THE COMPATRIOT-WIN EFFECT ON NATIONAL SALES OF A MULTI-COUNTRY LOTTERY

by

Rose Baker[a]

David Forrest[b]

and

Levi Pérez[c]

[a] Salford Business School, University of Salford, Salford, M5 4WT, UK, e-mail- rose.baker@cantab.net, tel. +44 161 295 3861

[b] University of Liverpool Management School, University of Liverpool, Chatham Street, Liverpool, L69 7ZH, UK, e-mail- david.forrest@liv.ac.uk, tel. +44 151 795 0679

[c] *(corresponding author)* Departamento de Economía, Universidad de Oviedo, Edificio Departamental de Ciencias Jurídico-Sociales, Campus del Cristo, 33006 Oviedo, Spain, e-mail- lperez@uniovi.es, tel. +34 985105063

THE COMPATRIOT-WIN EFFECT ON NATIONAL SALES OF A MULTI-COUNTRY LOTTERY

Abstract

Euromillions is a lotto game played across nine countries. We examine more than eight years of sales data for individual states to assess whether a jackpot win in a country increases subsequent sales in that country. We propose a novel test for the presence of such a 'compatriot win' that has as its only assumption that the lottery draw is random. Results suggest elevated sales over twelve draws following a national win. When we model the size of the impact, it proves to be modest in size for average jackpot wins but much larger and longer-lasting for the highest pay-outs.

Keywords: lotto; Euromillions; compatriot-win effect; time series; statistical test.

1. Introduction

Euromillions is an example of a multi-state lotto game and is currently played twice-weekly across nine countries. The idea of a multi-state lottery is to allow larger prize pots to be gathered than would be possible in any one jurisdiction. As a result, playing is made more attractive because individuals are able to win bigger sums (dream bigger dreams) and per capita sales across the lottery bloc should thereby be enhanced. Essentially, multi-state games seek to exploit "the peculiar scale economies of lotto" (Cook and Clotfelter, 1993).

Multi-state games in America (Powerball and Mega Millions) and elsewhere have successfully combined the markets of the separate lottery jurisdictions which make up the country. However, Euromillions is a more ambitious project in the sense that it involves the lottery agencies of separate nation states offering a common game across national boundaries where countries have their own languages and own media. In this context, the idea suggests itself that, when jackpot wins in a particular country occur (and these will be rare since there are nine countries and no jackpot winner at all in most draws), this may change the demand for tickets in that country. We call any effect on subsequent sales ‘the compatriot win effect’. Such an effect, were it to exist, would be of interest to national lottery agencies since it could be exploited in marketing strategies. Further, it should be taken into account in future statistical modelling of sales.

It might also be of wider interest. Our exploration of the compatriot win effect is somewhat analogous to the identification of a ‘lucky store’ effect in the American lottery market. Guryan and Kearney (2008) found that sales of Texas lotto tickets at a given outlet were elevated for up to nearly a year following a jackpot win by a customer of that outlet, with the effect more marked in disadvantaged areas. Part of the effect was due to substitution of purchases previously made at other stores. This may reflect a cognitive misperception among some potential buyers that the probability of winning varies with where one buys a ticket. However, part of the effect reflected increased sales in the district rather than cannibalisation and this is easier to interpret without recourse to assuming irrationality. Other behavioural biases may also be in play. Chen and Chie (2008) note that aversion to regret may be a powerful psychological driver of lottery sales. A local win may induce regret among local residents because they think “it could have been me”. They then play subsequently to avoid the regret they had experienced this time. Alternatively, one interpretation of the motives for buying lottery tickets is that players ‘buy a dream’ (Forrest, Simmons and Chesters, 2002): they enjoy daydreaming about winning a fortune. This dream may be more vivid if they observe people like themselves achieving it.

Similarly, one interpretation of a compatriot win on Euromillions sales would be that a win changes the perception in a country about how probable it is to win the game even though the true odds against a ticket winning the jackpot are unaltered (at, currently, about 116m to 1). However, a compatriot win may also stimulate national sales through the emotional traits and biases noted above with reference to the increase in neighbourhood sales when there is a local winner.

Alternatively, it could be that sales simply respond to national media coverage of a national winner, for example because some consumers had previously forgotten about the product. A national win might simply increase awareness of the game.

The traditional approach to any issue in econometrics is simultaneously to test for the existence of a proposed effect of variable *x* on variable *y* and to estimate the magnitude of that effect. In the present context, researchers could follow this approach by applying the standard lotto demand model in each country, with dummy variables to capture the effect on national sales of a recent Euromillions jackpot win. The standard lotto demand model exploits the substantial draw-to-draw variation, in either the expected value of a ticket (Gulley and Scott, 1993 and many subsequent papers) or the size of jackpot (Forrest *et al*., 2002), which arises because jackpots not won in one draw are 'rolled over' to the jackpot pool for the following draw. Typically sales at draw *t* are modelled as dependent on expected value/ jackpot in draw *t* (and control variables). However, since the level of sales itself influences expected value/ jackpot, these are not known ex ante and there is a problem of endogeneity. This is addressed by using two-stage least squares estimation, with the size of rollover serving as an instrument for expected value/ jackpot.

Such a standard approach could in principle be employed in searching for a possible compatriot effect in sales in each country in the Euromillions bloc. This would not be without practical difficulties, e.g. the control variables would need to be different in each country because there are different competing (domestic) lotteries and these too will offer varying returns from week to week. More fundamentally, any results from hypothesis tests for the existence of a compatriot win effect and coefficient estimates to illustrate the size of the effect would be subject to the qualification that the findings would be only as sound as the assumptions made in the model. One example of an assumption implicit in the instrumental variables procedure is that buyers exhibit rational expectations (Forrest and Gulley, 2014).

In this paper, we separate out testing for the existence of a compatriot win effect from estimation of its magnitude. We devised a novel statistical test for the existence of an effect which had minimal and uncontroversial assumptions. Our evidence will show that there is indeed a statistically significant effect and the principal contribution of the paper is to put this on the agenda as something to be included in future modelling. However, we do go on to offer a simple model to help gauge the economic significance of the effect albeit this is dependent on more assumptions than the test for existence. The two steps, of hypothesis testing and estimation, are separated out because the former can be executed with fewer assumptions than the latter.

2. Data

The Spanish lottery agency, Loterías y Apuestas del Estado, provided us with time series of the number of tickets sold in each country where Euromillions was available. The time series extend from the inception of the game in February, 2004 to the 465th draw in March, 2012. From draw 379, it was played twice- rather than once-weekly. Initially, the game was offered just in France, Spain and the United Kingdom (UK) but six other countries (Austria, Belgium, Ireland, Luxemburg, Portugal and

Switzerland) joined from draw 35. In the case of Switzerland, the data provided sales for each of two lottery agencies, Swisslos Landeslotterie and Loterie Suisse Romande; however, for our analysis, we combined these into one national figure.

Table 1 provides information on variation in the size of jackpot. The maximum is at the €185m cap introduced at draw 265 and other very high values occurred at the end of each of two separate sequences when there was no jackpot winner for eleven draws and, consequently, the accumulation of a dramatically high first prize for the twelfth. The minimum value is from the organisers' guarantee that the jackpot will never fall below €10m (€15m from April, 2005). Table 1 also provides summary statistics of the share of sales accounted for by individual countries, a key variable in our analysis below.

Table 1. Summary statistics for ticket sales per draw by country

	mean	standard deviation	minimum	maximum	coefficient of variation
jackpot size (€m.)	39.43	34.19	10.00	185.00	
Austria	0.039	0.007	0.025	0.106	0.18
Belgium	0.065	0.010	0.040	0.115	0.61
France	0.231	0.028	0.164	0.347	0.27
Ireland	0.021	0.005	0.008	0.035	0.24
Luxemburg	0.006	0.001	0.003	0.015	0.50
Portugal	0.204	0.041	0.108	0.300	0.30
Spain	0.212	0.028	0.112	0.310	0.30
Switzerland	0.057	0.007	0.033	0.084	0.12
United Kingdom	0.166	0.061	0.048	0.407	0.38

data for individual countries describe the share of aggregate sales in a draw accounted for by the named country; statistics relate to the period from draw 35 when all nine countries were participating in the game

In preliminary analysis, we performed a Pearson chi-squared test to check whether there was any reason to doubt that jackpot wins were equally likely to occur in any country conditional on level of country sales. With N countries, let the total number of winning tickets in the ith country be M_i, the total proportion of sales accounted for by the country be q_i and the total number of winning tickets be n. Then $\mathrm{X}^2 = \sum_{i=1}^{N} \frac{(M_i - nq_i)^2}{nq_i}$.

Countries for which nq_i<5 were lumped into one, to give $\mathrm{X}^2 = 3.713$, on 7 degrees of freedom, which is not statistically significant and shows that there is no evidence that the lottery is unfair. More importantly, it shows that the sales figures in the database are believable.

When we looked at the time-series of sales in any one country and compared it with the incidence of national jackpot wins, it seemed unlikely that a compatriot win effect could be detected, certainly without a regression model. Figure 1 shows draw-by-draw sales in the UK with crosses to indicate draws immediately following a national win. As always in lotto games, the time-series shows wild fluctuations in sales because the attractiveness of the jackpot prize varies so much. The variation in pattern near the end of the plot coincides with the introduction of bi-weekly draws and reflects that sales are typically lower in Tuesday than in Friday draws. As might be expected, national winners are observed more often in the draws with highest sales; but there is no immediately obvious relationship between the incidence of winners and subsequent sales.

Figure 1. UK sales (number of tickets in millions)

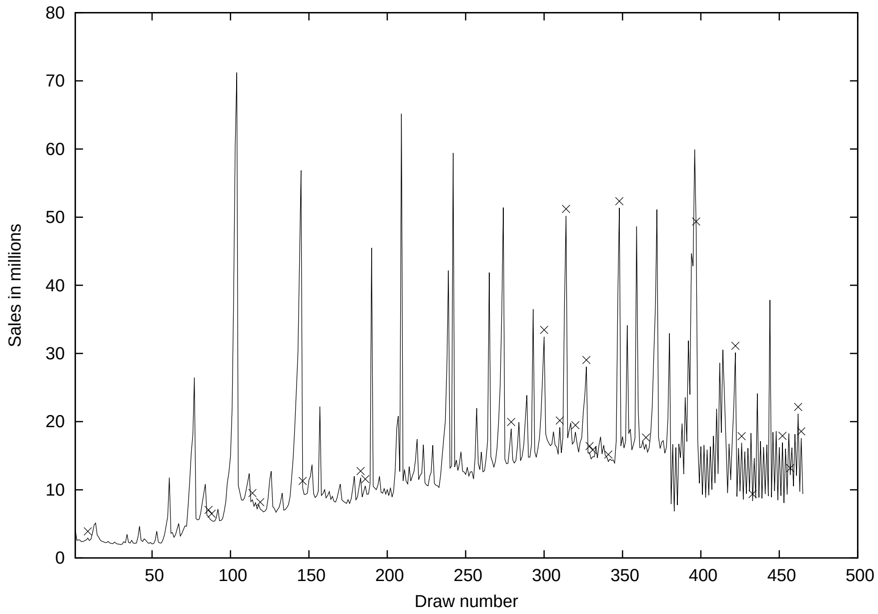

crosses indicate draws for which there was a UK jackpot winner

Figure 2. UK sales at each draw as a percentage of aggregate sales at that draw

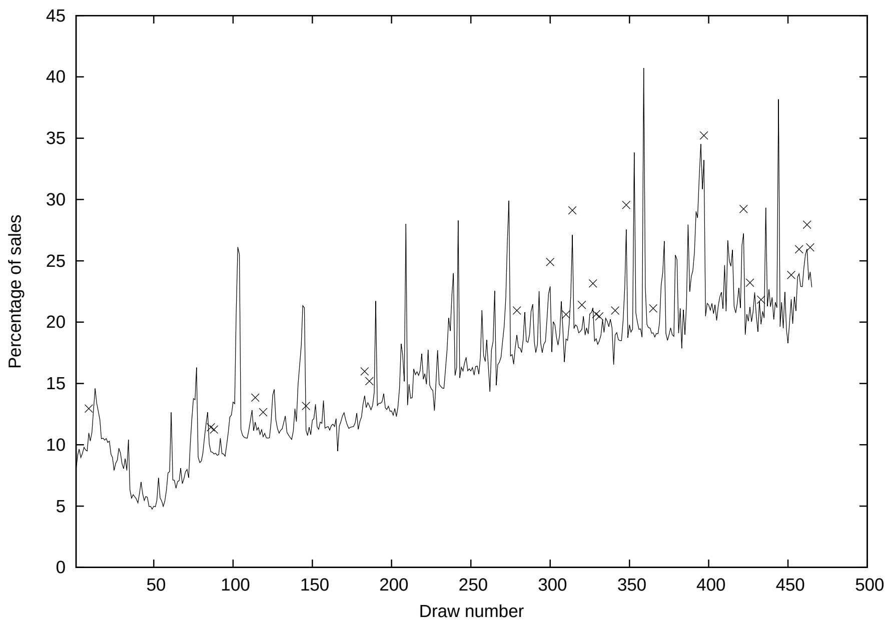

crosses indicate draws for which there was a UK jackpot winner

The proportion of total sales from any particular country (relative sales) does not fluctuate so much from draw to draw because large jackpots and any special events tend to increase sales everywhere. Nevertheless, figure 2, displaying relative sales for the UK, still shows some variability (such that this country purchases a greater share of tickets sold when the jackpot is high). Against this variability, there is still no obvious sign of a compatriot win effect.

Figures 3 and 4 show corresponding data for Spain; here relative sales does not show peaks at large jackpots, but rather dips. This is because UK sales have a large elasticity with respect to jackpot size (0.6) whereas Spanish sales have smaller elasticity (0.2). Hence the proportion of UK sales increases with a large jackpot, while the proportion of Spanish sales decreases (elasticities were calculated by a simple regression of log sales on log jackpot). This is consistent with the observation by Roger (2011), based on regression of sales on jackpot and other variables, that the UK had the highest and Spain the lowest sensitivity of sales to jackpot size among the nine member countries. He speculated that explanations might include different income levels, different degrees of competition in the gambling market and a strong tendency of Spanish players to participate through syndicates. Whatever the explanation, Spain is like the UK in that the charts point to no readily observable link between either absolute or relative sales and preceding compatriot wins.

Figure 3. Spanish sales (number of tickets in millions)

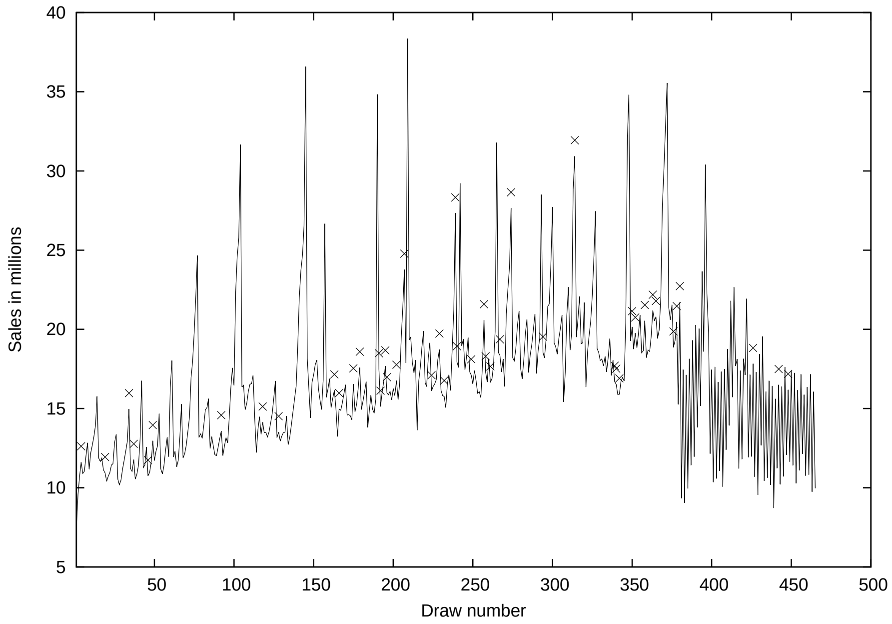

crosses indicate draws for which there was a jackpot winner in Spain

Figure 4. Spanish sales at each draw as a percentage of aggregate sales at that draw

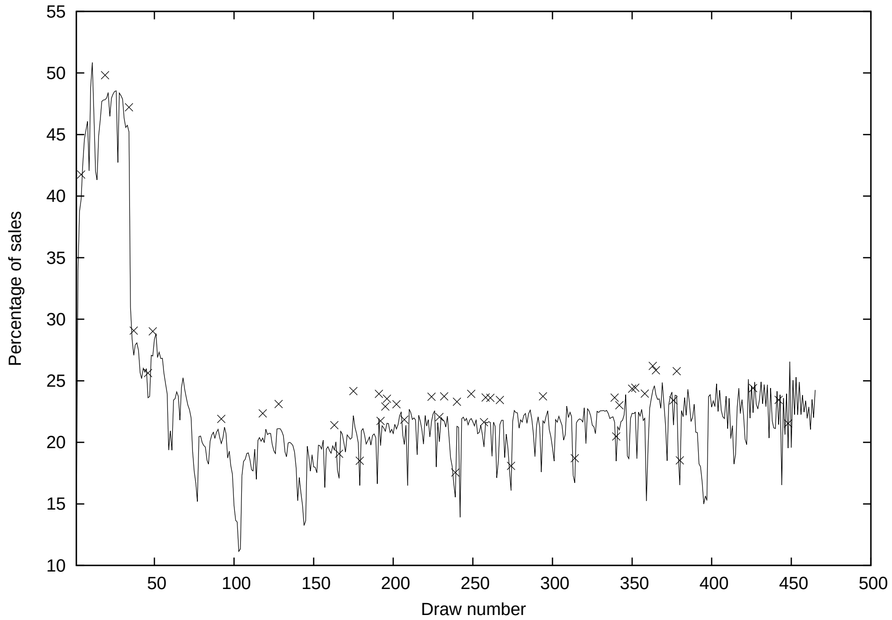

crosses indicate draws for which there was a jackpot winner in Spain

Detecting the tiny signal of the compatriot-win effect amidst the 'noise' of the large fluctuations in sales seems, from consideration of these plots, almost akin to extracting radium from pitchblende, or finding the Higgs boson. Without recourse to a full econometric model with all its assumptions, what we sought was a statistical test, based on relative sales data, which made as few assumptions as possible.

3. A statistical test for the compatriot win effect

3.1 Method

Let there be N_t countries at draw *t*, and M_{it} jackpot winners for the *i*th country at the *t*th draw. Let sales be Q_{it} tickets, and define $n_t = \sum_{i=1}^{N_t} M_{it}$, $V_t = \sum_{i=1}^{N_t} Q_{it}$, and the proportions of tickets sold in country *i* at draw *t* (relative sales) be $q_{it}=Q_{it}/V_t$.

Let q_{it+} be the same proportion Δ draws ahead, and $r_{it}=(q_{it+}-q_{it})/q_{it}$ be the proportional increase in relative sales from draw *t* to draw $t+\Delta$. Thus, if sales increase by 10% after a win, r_{it}=0.1. Then consider the statistic

$$S_t' = \sum_{i=1}^{N_t} r_{it}(M_{it} - n_t q_{it}) \tag{1}$$

This is large if wins lead to increased sales, or if lack of a win (in a draw where someone from another country wins) decreases sales. Conditioning on the total number of wins n_t, $E(M_{it}) = n_t q_{it}$, so S_t' has zero expectation. The statistic S_t' is the covariance of the number of wins and the proportional increase r_i, as can be seen by rewriting it as

$$S_t' = \sum_{i=1}^{N_t} (r_{it} - \overline{r}_t)(M_{it} - n_t q_{it}), \text{ where } \overline{r}_t = \sum_{i=1}^{N_t} r_{it} / N_t .$$

In fact, since $\sum_{i=1}^{N_t} r_{it} q_{it} = 0$, we can write simply $S_t' = \sum_{i=1}^{N_t} r_{it} M_{it}$.

In using the statistic S_t', we had to consider the issue of conscious selection. 'Conscious selection' refers to the propensity of players to choose numbers non-randomly and for their choices to be correlated with each other. For example, a culture may favour certain numbers as 'lucky' or even 'easy to remember' and a disproportionate number of entries will then feature these number choices. Because entries therefore cluster on certain number combinations, the effect will be to increase the probability of there being either zero or multiple winners among the players from that culture.

In the absence of country-specific conscious selection, under the null hypothesis that the distribution of wins among countries does not affect subsequent sales, numbers of wins $\mathbf{M}_t$ are random variables from the multinomial distribution with n_t wins, where $\mathbf{M}_t$ is the vector of random variables $(M_{1t}...M_{N_t t})^T$. An exact test can be performed by simulating a large number of sets of the $\mathbf{M}_t$. The p-value of a 1-sided test is then computed as the proportion of simulations for which the simulated test statistic is at least as large as S_t'. This is a use of the Monte-Carlo method to compute a difficult integral.

We need to cumulate values of the statistic across all T pairs of draws that are Δ draws apart, for which some country won at the first draw. There is a slight complication because sometimes new countries have entered at draw $t+\Delta$; in such cases, only the N_t countries common to both draws are used. How should these statistics be cumulated to give the best combined test statistic?

The covariance matrix V_{ij} for the M_{it} is $V_{ij} = n_t \left\{ q_{it}\delta_{ij} - q_{it}q_{jt} \right\}$ where δ_{ij} is the Kronecker delta, giving the variance of S_t' as

$$\sigma_t^2 = n_t \left\{ \sum_{i=1}^{N_t} q_{it} r_{it}^2 - \left(\sum_{j=1}^{N} q_{it} r_{it} \right)^2 \right\} = n_t \sum_{i=1}^{N_t} q_{it} r_{it}^2 \ .$$

One could average the S_t' regardless of variance, average them weighted by σ_t^{-2} (the minimum mean-squared error choice) or weight by σ_t^{-1}, etc.. We adopt this last choice, because the r_{it} vary wildly from draw to draw, and σ_t varies on the same scale, so that from the data, S_t'/σ_t is roughly normally distributed with variance near unity. This means that we will not have a few pairs of draws dominating the analysis. Hence we have finally the test statistic

$$S = \sum_{t=1}^{T} \sum_{i=1}^{N_t} r_{it} \left(M_{it} - n_t q_{it} \right) / \sigma_t = \sum_{t=1}^{T} \sum_{i=1}^{N_t} r_{it} M_{it} / \sigma_t \tag{2}$$

The test is based on the assumption that the Euromillions draw is random. This is an uncontroversial assumption, supported by our confirmation above that total numbers of winners by country reflect the pattern of sales. Our only other assumption, that conscious selection operates identically in all countries, is more questionable. Cultural differences might be reflected in different number preferences or else there might be differences in the proportions of tickets sold where the player has opted to have the computer choose his numbers (randomly) for him. Any heterogeneity of conscious selection across countries would invalidate the test, because the numbers of wins would no longer be random variables from the multinomial distribution. Multiple wins would tend to occur mainly within single countries, where bettors in a country would be disproportionately likely to use the same combinations of numbers.

The problem is overcome by restricting the pairs of draws to pairs where the earlier draw showed only one win, n_t=1. However, it is interesting to know whether this heterogeneity really exists. Therefore, a randomisation test for heterogeneity of conscious selection was performed, using as test statistic

$$U = \sum_{t=1}^{T} \sum_{i=1}^{N_t} \left\{ \left(M_{it} \left(M_{it} - 1 \right) / q_{it} \right) - n_t (n_t - 1) \right\} \tag{3}$$

where $E(U)$=0 and $\mathrm{var}(U) = \sigma_t^2 = 2n_t(n_t - 1)^2$. The term $M(M-1)$ is zero unless $M \geq 2$, when it is large if wins cluster in countries. This test statistic was arrived at by starting with $M_{it}(M_{it}-1)$ and then choosing the simplest possibility. As before, the p-value of the test is found by generating a large sample of values of U under the null hypothesis that the $\mathbf{M}_t$ derive from a multinomial distribution.

3.2 Results

In table 2, p-value 1 relates to the test for the existence of a compatriot win effect. The null hypothesis, that there is no relationship between the distribution of winners across countries and the subsequent distribution of sales across countries, is rejected for up to 12 subsequent draws (figure 5 illustrates the pattern of p-values in graphical form). The asymptotic version (not shown) that uses mean and variance predicted under H_0 and assumes normality gives similar p-values, for example for the subsequent draw, the randomisation test gives p=0.00214 and the asymptotic test gives p=0.00199. Here $S/\sqrt{T}$ is assumed standard normal, T being the number of pairs of draws.

Next, the test was repeated but restricting to n_t=1 to avoid the problem of potential conscious selection heterogeneity (p-value 2). Some p-values are higher but generally the results remain highly significant up to 12 draws.

Finally, p-value 1A and p-value 2A show corresponding results where the randomisation test was based only on pairs of data within the period when the game was offered just once a week. The pattern of results remains the same.

There is a risk that multiple hypothesis testing can lead to spurious conclusions. This can be addressed by applying a Bonferroni correction in which the significance level α required for n independent tests reduces to $\frac{\alpha}{n}$. If we take the first ten draws in table 2 as being of particular interest, for significance at the 5% level, we require a p-value of 0.005 or lower. We still have four p-values below this level, two of them (draws 5 and 6) very much below it.

Table 2. Results of randomisation tests for a compatriot win effect

Draws after a win	p-value 1	p-value 2	pairs	p-value 1A	p-value 2A	pairs
1	0.00214	0.00786	158	0.00139	0.00648	137
2	0.00420	0.04654	157	0.00708	0.05283	137
3	0.04275	0.00851	156	0.04974	0.01350	136
4	0.00509	0.00774	155	0.00868	0.00529	136
5	0.00010	0.00005	155	0.00023	0.00018	136
6	0.00001	0.00057	154	0.00001	0.00009	135
7	0.00441	0.00606	153	0.00574	0.01263	134
8	0.04778	0.18015	153	0.04364	0.16641	134
9	0.00823	0.02694	152	0.01333	0.06116	134
10	0.03760	0.13153	152	0.11212	0.21164	134
11	0.00735	0.01069	152	0.01380	0.01950	134
12	0.01678	0.02337	152	0.01633	0.03943	134
13	0.49185	0.42041	152	0.43701	0.45979	134
14	0.03931	0.13101	150	0.03838	0.14774	133
15	0.06331	0.08383	150	0.06605	0.12477	133
16	0.23386	0.58399	149	0.10774	0.42657	132
17	0.22842	0.31978	149	0.15599	0.28621	132
18	0.11048	0.08866	147	0.17432	0.11456	131
19	0.11007	0.07886	146	0.06745	0.06856	130
20	0.48612	0.67798	145	0.29127	0.50152	129
21	0.06023	0.09979	144	0.05097	0.12323	128
22	0.12508	0.15872	144	0.21573	0.21718	128
23	0.20068	0.10453	144	0.20893	0.14636	128
24	0.17776	0.30293	142	0.17164	0.35262	127
25	0.27283	0.26851	142	0.29236	0.32883	127

p-value 1 is for the randomisation data based on all data and p-value 2 for the randomisation data using only single-win draws; the suffix A indicates that the test was carried out only for the period when draws were once- rather than twice-weekly

The figures for the number of pairs used in each test relate to p-values 1 and 1A. Tests with p-values 2 and 2A used only draws with a single jackpot winner and therefore the number of pairs used was lower. However, the numbers are not reported to avoid clutter (24.7% of draws with a winner had a single winner, so the number of pairs was then typically about one-quarter lower than shown)

Figure 5. p-values from test for the presence of the compatriot win effect (logged scale) by draws from win, with the 5% significance level (horizontal line)

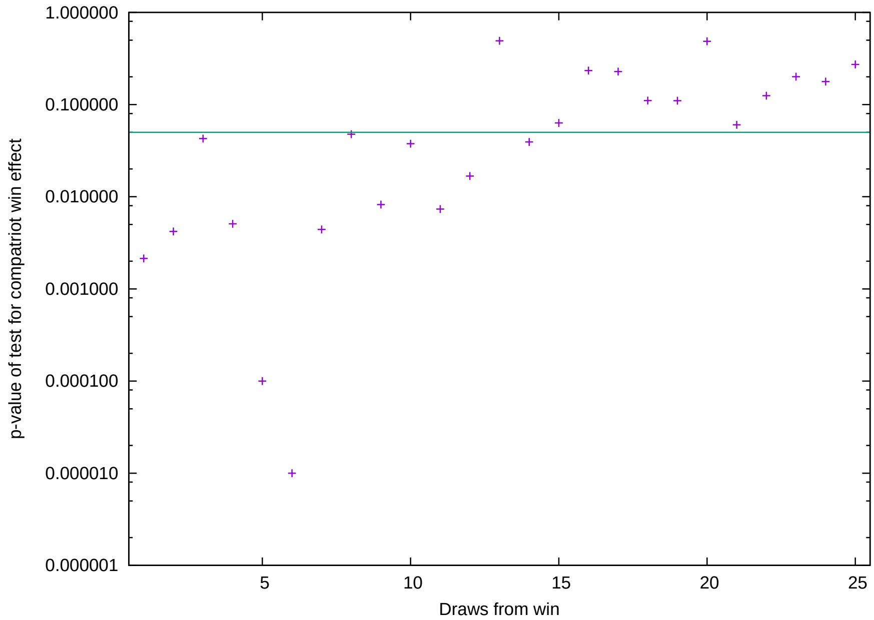

Hence the randomisation tests show clearly that there is a compatriot-win effect up to at least twelve draws after a compatriot win. The version of the randomisation test which is restricted to single prize winner draws has the randomness of the lottery as its only assumption. It must be rare in econometrics to have a test that does not make some debatable assumptions and that is an exact test, i.e. does not rely on asymptotic assumptions. This is possible here because of the complete randomness of the lottery, which is akin to a randomised controlled trial: at each draw, a country is randomly chosen to win, and the proportionate increase in relative sales is measured.

The test for heterogeneity of conscious selection rejected the null hypothesis with p=0.00573, showing that there is indeed heterogeneity of conscious selection among countries. This means that multiple wins tend to cluster by country.

It may be wondered whether the test for the presence of the compatriot-win effect presented here can really be valid. Of course, a win is more likely if sales are very high in a country for some reason, and they will usually decrease again for the next draw. So will there be a bias? The answer is no because the r values for other countries will be positive, because they have a lower proportion of sales this draw, giving a zero mean. On average, the test statistic is still zero. In fact, sales for the draw with a win could be fixed at any chosen values arbitrarily; only if the effect is present will higher values of r_{it} correlate with higher values of numbers of wins M_i giving a positive expected value.

We can also carry out an empirical check of the methodology. Reversing the order of draws and then applying the test, we are in effect testing that an own-country win could produce a change in sales for the *preceding* draw, which is clearly impossible. Doing this, there is no significant effect for previous draws. For the immediately preceding draw, we find a p-value of 0.63. There is no significant effect until the 12th previous draw, when p=0.017, and then there are no significant effects up to the 50th draw. One significant p-value in 50 is quite consistent with chance. Thus, when there is no compatriot-win effect, the test does not find one. This is encouraging.

We also sought to apply the test to individual countries using the statistic (2). However, these tests did not give significant results: evidently the effect was not strong enough to be detected with these smaller data sets.

We also repeated the test using only data from a subset of countries, the four largest contributors to Euromillions sales, France, Portugal, Spain, and UK. The pattern of significance was weaker than when data from all nine countries were employed but the randomisation test did show significance for the first, fifth and sixth draws with *p*-values of .0187, .0091 and .0288 respectively. When we reversed the draw ordering and applied the test, no effects close to significance were found.

Finally, we note that, in principle, because we model relative sales, the finding that the share of sales in a country tends to increase following a national win could reflect disillusion with the game (and falling absolute sales) in countries without a winner. However, it seems intuitively more plausible to associate the changes in shares of sales with a positive effect in winning countries, where a national winner will receive media publicity. In losing countries, no special news has been generated since it is normal for there to be no jackpot winner and therefore we do not expect there to be sufficient media coverage to generate negative sentiment. For these reasons, we refer to the impact on relative sales in a winning country as a 'compatriot win effect'.

4. Estimation of the magnitude of the compatriot win effect

4.1 Method

The demonstration that there is a statistically significant compatriot win effect naturally leads on to the question of whether the effect is large enough to be economically significant. Obtaining estimates unavoidably involves making more assumptions than was necessary for the existence test.

We expected that the impact of a national jackpot win might be greater where the amount collected by the winner was large since then the win is likely to receive more media attention and capture greater the public imagination. On the other hand, media focus when there is a large win might remind residents of that country that a large jackpot has gone and that the next draw will begin the rollover cycle again with a comparatively modest prize. This could depress relative sales. Allowing for either direction of effect, our modelling included both a constant term and an interaction term, number of winners multiplied by the size of interacting win. Our calculation of the estimated magnitude of the size of the effect was then based on the values of $\hat{\beta}_1$ and $\hat{\beta}_2$ that minimised the weighted sum of squares

$$SS = \sum_{t=1}^{T} w_t \sum_{i=1}^{N_t} \left\{ r_{it} - \beta_1 \left(M_{it} - n_t q_{it} \right) - \beta_2 \left(M_{it} - n_t q_{it} \right) \left(x_t - \overline{x} \right) \right\}^2 \quad (4)$$

Note that the mean jackpot per winner $\overline{x}$ is subtracted in the second factor of the coefficient of β_2. This does not affect the model being fitted, but it changes the interpretation of β_1, which is to be understood as estimating the proportional increase in relative sales when the payout per winner is at its mean. β_2 adjusts this estimate to take account of the size of the payout.

4.2 Results

The left columns in Table 3 presents estimates of the size of the compatriot win effect given by $100\hat{\beta}_1$, which is the percentage change in national relative sales predicted when there is a country winner of an amount equal to the mean jackpot payout in the data set. Figure 6 presents the estimates for up to 50 subsequent draws in graphical form. Estimation used the whole period of the data; results were very little changed when the period with twice-weekly draws was excluded.

Consistent with the findings from the existence test (where no account was taken of the size of win), the 'lucky' country experiences elevated relative sales for eleven draws but, beyond that, the estimated effect, though usually positive, is usually far from statistical significance. The size of the estimated effect on the country's share of sales from an 'average' jackpot win is modest, never more than 2% (of itself). It is scarcely surprising that such a small effect was undetectable from our initial inspection of the raw time series. However, this outcome changes when we consider very large jackpot payouts since the estimate of β_2 (which proved to be positive) now makes an important contribution to the predicted impact on relative sales. The final two columns of Table 3 show estimated effects on relative sales when the value of the payout is at the maximum observed in the data set. Figure 7 shows these effects for a more extended period.

Table 3. Estimates of the size of the compatriot win effect on subsequent relative sales (%) with payout size at its mean and at its maximum value

	jackpot at its mean		jackpot at its maximum	
Draws after a win	effect size	standard error	effect size	standard error
1	0.45	0.26	2.49	4.89
2	0.44	0.28	-4.90	4.01
3	0.43	0.43	8.02	5.20
4	0.90	0.45	3.19	4.82
5	1.91	0.51	7.71	5.68
6	1.73	0.48	3.26	4.60
7	1.22	0.57	9.18	5.23
8	0.85	0.54	0.35	5.80
9	1.34	0.60	11.07	6.18
10	0.93	0.56	1.31	5.57
11	1.40	0.64	11.04	6.27
12	0.53	0.55	4.59	4.89
13	-0.13	0.61	6.89	5.85
14	0.49	0.56	5.93	5.39
15	0.85	0.61	4.27	6.17
16	0.15	0.59	-0.80	5.90
17	0.30	0.64	-4.51	6.65
18	0.58	0.63	-2.81	6.31
19	0.62	0.61	10.45	6.52
20	-0.27	0.62	-0.44	6.20
21	0.94	0.67	12.00	6.00
22	0.67	0.64	2.66	5.50
23	0.61	0.73	8.50	6.74
24	-0.06	0.69	2.17	6.39
25	0.26	0.70	6.36	6.96

The initial impact on the country's share of sales is typically now 5-10% in the first twelve draws. Moreover, the effect then grows over time, to around 20% fifty draws after the happy event. From the way the model is constructed, it is relative sales which are predicted. A big jackpot in this member country would lead to declining market shares being predicted for all the other countries: for 'unlucky' countries at this draw, M_{it}-n_tq_{it} is negative and this leads to a fall in predicted market share, which is as it

Figure 6. Estimate of percentage increase in relative sales following a compatriot win with payout set to its mean

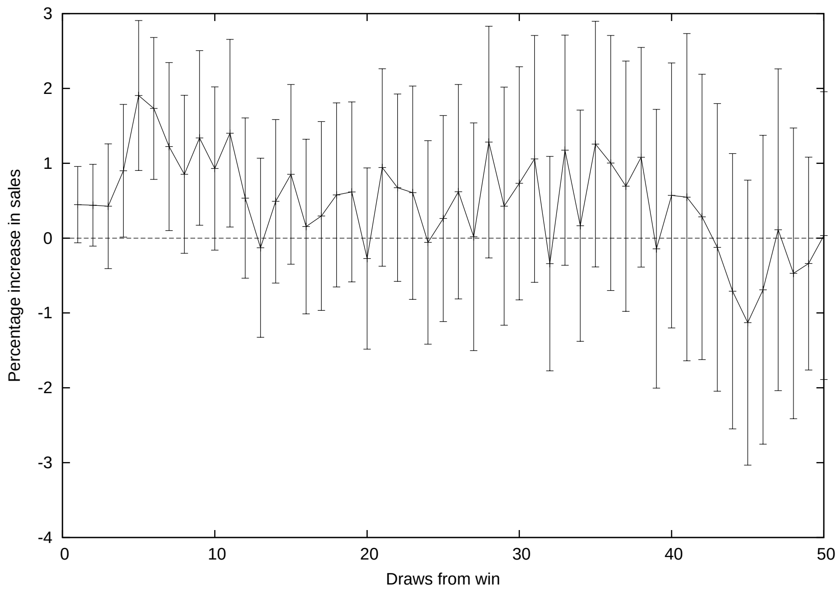

bars indicate 95% confidence intervals

Figure 7. Estimate of percentage increase in relative sales following a compatriot win with payout set to the maximum observed in the data

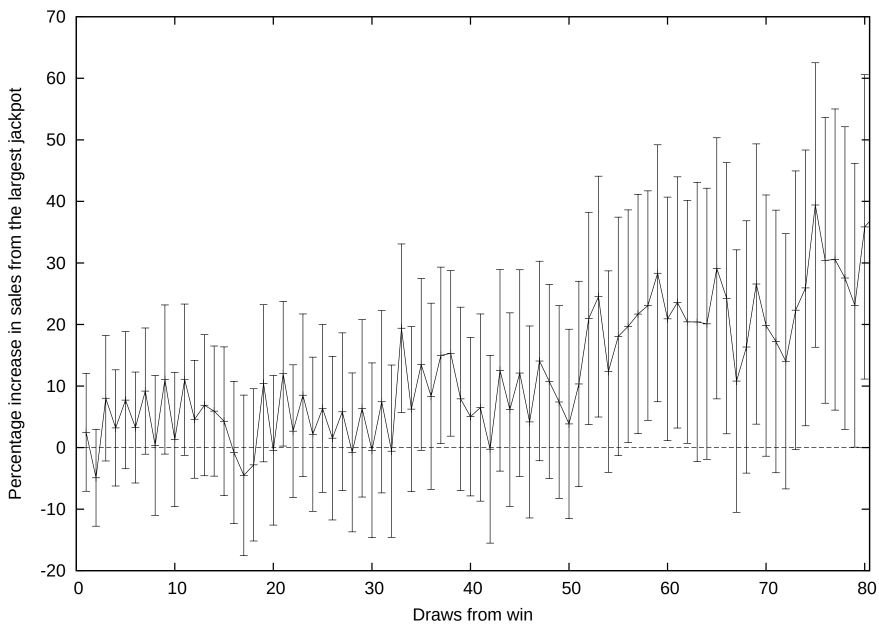

bars indicate 95% confidence intervals

should be because these countries have not produced one of the winners. If there is in fact no disappointment or national envy effect in other countries, a 20% increase in relative sales in the winning country would then imply a greater than 20% increase in absolute sales in the winning country. Clearly impacts of these orders of magnitude are likely to be economically significant.

As interesting as the large effect from the biggest jackpot wins is the long period over which it builds up. Economists typically explain that long-run effects may exceed short-run effects because of the phenomenon of habit formation, captured in econometric models by including lagged sales in the sales equation. Mostardinha *et al*. (2006) provide an alternative perspective when they discuss ‘exponential bursts’ in Euromillions sales during sequences of consecutive rollovers. They hypothesise that the population of potential buyers includes a proportion of ‘susceptible’ individuals who may be persuadable to become players either from media coverage or by the example or word of mouth recommendations of others. In this framework, media coverage of a big jackpot payout would trigger some new participants to join the market and their behaviour would itself draw in others in further rounds. A small lottery payout would not perhaps be a powerful enough trigger to initiate a long period of expansion of the market; but a very big payout with front page coverage just might be.

Whereas lottery economists have usually proceeded from the notion of potential players being individual decision takers, Mostardinha *et al*. (2006) essentially modelled herding behaviour and a substantial expansion of the Euromillions market following a major win in a country could be predicted from thinking about it as being a similar process as that behind an idea ‘going viral’ in the online World. Whether this story has relevance in the present context cannot be determined by our data. However, it might be interesting, were the data available, to look at time series of the proportion of the national population playing Euromillions in the period following a major national jackpot win.

5. Concluding remarks

We set out to establish whether there was a compatriot win effect in the market for Euromillions lottery tickets. We devised a test which was close to assumption-free and it demonstrated the typical presence of a compatriot win effect in the twelve draws following a win. When we modelled the size of the effect, it proved to be modest for average payouts but large and more durable for the very highest payouts. This should enhance understanding of how the market for Euromillions, now more than a decade old, has evolved. It should therefore aid decision taking by managers of lottery agencies in the Euromillions bloc, to the ultimate benefit of the governments and 'good causes' which claim the net revenue from the game.

For lottery agencies Worldwide, the findings add to those from Guryan and Kearney (2008) in providing evidence that jackpot wins in particular geographical areas have the potential to elevate sales for an extended period. Of course, the paper by Guryan and Kearney (2008) relates to a different game, adopts a different statistical methodology and focuses on much smaller sales areas. Nevertheless their paper and ours may be argued to have the common implication that geographically focused marketing campaigns to exploit and extend the duration of a winner effect might be justified.

The results could also be further explored in the context of choice of game design. Euromillions is just one of several prominent lotto games where it has been made harder to win the jackpot (the odds have been lengthened from about 80m to 1 to 116m to 1). Such redesigns are intended to lengthen the typical gap between draws which produce a winner, with consequent accumulation of larger jackpots through multiple rollovers. High jackpots are argued to attract attention and additional sales; but there is a trade-off because there may also be the risk of disillusion if it comes to be perceived as very unlikely that the current draw on sale will make even a single player rich. In evaluating the trade-off, the present paper may be relevant because it points to potential players in a country responding to (and, by implication, remembering) a national win for a period. At the end of that period, sales return to 'normal' but could be boosted again by another national win. Game design could exploit this pattern by seeking to induce a pattern of wins that would typically (in the larger countries) produce a national winner at appropriately spaced intervals.

One interesting feature in the evolution of the game has been the extent to which trends in sales have differed across the nine member countries. Limitations of space preclude the presentation of the plots for each of them (these plots are available as supplementary material) but among the most striking features has been the clear upward trend in the share of sales accounted for by the UK (figure 2). This could be related to the effect of national jackpot wins in the UK market. Of course, our preliminary analysis did show that countries received their 'fair share' of jackpot wins given their contribution to aggregate sales. However, because the sensitivity of UK sales to jackpot size is particularly high, then it would be expected that its jackpot wins would be disproportionately likely to be in high payout draws. Such indeed has been the case. Of the eight highest payouts observed in our data set, four were for tickets purchased in the UK. It seems likely therefore that the greater willingness of UK players to increase their purchases when the jackpot is high will have generated relatively frequent headline-making payouts to UK players (Roger, 2011, noted an' implicit subsidy' to UK players from players in other countries to the extent that the average expected loss per ticket was lower in the UK since a greater proportion of UK sales occurred when the draw offered better value). We have shown that each such payout has the potential to have a strong and long-lasting effect on subsequent sales.

What the findings say says about rationality of players is less clear. Of course, the official announcement of the result will always include the nationality of the winner(s) of the jackpot prize**.** That some impact on relative sales across countries is observed could partly be explained by irrational response: there may be some people whose assessment of the probability of a win is affected by past results. However, large behavioural changes are only predicted when the prize that has been won is very high. National media coverage will then take the information that there has been a national winner to a much wider audience than those who seek out the official results. When they become players (or buy more tickets), it cannot be determined from the data whether their response is based on modifying their perception of the chance of winning or on emotional feelings of regret that they have missed out on the chance of winning or on an advertising effect which works simply by reminding the public about what they might have forgotten: that there is a product where spending €2 on a ticket just might make them rich beyond the dreams of avarice.

Acknowledgements

We are grateful to Professor O.D. Gulley (Bentley University) for helpful comments on an earlier draft.

References

Chen, S. H. and Chie, B. T. (2008) Lottery markets design, micro-structure and macro-behavior: An ACE approach. *Journal of Economic Behavior and Organization*, 67, 463-480.

Cook, P. J. and Clotfelter, C. T. (1993) The peculiar scale economies of lotto. *American Economic Review*, 83, 634-643.

Forrest, D. and Gulley, O. D. (2014) The efficiency of lottery markets. In *The Oxford Handbook of the Economics of Gambling* (eds L. Vaughan Williams and D. Siegel). Oxford, UK: Oxford University Press.

Forrest, D., Simmons, R. and Chesters, N. (2002) Buying a dream: Alternative models of demand for lotto. *Economic Inquiry*, 40, 485-496.

Gulley, O. D. and Scott, Jr., F. A. (1993) The demand for wagering on state-operated lottery games. *National Tax Journal*, 45, 13-22.

Guryan, J. and Kearney, M. (2008) Gambling at lucky stores: Evidence from state lottery sales. *American Economic Review*, 98, 458-473.

Mostardinha, P., Durana, E. J. and Vistulo de Abreu, F. (2006) The econophysics in the Euromillions lottery. *European Journal of Physics*, 27, 675-684.

Roger, P. (2011) La demande de grilles d'Euromillions: Une comparison internationale. *Revue Économique*, 62, 29-55. An earlier (2008) version in

English is available through the Social Science Research Network at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1300436.