined in [2], the sizes are as shown in Table 10. Notice that the 10-Line Ultimate X game was easier to solve than the 3-Line game of Bonus Streak Ultimate X.

	1-Line	3-Lines	5-Lines	10-Lines
$\lvert\Omega\rvert = \lvert M\rvert^L$	7	343	16,807	282,475,249
$\lvert C\rvert = \binom{\lvert M\rvert + L - 1}{L}$	7	84	462	8,008
$\lvert D\rvert$	7	29	51	106

Table 10: Size of Sets in [2] for Ultimate X Deuces Wild

In short, without some massively parallel computing platform, some new insights are needed to solve the Bonus Streak versions of Ultimate X for 10-Line games. 5-Line games are within reach but will take weeks to solve.

3. Possible Speed-ups

Some obvious computational speed-ups include precomputing the following values which don't change from iteration to iteration:

1. $P_H m(\pi) R_{H_i} \qquad \forall H \in \bar{\mathbb{H}}, i = 1, \ldots, 32$
2. $P_H P_{\pi,\gamma}(H_i) \qquad \forall H \in \bar{\mathbb{H}}, \pi, \gamma \in C, i = 1, \ldots, 32$

The second suggestion above may be impractical because C grows so fast and $\bar{\mathbb{H}}$ is large.

Similarly, dividing the iterations to parallel computations over $\bar{\mathbb{H}}$ and C are easily done. With most processors implementing multiple cores and hyper-threading, parallel computing is possible[4].

As mentioned when discussing state-space reductions, it was noted we can have a small reduction of states by collapsing those states having all single-length streaks and equal $m(\pi)$ values. The impact is minimal, however. For example, in the Jacks or Better game shown in Section 4 below, the 3-Line game has 560 states in C and only 17 can be reduced using this

[4] We used 10 of our 12 cores on a Xeon E5645 Intel processor.

equivalence. The overhead to implement this reduction hardly covers the slight reduction in state space size.

Another possible speed-up can be achieved using a termination criterion first suggested by Odoni [3]. He showed that

$$\overline{L}^n \geq \overline{L}^{n+1} \geq g \geq \underline{L}^{n+1} \geq \underline{L}^n$$

$$\overline{L}^n = \max_{\pi} e_{\pi}^{n+1} - v_{\pi}^n$$

$$\underline{L}^n = \min_{\pi} e_{\pi}^{n+1} - v_{\pi}^n$$

So, stopping when $\overline{L}^{n+1} - \underline{L}^{n+1} < \varepsilon$ will provide a good estimate of g for small enough ε. For examples, for the first Jacks or better game shown later using ε values shown in the Table below, we found the following number of iterations needed to achieve the stopping condition:

Lines	**1**	**2**	**3**
Iterations with Condition (4)	**28**	**29**	**30**
Iterations with $\varepsilon = 10^{-8}$	**25**	**25**	**26**
Iterations with $\varepsilon = 10^{-7}$	**22**	**24**	**24**
Iterations with $\varepsilon = 10^{-6}$	**21**	**21**	**21**

This stopping criterion may not leave us with as accurate estimates of the steady state probabilities or relative bias values as the stopping criterion discussed earlier with Equation (4), but it could save iteration rounds if we are interested in just computing the gain of a game.

In [2] we discussed some additional computational reductions. One was to use other forms of iteration where both storage requirements and rate of convergence improved when applicable. Such methods exist for solving discounted, infinite-horizon, Markov Decision problems. However, we know of no way to implement these for the non-discounted problem without first converting it to a form where they can be applied (as done by Koehler et al. in [1]) which itself required solving a Markov decision problem.

We also mentioned it is possible to permanently eliminate sub-optimal decisions as the iteration proceeds, thus, in principle, reducing the problem size. In our explorations of this approach, the overhead introduced did not justify the improvement in convergence speed.

4. Results

Below are the results we found for a selection of games, pay tables and bonus streaks for 1-Line and 3-Line versions of the game.

Game	Pays	Streak			EVs		
		3K	STR, FL	HIGHER	Regular	1-Line	3-Line
Double Double Bonus	9-5	2,4	2,4,8	2,4,8,10,12	0.978729	0.987373	0.984091
Double Double Bonus	8-5	2,4	2,4,8	2,4,8,10,12	0.967861	0.976225	0.972826
Double Double Bonus	7-5	2,4	2,4,8	2,4,8,10,12	0.957120	0.965194	0.961655
Double Double Bonus	6-5	2,4	2,4,8	2,4,8,10,12	0.946569	0.954333	0.950554
Triple Double Bonus	9-6	2,4	2,4,8	2,4,8,10,12	0.981540	0.993189	0.990460
Triple Double Bonus	9-5	2,4	2,4,8	2,4,8,10,12	0.970204	0.978236	0.974948
Triple Double Bonus	8-5	2,4	2,4,8	2,4,8,10,12	0.959687	0.967222	0.963846
Triple Double Bonus	7-5	2,4	2,4,8	2,4,8,10,12	0.949178	0.956277	0.952820
Double Bonus	9-6-5	2,4	2,4,7	2,4,7,11,12	0.978062	0.982587	0.980935
Double Bonus	9-6-4	2,4	2,4,8	2,4,8,10,12	0.963754	0.976847	0.974719
Double Bonus	9-5-4	2,4	2,4,8	2,4,8,10,12	0.952738	0.962197	0.959335
Double Bonus	8-5-4	2,4	2,4,8	2,4,8,10,12	0.941897	0.950919	0.947926
Bonus Poker	7-5	2,4	2,4,8	2,4,8,10,12	0.980147	0.987757	0.984631
Bonus Poker	6-5	2,4	2,4,8	2,4,8,10,12	0.968687	0.976217	0.973129
Jacks or Better	9-5	2,4	2,4,8	2,4,8,10,12	0.984498	0.992208	0.989064
Jacks or Better	8-5	2,4	2,4,8	2,4,8,10,12	0.972984	0.980650	0.977559
Jacks or Better	7-5	2,4	2,4,8	2,4,8,10,12	0.961472	0.969092	0.966057
Jacks or Better	6-5	2,4	2,4,8	2,4,8,10,12	0.949961	0.957538	0.954556
		3K STR	FLUSH	HIGHER			
Bonus Poker Deluxe	8-6	2,5	2,5,7	2,5,7,11,12	0.984928	0.995215	0.987909
Bonus Poker Deluxe	8-5	2,5	2,5,8	2,5,8,10,12	0.974009	0.980375	0.977853
Bonus Poker Deluxe	7-5	2,5	2,5,8	2,5,8,10,12	0.962526	0.969092	0.966412
Bonus Poker Deluxe	6-5	2,5	2,5,8	2,5,8,10,12	0.953611	0.958301	0.955063
		FL, FH, 4K	HIGHER				
Deuces Wild	20-12-10-4-4-3	2,2,4	2,4,4,11,12	n/a	0.975791	0.984147	0.981442
Deuces Wild	20-12-10-4-4-3	2,2,4	2,4,4,10,12	n/a	0.975791	0.981346	0.978687
Deuces Wild	25-16-13-4-3-2	2,2,4	2,4,7,10,12	n/a	0.967651	0.973327	0.969721
Deuces Wild	20-10-8-4-4-3	2,2,4	2,4,5,10,12	n/a	0.959638	0.966627	0.963748
Deuces Wild	25-15-10-4-3-2	2,2,4	2,4,8,10,12	n/a	0.948182	0.954816	0.950898
Bonus Deuces Wild	10-4-3-3	2,2,4	2,4,5,10,12	n/a	0.973644	0.988510	0.983837
Bonus Deuces Wild	12-4-3-2	2,2,4	2,4,8,10,12	n/a	0.962183	0.975882	0.971701
Bonus Deuces Wild	12-4-3-2	2,2,4	2,4,6,10,12	n/a	0.962183	0.969432	0.965357
Bonus Deuces Wild	10-4-3-2	2,2,4	2,4,6,10,12	n/a	0.953368	0.958227	0.953696

Tables from actual 3- and 5-Line Games

Game	Pays	Streak		EVs		
		3K, STR, FL	HIGHER	Regular	1-Line	3-Line
Double Double Bonus	9-5	2,3,4	2,3,4,8,12	0.978729	0.986802	0.984845
Double Double Bonus	8-5	2,3,4	2,3,4,8,12	0.967861	0.975780	0.973698
Triple Double Bonus	9-6	2,3,4	2,3,4,8,12	0.981540	0.991724	0.990154
Triple Double Bonus	9-5	2,3,4	2,3,4,8,12	0.970204	0.977546	0.975763
Double Bonus	9-6-5	2,3,4	2,3,4,8,12	0.978062	0.991708	0.990513
Double Bonus	9-6-4	2,3,4	2,3,4,8,12	0.963754	0.975194	0.973749
Bonus Poker	7-5	2,3,4	2,3,4,8,12	0.980147	0.986813	0.985326
Bonus Poker	6-5	2,3,4	2,3,4,8,12	0.968687	0.975253	0.973772
Jacks or Better	8-6	2,3,4	2,3,4,7,12	0.983927	0.989636	0.988314
Jacks or Better	8-5	2,3,4	2,3,4,8,12	0.972984	0.979645	0.978164
Bonus Poker Deluxe	8-5	2,3,4	2,3,4,8,12	0.974009	0.983523	0.981652
Bonus Poker Deluxe	7-5	2,3,4	2,3,4,8,12	0.962526	0.972368	0.970263
		FL, FH, 4K	HIGHER			
Deuces Wild	20-12-9-4-4-3	2,2,4	2,4,8,10,12	0.970554	0.988058	0.985117
Deuces Wild	25-16-13-4-3-2	2,2,4	2,4,8,10,12	0.967651	0.976643	0.972999
Bonus Deuces Wild	10-4-3-3	2,2,4	2,4,7,10,12	0.973644	0.993026	0.990081
Bonus Deuces Wild	12-4-3-2	2,2,4	2,4,8,10,12	0.962183	0.975882	0.971701

Tables from actual 10-Line Games

5. Vulturing

Vulturing refers to the process of scavenging left-over multipliers from previous players. For Ultimate-X games, if there are any multipliers greater than one, the expected value of playing a hand at a 5 coin bet is positive. In Bonus Streak games, that strategy doesn't work because the multipliers are disabled for any bet less than the maximum bet. However, the left-over multipliers may still lead to a positive expected value.

Before proceeding further, we need a firm definition of what vulturing might encompass for any game leaving something of potential player value. In short, vulturing dictates in what states we should play a game, and how, and when we should stop playing the game because it is no longer of any advantage to us. We distinguish this from non-vulturing where the decisions to play a game and when to stop are dictated by other concerns such as game likes and dislikes, how much money one has with them, time available, and so on.

States: For both Ultimate X and Bonus Streak, the "states" are defined by the current multipliers and, for Ultimate X, the state can be simplified to just a sum of the multipliers. Which states should be played and which not played has to be determined.

How to play: There are several issues associated with how to play a game. The first is the bet size. For both Ultimate X and Bonus Streak there are three possibilities to consider. As with most video poker games, there is often a distinction between playing 1-4 coins versus five coins, usually in how much a Royal Straight Flush pays. Five coins almost always pays-out more per coin-in than does 1-4 coins. No serious player plays less than 5 coins per bet in such cases. For both Ultimate X and Bonus Streak, playing 10 coins per line activates the multiplier generation feature. That is, the outcome of the hand generates multipliers that are applied to future hands. Unlike Ultimate X, Bonus Streak multipliers do not apply to future hands unless they too are played at the maximum (10 coin per line) level.

The second issue, once a bet size is chosen, is how to play a dealt hand? This has a lot to do with what we do at the end of the hand of play. If we have placed a maximum bet (10 coins per line), we should consider the impact our play has on potential new multipliers. Unlike non-vulturing players, we do not have to endure unattractive states (we simply stop playing) or we can take actions that increase our short-term results at the expense of long-term ones. In Ultimate X, we can use the multipliers on a hand at less than the max bet (5 coins) and play using the optimal strategy for the underlying game.

Stopping: In short, vulturing entails starting at a desirable state and stopping at an undesirable state. What entails a desirable state and undesirable one needs some discussion. One might describe a desirable state as one that has a positive expected value for the next hand of play, and then stopping occurs when a state is encountered where the expected value of the next hand of play is negative. Another might describe one as having a positive expected value across a sequence of hands played until the player decides to stop playing. There may be hands played with negative expected returns for one hand that are covered by hands already having good multipliers.

Suppose one finds the following left-over multipliers in a 9-5 Jacks or Better game: $\{(2,4),(1),(1)\}$. Should one "vulture" this? If only one hand will be played (at a max bet) for this 3-line game, the expected value is 0.984498 * 4/6 = 0.656332. The 0.984498 is the expected value for the normal 9-5 game (since we are playing just one hand). The 4/6 is the average multiplier per coin-in. So this is not attractive. However, since we are playing at the max bet, we have the potential of new streaks and the next state may compensate for the expected loss for the current state (see the examples given shortly). So, like normal play in Ultimate X and Bonus Streak, we must anticipate future hands, even in vulturing.

So which states should we vulture in Bonus Streak? Two conditions might be considered. From a solution to (1) we have one condition. Vulture state π if

C1: $v_\pi + g \geq 2L$.

This rule takes into account future hands. Another condition is also obvious, vulture a state if

C2: $m(\pi)\sum_{H\in\mathbb{H}} P_H \max_i R_{H_i} \geq 2L$.

This latter condition just considers only one hand (played perfectly for the underlying game) and ignores any future possibilities.

It is possible that a state satisfies the second condition without satisfying the first one. For example, in the Deuces Wild game used throughout this paper, the state $\{(1),(4),(2,4)\}$ has

$$m(\pi)\sum_{H\in\mathbb{H}} P_H \max_i R_{H_i} = 7 \bullet 0.9676505 = 6.7735535 > 6$$

but

$$v_\pi + g = \text{-}0.8098685+5.8183247=5.0084562<6 .$$

This means that playing the hand myopically for one hand is better than using perfect play for the regular Bonus Streak game. Of course, one may get lucky with a new set of attractive multipliers.

Likewise, it is possible a state satisfies the first condition without satisfying the second. For example, the state $\{(1),(1),(4,7,10,12)\}$ has

$$m(\pi)\sum_{H\in\mathbb{H}} P_H \max_i R_{H_i} = 6\bullet 0.9676505 = 5.805903 < 6$$

but

$$v_\pi + g = 18.4245963+5.8183247=24.24292>6 .$$

One can see that, although the first hand will be played with a negative expected value, the subsequent three hands will all have a positive expected value.

Strictly speaking, condition (C1) assumes play will follow the normal optimal play for the Bonus Streak game. However, we won't be playing a normal Bonus Streak game but rather one that terminates with unattractive states. Likewise, (C2) assumes we will play the hand myopically, ignoring any future hands (at least until we see the next state which might be good). We propose using C1 where the gain and relative bias values are determined by what we call the Optimal Vulturing problem and C2 only when a state does not satisfy C1 but does satisfy C2.

Here we give a formulation for the Optimal Vulturing problem. For any state $\pi \in C$ let $\delta(\pi)$ be 1 if we should vulture the game under rule C1 in state π and 0 otherwise.

$$\max_\delta g$$

$$v_\pi + g = \sum_{H\in\mathbb{H}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma\in\Omega} P_{\pi,\gamma}(H_i) v_\gamma \delta(\gamma) \right) \qquad \pi \in \Omega \tag{4}$$

$$\sum_\pi P_\pi v_\pi = 0$$

These gain and bias values may be different from those determined by (1).

Define the following for a fixed set $\Delta \subseteq C$.

$$v_\pi + g = \sum_{H\in\mathbb{H}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma\in\Omega} P_{\pi,\gamma}(H_i) v_\gamma \delta(\gamma) \right) \qquad \pi \in \Omega$$

$$\sum_\pi P_\pi v_\pi = 0 \qquad (\Delta)$$

$$\delta(\pi) = \begin{cases} 0 & \pi \notin \Delta \\ 1 & \pi \in \Delta \end{cases}$$

Theorem 1

Given a solution to (Δ) then for any $\pi \in \overline{\Delta}$

$$v_\pi^*(\Delta) < 0 \Rightarrow \delta(\pi) = 0$$

$$v_\pi^*(\Delta) > 0 \Rightarrow \delta(\pi) = 1$$

Proof:

We have for any state s

$$v_s + g = \sum_{H \in \mathbb{H}} P_H \max_i \left(m(s) R_{H_i} + \sum_{\gamma \in \Omega} P_{\pi,\gamma}(H_i) v_\gamma \delta(\gamma) \right)$$

$$= \sum_{H \in \overline{\mathbb{H}}} P_H \max_i \left(m(s) R_{H_i} + \sum_{\gamma \in \Omega / \{\pi\}} P_{s,\gamma}(H_i) v_\gamma \delta(\gamma) + P_{s,\pi}(H_i) v_\pi \delta(\pi) \right)$$

$$= v_\pi \delta(\pi) \sum_{H \in \overline{\mathbb{H}}} P_H P_{s,\pi}\left(H_{i^*(s)}\right) + \sum_{H \in \overline{\mathbb{H}}} P_H \left(m(s) R_{H_{i^*(s)}} + \sum_{\gamma \in C/\{\pi\}} P_{s,\gamma}\left(H_{i^*(s)}\right) v_\gamma \delta(\gamma) \right)$$

So

$$\sum_s P_s v_s = 0$$

gives

$$g = \left\{ \begin{aligned} & v_\pi \delta(\pi) \sum_s P_s \sum_{H \in \overline{\mathbb{H}}} P_H P_{s,\pi}\left(H_{i^*(s)}\right) \\ & + \sum_s P_s \sum_{H \in \overline{\mathbb{H}}} P_H \left(m(s) R_{H_{i^*(s)}} + \sum_{\gamma \in C/\{\pi\}} P_{s,\gamma}\left(H_{i^*(s)}\right) v_\gamma \delta(\gamma) \right) \end{aligned} \right\}$$

Now, if $v_\pi^* < 0$ and $\delta(\pi) = 1$ or $v_\pi^* > 0$ and $\delta(\pi) = 0$ we reach a contradiction of the optimality of our solution since with all the other decisions held constant, we could achieve a better gain value since

$$\sum_s P_s \sum_{H \in \overline{\mathbb{H}}} P_H P_{s,\pi}\left(H_{i^*(s)}\right) > 0.$$

□

Theorem 1 suggests a greedy algorithm to solve (4). Let $\Delta = C$. Given a solution to (1), if any conditions of Theorem 2 are not satisfied for some $\pi \in \overline{\Delta}$ set

$$\delta(\pi) = \begin{cases} 0 & v_\pi^* < 0 \\ 1 & v_\pi^* \geq 0 \end{cases}$$

Then use the following greedy algorithm:

1. Solve for the steady state values using the following iterative approach:

$$v_\pi^{n+1} + g^{n+1} = \sum_{H \in \mathbb{H}} P_H \max_i \left(m(\pi) R_{H_i} + \sum_{\gamma \in C} P_{\pi,\gamma}(H_i) v_\gamma^n \delta(\gamma) \right) \qquad \pi \in C$$

$$g^{n+1} = \sum_\pi P_\pi^{n+1} e_\pi^{n+1}$$

$$P_\pi^{n+1} = \sum_{\gamma \in \Omega} P_\gamma^n P_{\gamma,\pi}\left(H_{i \in S_{\gamma,H}}\right) \qquad \pi \in \Omega$$

2. If any values for a state $\pi \in \overline{\Delta}$ do not meet the conditions of Theorem 1, pick one and change its δ value and return to Step 1. Otherwise stop.

Since the gain increases with each cycle, the solution monotonically increases until no further opportunities exist. This does not guarantee that the greedy algorithm stops with an optimal solution. However, we have not seen any solutions better than the ones we have found using the greedy algorithm.

Here are the steady-state results for the Deuces Wild game highlighted in this paper.

Deuces Wild Video Poker	**EV**
1-Line	**1.658137**
2-Lines	**1.527503**
3-Lines	**1.513952**

These values are not indicative of one's vulturing EV since no simple scheme dictates what collection of multipliers a person might abandon. The reasons one stops playing a game and leaving unused multipliers are varied and indeterminate and could easily include factors like fatigue, alcohol consumption, financial resources, superstition, other obligations, unacceptable conditions (like an obnoxious player, too cold, too much noise, etc.), and so forth. So not knowing the probability of finding an abandon game state, computing an overall expected value is impossible.

So, assuming we have an optimal solution to (4), we have the following vulturing rules. Vulture a state if

C1: $v_\pi + g \geq 2L$ and play according to optimal decisions using (4) values.

Otherwise, vulture a state if

C2: $m(\pi)\sum_{H\in\mathbb{H}} P_H \max_i R_{H_i} \geq 2L$ and play it myopically – using perfect play for the underlying video poker game.

Here is an example contrasting the use of using an optimal solution to (4) instead of the normal solution to (1) that doesn't take into account a stopping condition. Suppose in our example Deuces Wild game we are at the state $\{(1),(2,4),(2,4)\}$. We have

Using Model (1):

C1: $v_\pi + g =$ 0.1763380+5.8183247=5.9946627<6 so don't vulture (just barely).

However,

Using Model (4):

C1: $v_\pi + g =$ -0.9836034+8.7346019=7.7509984>6 so vulture.

On the surface, the result using (1) seems counter-intuitive. Clearly the first hand would be played with a negative expected value (5.8183247/6 * 5/6) and the second will have a positive expected value (at worst, 5.8183247/6 * 9/6). However, (1) assumes we will keep going so further hands would impact the expected return, and since the overall Bonus Streak game has a negative expected value, the second hand return isn't enough, on average, to overcome the expected long-term play. Model (4) takes into account that we will stop after the second hand unless it gets us to a "good state".

Of some interest is the size of $\overline{\Delta} \equiv \left\{\delta : m(\delta)\sum_{H\in\mathbb{H}} P_H \max_i R_{H_i} < 2L\right\}$. Table 11 shows the sizes of $\overline{\Delta}$ for the Deuces Wild Bonus Streak game used earlier.

	1-Line	3-Lines	5-Lines	10-Lines
$\lvert\Omega\rvert = \lvert M\rvert^L$	12	1,728	248,832	61,917,364,224
$\lvert C\rvert = \binom{\lvert M\rvert + L - 1}{L}$	12	364	4,368	352,716
$\lvert\bar{\Delta}\rvert$	4	22	76	?

Table 11: Size of Sets for Ultimate X Bonus Streak Deuces Wild

6. Summary

This paper presented an analysis of Ultimate X Bonus Streak games. This generalizes the results of Ultimate X games [2] since Ultimate X can be considered as a special case of Ultimate X Bonus Streak. However, Ultimate X can be solved faster using reductions that can't be used with Bonus Streak games.

At the present time, we are unable to solve Bonus Streak games with 10-Lines because the state space is so large. 5-Line games are within reach, but we have not solved them yet. We are working on new insights and algorithmic improvements.

Lastly, various conditions for determining profitable vulturing states were determined.

7. Acknowledgements

We appreciate the many e-mail discussions with Michael Shackleford, The Wizard of Odds. We thank also Rick Percy from Columbus, Ohio who caught a typo in one pay table and an inconsistency in another table where our Streaks didn't match what we used in our code. Thanks go also to Neil Shatz whose comments on vulturing inspired the new section on vulturing.

References

[1] "An Iterative Procedure for Non-Discounted Discrete-Time Markov Decisions," G. J. Koehler, A. B. Whinston, and G. P. Wright, *Naval Research Logistics Quarterly*, pp. 719-723, December, 1974.

[2] Koehler, G. J., 2010. Ultimate X Poker Analysis. http://playperfectllc.com/uploads/3/4/9/0/34902374/ultimatex.pdf

[3] Odoni , A.R., "On Finding the Maximal Gain for Markov Decision Processes," *Operations Research*, 17, pp. 857-860 (1969) .

[4] Shackleford, M., 2010. http://wizardofodds.com/ultimatex