Memo

Re: Analysis of Belgian instant scratch Bingo game
Date: April 19, 2011

To: Floris Olsthoorn, Natuur & Techniek
From: R. Mohan Srivastava, FSS Canada
cc: Project files (Toronto)

These notes summarize analysis of a Belgian version of a popular instant scratch game that uses Bingo as the basis for the game's mechanics.

Data

This analysis is based on data from 30 cards provided by Floris Olsthoorn of *Natuur & Techniek* as part of the background research for an article. The data file included the numbers visible on the face of the card prior to its being played (the 96 numbers on the four Bingo boards), as well as the numbers that are hidden under a latex coating (the 24 numbers that I will refer to here as the "scratch list").

While there has been no concerted effort to verify the data entry, the software used for the analysis presented here does check that the data is consistent with the Bingo game mechanic. These checks did identify one data entry error, and this error was corrected. Having done several similar studies on other Bingo-style instant scratch games, and using the same software, my impression is that the data entry was very good in this case ... better than it usually is. And I do not feel that the main observations presented here are the result of data entry errors. Specifically, such errors tend to be random and to degrade the strength of the type of observations made here; random errors do not tend to strengthen such observations.

Purpose

The purpose of this analysis is to determine whether or not there might be some procedure for separating winning tickets from losing tickets.

In my opinion, an instant scratch game is flawed if this type of procedure (or, indeed, any other) allows one to separate winners from losers with a probability significantly higher than the "dumb luck" odds of the game. With the published odds of winning in this game being 27.7% (1 in 3.61), my opinion is that a significant departure from "dumb luck" would be to win about 55% of the time, i.e. to be able to double one's chances of winning. In this particular game, with the vast majority (roughly 90%) of the winning tickets having a prize value of €6 that is twice the cost of the ticket (€3), a success rate of 55% would allow one to play the game at a modest profit.

Procedure 1: Bayes' Theorem

In Appendix A of a 2003 report of mine entitled *PROBLEMS WITH AN INSTANT SCRATCH LOTTERY GAME: An analysis of why the OLGC's TicTacToe game was exploitable*, I outlined a procedure that uses Bayes' Theorem to rank the likelihood of each card being a winner, using the frequency of occurrence of the visible numbers to update prior probabilities.

This Bayesian updating procedure requires three pieces of information:

1. The possible X and O configurations† that might be present on each board.
2. The prior probability of each of the possible configurations.
3. The conditional probabilities of a cell containing a number that repeats n times given that it the cell is an X, and given that the cell is an O.

The possible configurations

For the analysis presented here, the set of possible configurations was limited to those that that were either losers, €3 "winners", or €6 winners. There are two reasons for limiting the possibilities to these, and not including higher-valued winners:

- the 30 cards do not include any higher-valued winners, so it is more difficult to make inferences about what the X–O configurations might look like on these cards; and,
- cards with a value of €0, €3 or €6 cover over 97% of the cards in the game.

From an analysis of the 30 cards provided, and from information published on the lottery website, several characteristics of the X–O configurations were identified:

1. A board‡ can not have more than one winning alignment. This is necessary to avoid confusion on the prize amount.

†In these notes, an "X" is a cell that has a number that appears on the scratch list; and "O" is a cell that has a number that does not appear on the scratch list.

‡A card can have more than one winning board, but any single board cannot have more than one winning pattern.

2. There are between 10 and 18 Xs on the board. The reason for the maximum is that, with more than 18 Xs, the board ends up with multiple winning alignments. The probable reason for the minimum is that, with fewer than 10 Xs, the player will feel cheated . . . that they got a "bad" card. It is normal in these instant scratch Bingo games that the game is designed to try to create a sense of almost having won, a sense that, with just one more number, you would have a winner.
3. There are no empty columns or rows. The reason for this is likely related to the previous observation, i.e. that boards with empty columns or rows look "bad".
4. There are always at least two Xs in columns 1 and 3, and in rows 3 and 5. This appears to be an attempt to provide more Xs in the columns and rows that are involved in certain winning alignments (e.g. the "L" winning pattern that uses column 1 and row 5, or the "line" winning patterns that use column or row 3).
5. Rows 1, 2, 4 and 5 and columns 1, 2, 4 and 5 cannot have five Xs unless these are critical to some other winning pattern (e.g. row 1 can be full only if the card has the "T" pattern).

With these constraints, there are 5,552,647 possible X–O configurations that have a value of either €0, €3 or €6. These were enumerated by a computer program that goes through all 2^{24} possibilities, and rejecting the ones that do not conform to the constraints given above.

The prior probabilities

The prior probabilities for each of the 5,552,647 X–O configurations can be calculated from the following table, which comes from the Belgian lottery website, and gives the exact number of tickets in each prize-winning category in each batch of 1,500,000.

PLAN DES LOTS DU BINGO (pour 1.500.000 de billets) :

Lots	Nbr de gains	Montant (Eur)	Total (Eur)	Probabilité (1 chance sur)
Bingo !	2	75.000,00	150.000,00	750.000,00
Carré	20	1.000,00	20.000,00	75.000,00
Croix	400	100,00	40.000,00	3.750,00
Plus	1.000	15,00	15.000,00	1.500,00
4 coins + 1 lettre	3.500	15,00	52.500,00	428,57
2 x lettres + 1 ligne	3.500	15,00	52.500,00	428,57
4 coins	5000	9,00	45.000,00	300,00
1 lettre + 1 ligne	12500	9,00	112.500,00	120,00
3 lignes	12500	9,00	112.500,00	120,00
Lettre	114000	6,00	684.000,00	13,16
2 lignes	258000	6,00	1.548.000,00	5,81
Ligne	5000	3,00	15.000,00	300,00
Total	**415.422**		**2.847.000,00**	**3,61**

This table enumerates the details of 6,000,000 boards (four on each of 1,500,000 cards). From this, we can calculate the number of boards that have each type of winning pattern.

The €3 winning "line" pattern appears three times on 12,500 cards ("3 lines" on the table above); it appears twice on 258,000 cards ("2 lines"); and it appears once on 3,500+12,500+5,000 cards ("2 letters + 1 line", "1 letter + 1 line" and "line"). This adds up to 574,500 boards with the €3 winning "line" pattern. So the probability of encountering this pattern is

$$P_3 = \frac{574{,}500}{6{,}000{,}000} = 0.0958$$

The €6 winning "letter" pattern appears twice on 3,500 cards ("2 letters + 1 line"); and it appears once on 3,500 + 12,500 + 114,000 cards ("4 coins + 1 letter", "1 letter + 1 line" and "letter"). This adds up to 137,000 boards with the €6 winning "letter" pattern. So the probability of encountering this pattern is

$$P_6 = \frac{137{,}000}{6{,}000{,}000} = 0.0228$$

Of the 6,000,000 – 574,500 – 137,000 = 5,288,500 boards that remain, a tiny percentage† of these have higher-valued winning patterns. But, for the purposes of this exercise, I have decided to call all of these €0 boards. So the probability of encountering no winning pattern is:

$$P_0 = \frac{5{,}288{,}500}{6{,}000{,}000} = 0.8814$$

Of the 5,552,647 X–O configurations, 1,323,069 have a value of €3. The sum of the prior probabilities for these 1,323,069 configurations has to be the 0.0958 value calculated above, so the prior probability for each one is approximately‡

$$\frac{0.0958}{1{,}323{,}069} = 7.2407 \times 10^{-8}$$

Of the 5,552,647 X–O configurations, 73,449 have a value of €6. The sum of the prior probabilities for these 73,449 configurations has to be the 0.0228 value calculated above, so the prior probability for each one is approximately

$$\frac{0.0228}{73{,}449} = 3.1042 \times 10^{-7}$$

† Less than 0.17%

‡ This is an approximation because we are assuming that the probability of a single board having a certain value is the same as the entire card having that value. This is a good approximation for the €6 "letter" winning patterns, but is not as good for the €3 "line" patterns because the game rarely has a single €3 winner. As the table shows, the €3 "line" patterns are almost always paired on a card, producing a card that has a value of €6. It is also an approximation because we don't know that our inventory of possible winning configurations exactly matches those used by the printer who printed the game ... we may be considering configurations that are, in fact, never used by the printer.

Of the 5,552,647 X–O configurations, 4,156,129 have a value of €0. The sum of the prior probabilities for these 4,156,129 configurations has to be the 0.8814 value calculated above, so the prior probability for each one is

$$\frac{0.8814}{4{,}156{,}129} = 2.1207 \times 10^{-7}$$

The conditional probabilities

The conditional probabilities $P(F = n|X)$ are calculated simply by counting the number of times on the available cards that an X cell has a number that occurs once, twice, three times or four times and dividing by the total number of X cells:

$$\begin{aligned} P(F = 1|X) &= 192 \div 1{,}595 = 0.1204 \\ P(F = 2|X) &= 391 \div 1{,}595 = 0.2451 \\ P(F = 3|X) &= 558 \div 1{,}595 = 0.3498 \\ P(F = 4|X) &= 454 \div 1{,}595 = 0.2846 \end{aligned}$$

The same approach is used to calculate the conditional probabilities $P(F = n|O)$:

$$\begin{aligned} P(F = 1|O) &= 493 \div 1{,}285 = 0.3837 \\ P(F = 2|O) &= 395 \div 1{,}285 = 0.3074 \\ P(F = 3|O) &= 255 \div 1{,}285 = 0.1984 \\ P(F = 4|O) &= 142 \div 1{,}285 = 0.1105 \end{aligned}$$

What these two tables show is that, in the X cells, the numbers that appear most frequently are those that appear three and four times on the card, while in the O cells, the most frequently-appearing numbers are those that occur once or twice.

Results

I started with the same program that I wrote in 2003 for the Ontario TicTacToe game, and modified it to the specifics of this Belgian Bingo game. The output of this program is a number that can be used to rank the cards in order of their likelihood of winning[†]

[†]I do not think of this as a "true" probability since there are other aspects of the game's creation that are not taken into account in the procedure. For example, in this Belgian Bingo game, the €3 patterns almost always occur in pairs, which is not something that would happen randomly; it is designed into the software. Even though the output is not a true probability, it can still can be used to rank the cards from worst to best (as long as there is some correlation between this ranking index and the true probability).

With winning odds of 1 in 3.61, a group of 30 cards should have about eight winners.† So I ranked the cards from top to bottom using the results of the Bayesian updating procedure, and took the top eight as my predicted winners. Sadly, in this group there were only three actual winners . . .

. . . so what went wrong?

Procedure 1A: Bayes' Theorem, take two

I was surprised that the procedure didn't work, especially because the conditional probabilities clearly point to the X cells containing more of the higher frequency numbers, and the O cells containing more of the lower frequency numbers.

My first guess was that my set of possible X–O configurations was not good . . . that my five million possibilities was very different from whatever set was actually being used. But I couldn't figure out any way to further constrain the set, so this possible explanation remains open to further investigation.

My second guess (I was plodding through the three pieces of input information) was that it might be a mistake to assign the same prior probability to each possible X–O configuration that had the same winning value. Perhaps some of my €3 winning configurations were more likely, a priori, than others? This is quite possible, given the tendency to have more Xs than Os on each board . . . this helps sustain the illusion of having almost won. I did try a couple of variations on the prior probabilities, but with little success. In one of these, I managed to get four winners in the cards that were ranked in the top eight; but four out of eight is not a remarkable success.‡

So I then turned to the conditional probabilities, and wondered what might be wrong with those. It occurred to me to check if $P(F = n|X)$ is the same if the X cell happens to be part of a winning pattern.

I divided the X cells into two groups:

- the X0 cells that are X cells that do not form part of a winning alignment; and,
- the X1 cells that are X cells that form part of a winning alignment

†As it turns out, this is exactly the number that appear in the set of 30 cards whose data was provided. Although this may appear like a stroke of good luck . . . and it is a little bit lucky . . . this may not actually as lucky as it might at first seem. In all of the North American instant scratch games, the game is created in such a way that the number of winners in each packet of tickets (100 to 200 tickets in a packet) is made much more uniform than pure randomness allows. This is done so that retailers don't get the feeling that they've been sold "bad" packets.

‡The chances of getting four winners in a randomly chosen set of eight is around 20% . . . uncommon, but not hugely surprising.

Here is the table of conditional probabilities for the X0 cells:

$$
\begin{aligned}
P(F = 1|X0) &= 15 \div 63 = 0.2381 \\
P(F = 2|X0) &= 17 \div 63 = 0.2698 \\
P(F = 3|X0) &= 21 \div 63 = 0.3333 \\
P(F = 4|X0) &= 10 \div 63 = 0.1587
\end{aligned}
$$

Compared to the earlier calculation of the $P(F = n|X)$ probabilities, the notable change is that, on X0 cells, the least likely frequency of occurrence is four. Previously, when I looked at all X cells, the least like frequency of occurrence was one. What this suggests is that the game has been tweaked a bit to make the frequency of occurrence misleading. This *might* be intentional (i.e. a response to knowledge that the frequencies of occurrence can betray useful information). Or it might be an accidental consequence of some other programming constraint.

Whatever the reason for the difference, it was clear to me why the first attempt at the Bayesian procedure failed: for boards with winning patterns, these were being declared unlikely as winners because the winning patterns often contain more 1-frequency numbers than 4-frequency numbers.

When I replace the old set of conditional probabilities with a new set that has values for $P(F = n|X0)$, $P(F = n|X1)$ and $P(F = n|O)$, the Bayesian procedure now produces ranks that get five winners in the top-ranked eight. The chance of this happening by pure dumb luck is about 3%, which is low enough that the norms of statistical analysis would deem this to be “statistically significant”.

Procedure 2: Exploiting the anomaly in $\sum_x F$ and $\sum_o F$ on the winning boards

As noted above, the initial attempt to use the Bayes' Theorem approach was thwarted by the fact that the frequencies of occurrence change on winning boards. This suggests that the algorithm that creates the cards does not use the same procedure for creating winning boards and losing boards. So I was curious about what other traces this might leave on the cards. Table 1 on the following page shows, board-by-board, the sum of the frequencies on all the X cells and the sum of the frequencies on all of the O cells. For each board, I simply add up the frequency of occurrence for the cells that are on the scratch list; that gives me $\sum_x F$. And I do the same for all the cells that are not on the scratch list; that gives me $\sum_o F$.

In Table 1, I have also included the ratio of $\sum_x F$ to $\sum_o F$. A quick skim up and down the Ratio column will confirm that this ratio is usually around 2: the sum of the frequencies on the X cells is usually about twice that on the O cells. This makes sense because: a) there are usually a couple more X cells than O cells, and b) the X cells usually carry the higher frequency numbers. But there are a few notable exceptions to this general observation; in a handful of cases, the ratio is much lower, down around 1 or less. The surprising thing (but entirely consistent with the

Table 1. The sum of frequencies, board-by-board, for all the X cells and all the O cells, with the ratio shown in red for the boards with winning patterns.

Card	Board	$\sum_x F$	$\sum_o F$	Ratio
1	1	35	23	1.52
1	2	35	22	1.59
1	3	38	22	1.73
1	4	43	16	2.69
2	1	33	23	1.44
2	2	34	24	1.42
2	3	39	20	1.95
2	4	34	25	1.36
3	1	41	15	2.73
3	2	47	13	3.62
3	3	43	17	2.53
3	4	39	21	1.86
4	1	39	16	2.44
4	2	42	19	2.21
4	3	39	16	2.44
4	4	35	22	1.59
5	1	27	35	**0.77**
5	2	32	32	**1.00**
5	3	38	19	2.00
5	4	35	28	1.25
6	1	43	22	1.96
6	2	38	22	1.73
6	3	29	33	**0.88**
6	4	24	31	**0.77**
7	1	39	20	1.95
7	2	35	19	1.84
7	3	34	22	1.55
7	4	29	28	**1.04**
8	1	39	21	1.86
8	2	41	20	2.05
8	3	41	17	2.41
8	4	38	21	1.81
9	1	40	22	1.82
9	2	40	17	2.35
9	3	41	21	1.95
9	4	38	19	2.00
10	1	32	28	1.14
10	2	33	23	1.44
10	3	31	25	1.24
10	4	37	21	1.76
11	1	46	14	3.29
11	2	39	18	2.17
11	3	44	15	2.93
11	4	39	15	2.60
12	1	40	22	1.82
12	2	39	20	1.95
12	3	39	16	2.44
12	4	45	17	2.65
13	1	34	21	1.62
13	2	43	18	2.39
13	3	50	11	4.55
13	4	39	18	2.17
14	1	26	31	**0.84**
14	2	22	38	**0.58**
14	3	35	27	1.30
14	4	31	30	1.03
15	1	22	36	**0.61**
15	2	38	22	1.73
15	3	36	26	1.39
15	4	42	20	2.10
16	1	41	16	2.56
16	2	41	17	2.41
16	3	41	20	2.05
16	4	48	14	3.43
17	1	28	26	1.08
17	2	24	38	**0.63**
17	3	33	21	1.57
17	4	36	24	1.50
18	1	43	17	2.53
18	2	39	21	1.86
18	3	38	22	1.73
18	4	38	14	2.71
19	1	27	29	0.93
19	2	39	23	1.70
19	3	39	22	1.77
19	4	44	17	2.59
20	1	31	22	1.41
20	2	42	20	2.10
20	3	42	19	2.21
20	4	42	16	2.63
21	1	41	18	2.28
21	2	40	18	2.22
21	3	38	21	1.81
21	4	44	14	3.14
22	1	36	23	1.57
22	2	34	23	1.48
22	3	43	22	1.96
22	4	36	25	1.44
23	1	40	20	2.00
23	2	36	23	1.57
23	3	35	24	1.46
23	4	40	16	2.50
24	1	25	33	**0.76**
24	2	34	23	1.48
24	3	24	36	**0.67**
24	4	32	23	1.39
25	1	35	25	1.40
25	2	31	27	1.15
25	3	32	31	1.03
25	4	42	17	2.47
26	1	45	12	3.75
26	2	34	25	1.36
26	3	44	20	2.20
26	4	38	20	1.90
27	1	31	30	1.03
27	2	35	22	1.59
27	3	31	30	1.03
27	4	39	28	1.39
28	1	41	20	2.05
28	2	45	22	2.05
28	3	41	21	1.95
28	4	36	24	1.50
29	1	41	17	2.41
29	2	34	21	1.62
29	3	42	15	2.80
29	4	41	19	2.16
30	1	45	16	2.81
30	2	23	32	**0.72**
30	3	46	11	4.18
30	4	39	16	2.44

earlier observation about the frequencies on the X1 cells) is that the lowest values are all on the winning alignments. When the ratio of $\sum_x F$ to $\sum_o F$ is below 1.10, the board is almost certainly a winner.[†]

While this is an interesting result, and confirms that the winning boards are generated using a different procedure than the losing boards, it cannot serve as a winner/loser separation procedure because calculating $\sum_x F$ and $\sum_o F$ requires first knowing where the X and O cells are, and we don't know this until we scratch off the card. But it does give us an anomaly that we can exploit.

I have looked into two possibilities here. The first is trying to predict the $\sum_x F$:$\sum_o F$ ratio using the information that we can see: the frequencies of occurrence. The second is checking for anomalies in $\sum_x F$ and $\sum_o F$ along the columns, rows and diagonals of the board.

Trying to predict the ratio

The best predictor I could find, using regression, has a correlation of about 0.5 with the true ratio. While this is encouraging, it is not nearly strong enough to serve as a good basis for separating winners from losers.

But what does work well, is trying to predict the spread in the ratios. On winning cards, the winning boards have very low ratios and the losing boards have much higher ratios. Even though my attempt to predict the ratio was not great, it does give me a good ability to rank the spread of the ratios on any card. Using this to rank the cards from best (biggest spread in the predicted ratios) to worst (lowest spread in the predicted ratios), I was able to get five winners in the eight cards that ranked at the top of the list.

Anomalies in $\sum_x F$ and $\sum_o F$ along columns, rows and diagonals

Even though we do not know exactly where the X cells are, we can still test each of the columns, rows and diagonals involved in the winning patterns to see if they have anomalous values for $\sum_x F$ or $\sum_o F$. There are three columns, three rows and two diagonals that figure prominently in the winning alignments. The sum of the frequencies can be calculated along each of these, and on the remainder of the board, providing a kind of pseudo-value for $\sum_x F$ or $\sum_o F$ ... I call these "pseudo" because they are calculations that assume that all the Xs are along a particular column, row or diagonal.

The reason that it occurred to me that this might work is that it's clear that the algorithm being used is doing something different for the winning boards, and is trying to deplete (or accidentally depletes) the frequencies along the winning alignments. So by check all of the possible winning lines, I might be able to pick up a slight weirdness in the sum of the frequencies.

I have 120 boards on 30 cards, I calculate the pseudo-values for $\sum_x F$ and $\sum_o F$ for each of the eight lines. This gives me 16 values for each board: a pseudo-value for $\sum_x F$ and for $\sum_o F$ for

[†]The only board that makes this rule less than a complete success is the fourth board on card #14, which has a ratio of 1.03 even though it does not have a winning pattern. But this board is on a card that has two other winning boards, and so the rule would still work, but not because that particular board has a low ratio but because some of its companion boards have even lower ratios.

each of eight lines. I then compared the pseudo-values for each line, giving them a "weirdness index" that is simply the absolute value of their rank minus 60. What this does is it gives numbers in the middle of the list a zero, and numbers at the top or bottom of the list a value of 60. So the more anomalous the value is, the more it departs from the average, the higher its "weirdness index" will be. I have 16 weirdness indexes for each board. I add these up for each card to create a total weirdness index, and then I use this to rank the cards.

Using this procedure, I again get five winners in the top eight spots.†

Procedure 3: Monte Carlo simulation

As one final kick at the can, I checked to see what would happen if used a Monte Carlo simulation procedure to randomly generate possible versions of the scratch list. I know that the 24 numbers on the scratch list all appear somewhere on the board.‡ So I can take all of the numbers that show up on the board and see if I can guess at a set of 24 that might form the scratch list. I can have my computer go through several thousand such guesses and count up the number of times that the card is a winner.

The reason that this occurred to me as a procedure worth exploring is that I noticed that the top prize in this game comes from a very peculiar pattern: the entire board is covered with Xs. So it occurred to me to check, on each of the 30 cards I was provided, what would happen if I covered each of the four boards with Xs and used this to define the scratch list ... what would happen on the *other* boards.

Figure 1 on the following page shows an example of what I was checking. This figure shows how the four boards on one particular card (this happens to be Card #12) if I define the scratch list as being the 24 numbers on each board. I noticed that it is actually quite hard to keep the entire card legitimate. For example, on the first row of Figure 1, the second and third boards have a row of Xs across the top row, and this can't happen unless the third column is also full (creating the "T" pattern). On the second row of Figure 1, we have the "Y" pattern on the first board, and this can't happen if the second board is a €75,000 winner (i.e. there are no €75,006 winners; if you win €75,000 on one board, you can't win anywhere else). On the third row, the first and second boards have empty rows, something that never happens (at least in the 120 boards that I've observed). And on the fourth row, the first board is a €3 winner, which can't happen if the fourth board is a €75,000 winner.

†Even though my top-eight procedure would not have looked any further, it turns out that the ninth spot was also a winner, which provides some additional support for the view that this "weirdness" ranking does accomplish something useful.

‡This is not a requirement of a Bingo game, but it is a common design constraint for instant scratch Bingo games. Players usually feel they got cheated if they scratch off a number and it doesn't appear anywhere.

Figure 1. The X–O patterns on all four boards of Card #12 when each of the boards is assumed to be a €75,000 winner.

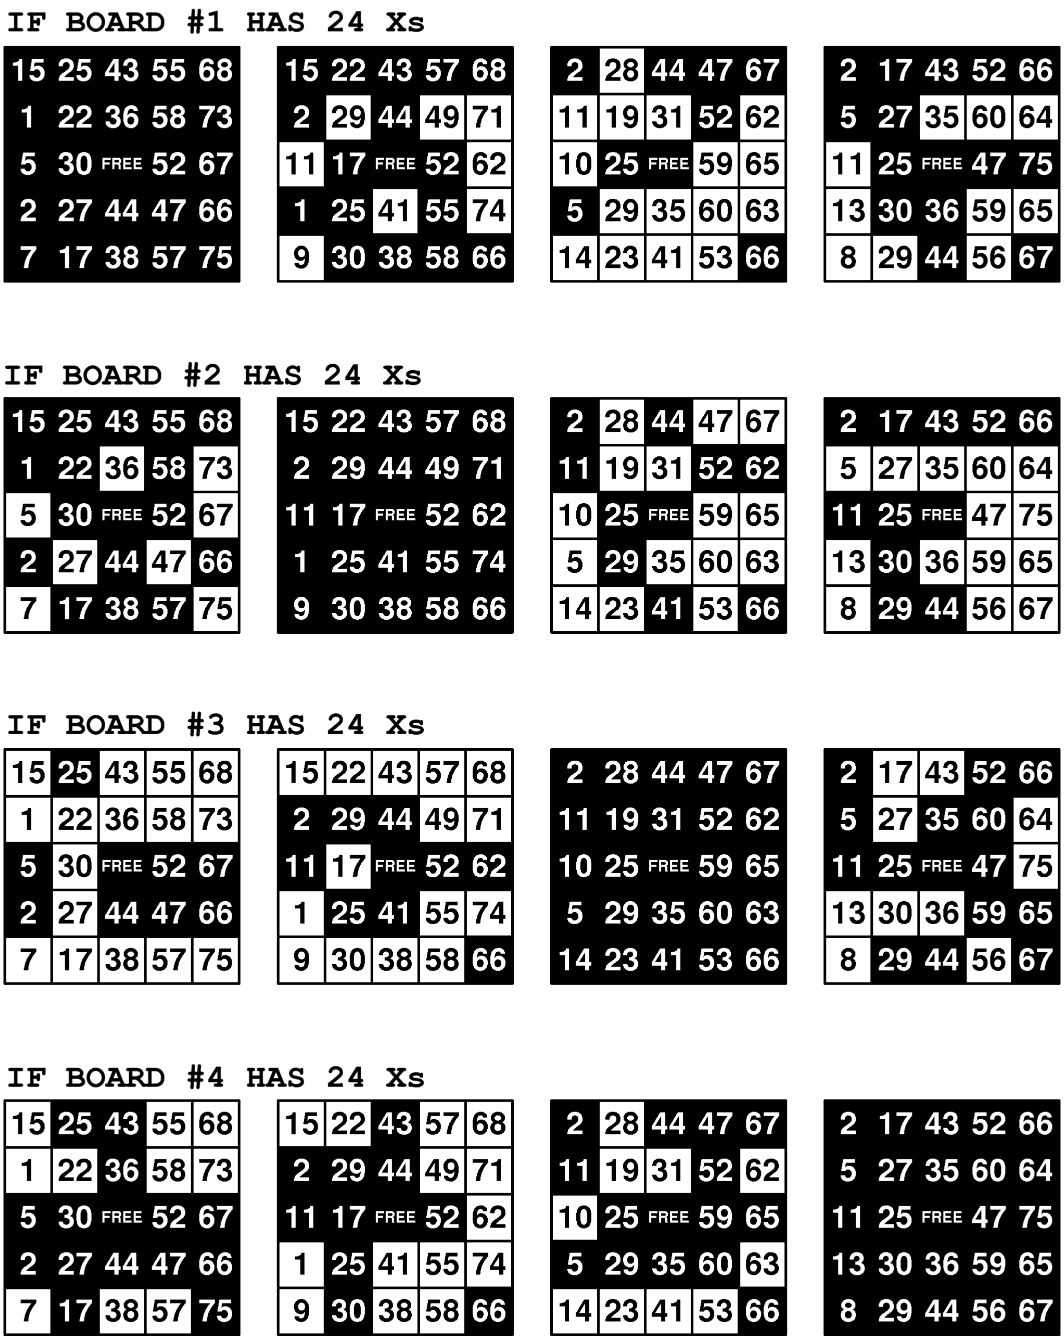

So the example in Figure 1 shows a situation where, just by studying the numbers on the face of the card, you can confirm that this cannot possibly be a €75,000 winner.† On all of the other cards in the set I was provided, you can do the same thing: you can confirm that none of these can be a €75,000 winner simply by checking to see if the card becomes illegitimate when you assume that all 24 numbers on one board form the scratch list.

What this observation led me to conclude is that it's actually very hard to create sets of 24 numbers for the scratch list and have the entire card remain legitimate. There are so many constraints on the X–O configurations on a board that most of the randomly-drawn sets of 24 will not work for the scratch list.

So I wrote a small program that creates random sets of 24 and then checks to see if the entire card remains legitimate. In addition to the constraints described at the beginning of these notes, I included one additional constraint: if a board was a winner, then the $\sum_x F$:$\sum_o F$ ratio had to be less than 1.10; this additional constraint comes from the observation shown in Table 1.

If the randomly drawn scratch list creates a card that is nor entirely legitimate (i.e. I violate some design constraint), then I reject that possibility. If the randomly drawn scratch list creates a card that is entirely legitimate, I keep that possibility, and check to see if the card is a winner or a loser. I keep drawing random scratch lists until I've managed to create a set of 100,000 legitimate cards (which entails millions of attempts, because most of them fail). From this, the number of winners I was able to create serves as my ranking criterion.

With this procedure, I was able to get six winners in the top-ranked eight, a success rate of 75%, and something that has a less than a 1% chance of occurring by pure luck.

Conclusions

With three entirely different procedures — Bayesian updating, exploiting anomalies in $\sum_x F$ and $\sum_o F$, and Monte Carlo simulation — all achieving a success rate well above 50%, it is clear to me that there is something worth taking a closer look at here.

For those ostriches who wish to stick their heads firmly in the sand and avoid further investigation, it is easy to dismiss all of this as a fluke, a bunch of lucky guesses, an accident that affects only the 30 cards that I was provided. But for those who have some ability to think statistically, then any statistically significant deviation from the "dumb luck" hypothesis deserves closer investigation.

To me, knowing how little time I put into this analysis (a total of about four hours on two train rides between Toronto and Kingston), I am quite sure that I could make improvements in any of the three successful procedures described in these notes. In my mind, these have easily met the minimum threshold of having an experimental success rate that is at least 55%, and that has a

† I think it's important to note the difference between making a good guess that the card is not a €75,000 winner (because almost none of them are) versus *knowing* that it can't be, because there is no board for which all 24 numbers could simulataneously be on the scratch list. Even if there's no generally successful winner/loser separation trick, I think the game is flawed if, for certain cards, a person is able to confirm that they cannot possibly be the top prize winner.

statistically significant chance of deviating from the fair success rate of about 27%.†

With the three procedures all using different approaches, and relying on different statistics (all of which can be gleaned without scratching off anything), the combination of all three procedures could be used to high-grade a set of cards, i.e. retain as likely winners only those cards that score well on all three procedures.

A closing cautionary note

It is worth stating the obvious here. None of the procedures described in these notes is the kind of thing that a human being could do in their heads; I have not shown a card trick that can be done in real time without the aid of a computer. In this sense, the Belgian Bingo game is different from the Ontario TicTacToe game from 2003. In that Ontario game, it was possible to leverage the Bayesian analysis to develop a simple rule (the 1s rule described in the Wired magazine article) that many people could do in their heads.

While it is true that a computer is a great help here, I am not sure that it is necessary. In the set of cards that I had, the only two cards that had frequency sums of five or less across the middle row of one of the boards were both winners. This observation is based on only two cards that had that characteristic; and two cards does not make this a statistically significant result. But it might turn out that, with more cards, this actually does turn out to be a useful rapid diagnostic procedure for winners.

Even if a computer is necessary, I do not think that this should be used as an excuse for ignoring the issue. With modern hand-held devices like smart-phones that are programmable and that have cameras, it would be relatively simple for someone to write an app that allows you to take a picture of the card and then do whatever computations are required to determine if it is a likely winner. As long as buyers have a choice in which cards they are allowed to buy‡, then the game can be exploited by anyone with the patience to write a smart-phone app.

The other group of people who *always* have the luxury of sorting through the available cards, and taking whatever computational time is needed to make their choices, are the retailers. Any retailer that chose to try to separate winners from losers has the opportunity to do so, an opportunity that is curtailed in the general public.

†I.e. the chance of the actual success rate being 27% is less than 5%.

‡Some jurisdictions limit this choice by using dispensing machines. Others, like Ontario, allow users to choose from a large group that are on display at the cash register.