



Københavns Universitet



Predicting Lotto Numbers

Jørgensen, Claus Bjørn; Suetens, Sigrid ; Tyran, Jean-Robert

Publication date:
2011

Document Version
Early version, also known as pre-print

Citation for published version (APA):
Jørgensen, C. B., Suetens, S., & Tyran, J-R. (2011). Predicting Lotto Numbers. Department of Economics, University of Copenhagen.

Discussion Papers
Department of Economics
University of Copenhagen

No. 11-10

Predicting Lotto Numbers

Claus Bjørn Jørgensen, Sigrid Suetens,
and Jean-Robert Tyran

Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K., Denmark

Tel.: +45 35 32 30 01 – Fax: +45 35 32 30 00

<http://www.econ.ku.dk>

ISSN: 1601-2461 (E)

Predicting Lotto Numbers

Claus Bjørn Jørgensen*, Sigrid Suetens[†] and Jean-Robert Tyran[‡]

March 2011

Abstract

We investigate the “law of small numbers” using a unique panel data set on lotto gambling. Because we can track individual players over time, we can measure how they react to outcomes of recent lotto drawings. We can therefore test whether they behave as if they believe they can predict lotto numbers based on recent drawings. While most players pick the same set of numbers week after week without regards of numbers drawn or anything else, we find that those who do change, act on average in the way predicted by the law of small numbers as formalized in recent behavioral theory. In particular, on average they move away from numbers that have recently been drawn, as suggested by the “gambler’s fallacy”, and move toward numbers that are on streak, i.e. have been drawn several weeks in a row, consistent with the “hot hand fallacy”.

JEL Codes: D03, D81, D84

Keywords: gambler’s fallacy, hot hand fallacy, representativeness, law of small numbers

*University of Copenhagen, e-mail: clausbjorn@gmail.com

[†]CentER, TIBER, TILEC, Tilburg University, e-mail: S.Suetens@uvt.nl, corresponding author.

[‡]University of Vienna and University of Copenhagen, e-mail: Jean-Robert.Tyran@univie.ac.at. We thank Dansk Spil for providing us with data. We also thank Dirk Engelmann, Lapo Filustrucchi, Uri Gneezy, Pedro Rey Biel, Adriaan Soetevent, Jan van Ours, Frederic Vermeulen, and participants in seminars at the University of California San Diego, the University of Copenhagen and Tilburg University for valuable comments.

1 Introduction

Predicting lotto numbers is a pointless exercise. Because lotto numbers are truly random, my guess is as good as yours, and any number picked is equally likely to win. Yet, some lotto players seem to believe that they can predict lotto numbers from previous draws. In this paper, we show that lotto numbers picked are systematically related to previous draws in the aggregate and that some players are particularly prone to the belief that lotto numbers can be predicted. We infer this belief from how players react to previous draws and show that the emerging patterns of number picking are consistent with recent behavioral theory (Rabin, 2002; Rabin and Vayanos, 2010). In particular, we show that players tend to avoid numbers that have been drawn in the previous week but tend to favor numbers that are “on streak”, i.e. have been drawn in several consecutive weeks. Using panel data from the Danish State Lottery allows us to track choices of individual players and to provide unusually clean field evidence for the “law of small numbers”. We show that the “gambler’s fallacy” and the “hot hand fallacy” occur in lotto gambling, that the fallacies are systematically related, and that being prone to these biases is costly. For example, players prone to the gambler’s fallacy spend 1.2 EUR more in an average week than other players.

Mounting evidence from the experimental laboratory and the field suggests that truly random processes are difficult to grasp for most people, and that many people tend to see patterns in data when in fact there are none. According to the “law of small numbers” (Tversky and Kahneman, 1971), people tend to mispredict random sequences because they expect small samples to “look like” large samples. For example, if asked to generate a random sequence as in a coin toss, many people predict too many switches between head and tails because they falsely believe that head and tails should appear in equal proportion even in a small sample (see, e.g., Bar-Hillel and Wagenaar, 1991; Rapoport and Budescu, 1997). Or when playing the roulette, people tend to expect that a black number is “due” after observing a sequence of red numbers (e.g. Croson and Sundali, 2005). The belief in frequent reversals in random sequences has been dubbed the “gambler’s fallacy” (Tversky and Kahneman, 1971).

A second type of misprediction of random sequences is the “hot hand fallacy”. People who believe in hot hands overpredict the probability that streaks continue. For example, bets on basketball players who scored unusually well recently tend to be too high (Gilovich et al., 1985; Camerer, 1989). The reason for this type of misprediction seems to be that people find it hard to believe that long streaks are the result of a random process because long streaks are perceived to be unrepresentative of a random sequence. Once a long streak is observed, believers in the

“hot hand” expect the streak to continue.

While the gambler’s fallacy and the hot hand fallacy seem to be contradicting examples of people seeing patterns in random data when there are none, Rabin (2002) and Rabin and Vayanos (2010) have recently developed models relating the two types of fallacies: the same people who fall prey to the gambler’s fallacy also tend to develop the hot hand fallacy as they observe a streak getting longer. The intuition behind the theory is straightforward. Suppose a person is prone to the gambler’s fallacy — because she believes in the law of small numbers — and suppose that the person is uncertain about the true probability underlying a random event. Such a person starts to doubt about the true probability when observing a long streak because this does not correspond to what she believes a random sequence should look like. As a consequence, such person starts to believe in the continuation of the streak, i.e. develops the hot hand fallacy. Uncertainty about the true probability is key for the argument. The authors argue that people who are prone to the gambler’s fallacy but are absolutely certain about the true probability underlying random events will continue to act according to the gambler’s fallacy even when observing long streaks.¹

Recently, Asparouhova et al. (2009) found support for the prediction that the gambler’s fallacy turns into the hot hand fallacy for some players in a laboratory experiment where participants are asked to predict the next observation in a random walk process. The authors also show that people are more likely to exhibit the hot hand fallacy as the perception of randomness of the data-generating process decreases. This result resonates well with findings in the psychology literature which suggest that the gambler’s fallacy is mostly observed when events are believed to be totally random while the hot hand fallacy arises when events are perceived to be less random (involving, for example, a human factor as suggested by Ayton and Fischer, 2004). Burns and Corpus (2004), for example, show that people who are confronted with the same streak of events but stemming from scenarios that differ in the degree of randomness respond differently to streaks: with less random scenarios, people were more likely to believe in the continuation of a streak, and vice versa for a reversal of a streak. This might explain why in the literature the gambler’s fallacy is mostly examined in the context of certainty about the true probability of random draws (like lotto, roulette play or coin flipping), while the hot hand fallacy is often studied and observed in a context with uncertainty (involving an element of control or skill, as

¹Rabin and Vayanos (2010) assume the gambler’s fallacy to be the “primitive” bias. The assumption seems plausible given that it arises in simple experiments involving fair coins. See Appendix A.1 for a summary of their theoretical model interpreted in terms of lotto play.

in basketball, or including the probability of a false coin as in Offerman and Sonnemans, 2004).

Using lotto data to investigate misperception of random events is particularly apt because the true randomness of the game is tightly controlled (often by government regulation) and made transparent to players (e.g., by drawing balls from an urn and by airing the draws on TV). Thus, previous studies have used lotto data to study the gambler's fallacy. Clotfelder and Cook (1993) investigate the fixed prize Maryland lottery and Terrell (1994) the parimutuel New Jersey lottery. Both studies find that after a number has been drawn, the amount bet on the number falls sharply and then gradually recovers. Both of these studies use aggregate-level data while we use individual-level panel data. The advantage of using panel data is that we are, to the best of our knowledge, the first to observe the *reaction* of individual players to the history of draws. We are therefore not only able to see if a bias is present for the average player but we can also test if players who are prone to gambler's fallacy are also prone to hot hand fallacy as streaks get longer, i.e. to test the theory by Rabin (2002) and Rabin and Vayanos (2010) in the field. Roulette shares some of the desirable properties of lotto, and a remarkable paper by Sundali and Croson (2006) uses individual-level data to investigate the two fallacies. These authors use data from videotapes to show that players who exhibit the gambler's fallacy, i.e. bet less on numbers that came out of recent spins, are more likely to increase their bets after winning and vice versa. They are thus more likely to believe to have a hot hand with respect to the amount bet.

The paper is organized as follows. Section 2 describes the lotto data and defines the main variables of interest. We are able to track lotto numbers picked by individual players over time because we investigate lotto play over the Internet in Denmark where lotto players are required by law to be uniquely identified. This allows us to construct variables measuring the change in numbers picked by an individual over time. Section 3 analyzes the aggregate reaction of lotto players to the recent drawing history of lotto numbers. We find that, on average, players place two percent fewer bets on numbers drawn in the previous week than on numbers not drawn, as long as those numbers are not on a streak. They also increase their bets on numbers drawn in the previous week, the longer the streak. In Section 4 we use individual-level data to show that 9.2% of the players are biased in the sense that they significantly react to the previous week's drawings. Of those players who change their bets significantly in reaction to previous weeks' drawings, 64% do so by decreasing their bets. We also find that among those avoiding numbers drawn in the previous week, 57% tend to increase their bets on numbers on streak as the streak gets longer. Section 5 shows that being biased is costly. We find that players prone to the gambler's fallacy lose more money because they buy more tickets and, because of the pari-mutuel structure of

lotto, win less money than a random player if they happen to win. Section 6 concludes.

2 The data

We analyze data from lotto played on Saturdays in Denmark over the Internet (henceforth lotto for short) covering 28 weeks in 2005. Lotto is organized by a state monopoly (Dansk Spil). Every Saturday, 7 balls are drawn from an urn containing 36 balls numbered from 1 to 36, which is aired on state TV. The price of a lotto ticket is about EUR 0.40 (DKK 3).²

The payout rate is set to 45% by law and the remainder of the revenues mostly goes for “good causes” and to a small extent to the general government budget. Lotto has a parimutuel structure as the payout rates are fixed per prize category and the prize money per category is shared among the winners in that category. One quarter of all payouts are reserved for the jackpot (7 correct numbers), and there are four graded prizes for having selected fewer correct numbers. If no-one wins the jackpot, it is rolled-over to the next week. In our data set (covering the second half of 2005), the average jackpot was about EUR 534’000 (4 million DKK), and the highest jackpot was 1.4 million EUR (10.2 million DKK). Prizes above DKK 200 are subject to a special tax of 15% but are otherwise exempt from income tax.

We investigate lotto numbers picked by players over the Internet in the last 28 weeks of the year 2005 (from week 25 to week 52). Lotto is normally played by purchasing hard-copy lotto tickets in vending booths like drugstores and supermarkets. Since 2002, lotto can also be played over the Internet (<http://danskespil.dk>). Lotto numbers can be picked in various ways in Denmark. Traditionally, players manually select 7 out of 36 numbers on each ticket they buy. However, we analyze numbers picked in “Systemlotto”. Here, players select between 8 and 31 numbers manually and let the lotto agency choose combinations of 7 out of these numbers.³ Our data has been provided directly by the state lotto agency and is unlikely to contain any error. All players in our dataset are identified by a unique ID-number, which allows us to track the choices of players over time. In total, 189’531 persons have played lotto over the Internet at least once in the second half of 2005. More than half of these (100’386) manually select their numbers in the traditional way, and 25’807 select numbers using Systemlotto.

²The following describes the rules of lotto at the time the data was collected. The prize structure has been modified since to yield higher jackpots.

³Other ways to play are “Quicklotto” where all numbers are selected randomly by the lotto agency and “Lucky-lotto” where players select up to 6 numbers manually and let the lotto agency choose the remaining numbers.

The reason for focusing on Systemlotto rather than the traditional manual selection is that Systemlotto is more appropriate to capture the belief in one’s ability to predict winning numbers and the reaction of number picking to recent draws. In fact, in Systemlotto players choose *numbers*, rather than *combinations* of numbers, as in the traditional manual selection. They choose fewer unique numbers than players who select in the traditional way which suggests that they are more likely to believe that a particular number is going to win. To illustrate, Systemlotto players pick less than half among the 36 available numbers (14 numbers in an average week, 8 in a modal week), while in manual selection players pick most available numbers (29 in an average week, 32 in a modal week). A particular focus of our study is on how players react to recent draws and in particular, streaks of draws. While (short) streaks of *numbers* are relatively likely, streaks of *combinations* of numbers are extremely unlikely and do not exist in our dataset.⁴

An advantage of our data set compared to laboratory data is that it reflects behavior of a heterogeneous pool of people and behavior is observed in a “natural” situation. In fact, lotto is quite popular in Denmark. For example, according to the lotto agency, about 75% of the adult Danish population have played lotto at least once. Yet, Systemlotto players are clearly not representative for the Danish population or even for the pool of internet lotto players. People playing Systemlotto — the name is suggestive — may be particularly confident that they can predict lotto numbers with some accuracy. In fact, Systemlotto players buy on average about twice as many tickets as other internet players (29 vs. 14 tickets per week; the medians are 19 and 10, respectively). Systemlotto is also especially popular with male players: 82% of the players are male compared to 73% for other selection devices.⁵ An advantage of our data compared to other studies on state lotteries is that we can track players in time, which makes it possible to investigate whether and how they change their number selection in response to recent drawing outcomes.

⁴While there are only 36 numbers, there are about 8 million ways to combine 7 out of 36 numbers ($36!/(7!(36-7)!) = 8'347'680$). The probability that the same combination occurs twice in a row is therefore about 1 in 70 trillion (7×10^{13}) in Danish lotto. Curiously enough, the same six numbers were drawn twice in a row in the Bulgarian lottery in September 2009. This event was considered to be so unlikely that it caused attention in the media and prompted the Bulgarian government to initiate an investigation for manipulation of the game.

⁵Male players buy significantly more tickets than female players, irrespective of the selection device, and Systemlotto players buy significantly more tickets than other players, irrespective of the gender ($p = 0.000$ in Mann-Whitney ranksum tests).

2.1 Independent variable

We relate numbers picked by players to winning numbers, i.e. numbers drawn by the lotto agency in previous weeks. The independent variable Streak_{jt} measures for each Lotto number j in week t whether (1) number j has been drawn in week $t - 1$ and (2) in how many consecutive weeks number j has been drawn:

$$\text{Streak}_{jt} = \begin{cases} 0 & \text{if number } j \text{ has not been drawn in week } t - 1, \\ k & \text{if number } j \text{ has been drawn in weeks } t - 1 \text{ to } t - k \text{ and not in week } t - 1 - k. \end{cases} \quad (1)$$

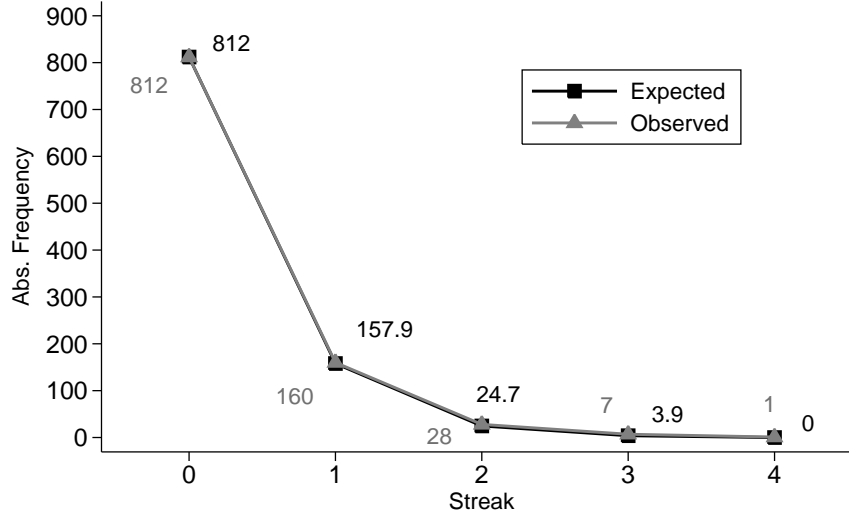
Figure 1 shows that the expected and observed frequency of Streak_{jt} are very similar for the 28 weeks of lotto drawings. For example, we observe 160 streaks of length 1 over the 1008 cases (= 36 numbers over 28 weeks) while 157.9 (= $29/36 * 7/36 * 1008$) such streaks are expected. Note that observing a streak of length larger than 4 is highly unlikely. For example, the probability of observing a streak of length 5 is 0.0002, and such streaks are therefore not observed in our dataset. Since lotto drawings are truly random, we find that observed and expected counts are not different according to a Chi-square test ($p = 0.743$).

2.2 Dependent variable

Our empirical strategy is to make inferences about the (unobservable) belief in the ability to predict winning lotto numbers more accurately than pure chance from observable reactions to previous drawings. That is, we infer that players think recently drawn numbers are more likely to win if they systematically prefer them and *vice versa* if they avoid them. A player is said to be more confident that a particular number is going to be drawn if he or she places more bets on it (i.e. buys more tickets including this number). More specifically, we define a variable “weight” showing how much money is bet on a number relative to other numbers. We then define a variable “reaction” measuring how much more a player is confident that a particular number is going to win conditional on previous draws. More specifically, “reaction” shows the change in “weight” placed on a particular number.

This empirical approach is rather straightforward but not perfect because it can only detect some of the fallacies that may be present. In other words, our approach tends to underestimate the true extent of fallacies. To illustrate, consider the gambler’s fallacy. Whether players are prone to the gambler’s fallacy, i.e. believe that a recently drawn number is less likely to be drawn again in the future, can only be measured for those players who have picked the number

Figure 1: Expected and observed frequencies of Streak_{jt}



Notes: The chart shows expected and observed Streak_{jt} for 1008 cases (i.e., $i = 1$ to 36 lotto numbers times $t = 1$ to 28 weeks). The expected frequency of $\text{Streak}_{jt} = k$ is $29/36 * (7/36)^k$.

previously in the first place.

Our empirical proxy for how confident a player is in his prediction is how much money he or she bets on a particular number relative to other numbers. Constructing this proxy is not entirely trivial because it depends on the number of “sets” of numbers, the number of tickets generated from each set, and the number of lotto numbers contained in each set. Recall that in Systemlotto, players pick one or more sets containing 8 to 31 numbers. For each chosen set, the lotto agency generates at least 8 tickets with different combinations (of subsets) of the chosen numbers.⁶ To illustrate the construction of our proxy, consider the following two examples.

Example 1 *Player A chooses a set of 10 numbers and Player B chooses a set of 24 numbers. Both A and B buy 120 tickets generated out of their chosen sets.*

Example 2 *Player C chooses a set of 10 numbers from 1 to 10 and a set of 10 numbers from 5 to 14. For each set, 8 tickets are generated. Player C thus buys 16 tickets in total.*

⁶The total number of tickets/combinations generated by Systemlotto out of a set of chosen numbers is positively related to the total number of lotto numbers a player chooses in the set. The exact relation depends on which of three “systems” players use to generate tickets/combinations. See Appendix A.2 for details.

In example 1, both players buy the same number (120) of tickets. Yet, it is plausible to assume that player A is more confident that (some of) his 10 numbers are going to win than player B who picks 24 numbers. Our proxy therefore gives each of the 10 numbers chosen by player A a larger weight (of 1/10) than each of the 24 numbers chosen by player B (1/24). In example 2, player C chooses two sets which partly overlap since the numbers 5 to 10 are elements of both sets. It seems plausible to assume that player C is more confident that one of the numbers chosen in both sets (5 to 10) is going to win than one of the numbers chosen in only one of the sets (1 to 4 and 11 to 14). Our proxy therefore gives numbers occurring in overlapping sets accordingly more weight than numbers occurring in only one set.

To compare weights across players, we normalize the number of times lotto number j is picked by player i in week t by the total number of lotto numbers picked by player i in week t across all sets. This variable sums to 1 across all lotto numbers for each i and each t and thus gives us information about the relative weight players put on particular numbers. We refer to this variable as Weight_{ijt} and define it as follows for lotto number j , player i and week t :

$$\text{Weight}_{ijt} = \frac{\# \text{ of times lotto number } j \text{ is picked in week } t \text{ by player } i}{\# \text{ of lotto numbers picked in week } t \text{ by player } i} \quad (2)$$

We are now ready to define the key variable Reaction_{ijt} which shows how players change relative weights on numbers from period to period. Again, the definition of this variable requires some care because some (i.e. low) numbers are generally more popular than others, and some (perhaps idiosyncratically “lucky”) numbers are more popular with particular players than others. In the construction of the variable Reaction_{ijt} , we therefore control for “baseline” choices of players, i.e. numbers that players choose irrespective of the recent history of drawings by differencing them out. To illustrate, consider a player who, for example, always chooses lotto number 2 in combination with other (time-varying) numbers. Suppose number 2 is drawn in week $t - 1$. If we do not correct for the player’s idiosyncratic preference, we would wrongly conclude that the player exhibits the hot hand fallacy. This consideration is especially important given that low numbers are more popular than high numbers. For example, the lowest 5 numbers (1 to 5) are picked more than 30% more often than the highest 5 numbers (32 to 36).⁷ Therefore, the dependent variable Reaction_{ijt} for lotto number j , player i and week t is defined as:

⁷The popularity of the lotto numbers is shown in Appendix A.3.

Table 1: Summary statistics of dependent variable Reaction_{ijt}

Minimum	-0.12500
Maximum	0.12500
Mean	0.00000
Standard Deviation	0.05098
Median	-0.00069
First Quartile	-0.03704
Third Quartile	0.03704
Number of data points	549'071

Notes: The summary statistics are based on data points where $\text{Reaction}_{ijt} \neq 0$. By definition, the mean is zero.

$$\text{Reaction}_{ijt} = \Delta \text{Weight}_{ijt} = \text{Weight}_{ijt} - \text{Weight}_{ijt-1} \quad (3)$$

Intuitively, a player is said to avoid or “move away from” a number in period t if $\text{Reaction}_{ijt} < 0$ and vice versa for favoring or “moving towards” a number.

In the regression analyses reported in Sections 3 and 4 we focus on data points that involve a change in weight on — a move toward or a move away from — numbers. The reason is that $\text{Reaction}_{ijt} = 0$ is difficult to interpret. Indeed, $\text{Reaction}_{ijt} = 0$ can have two totally different interpretations. It can refer to player i *staying away* from number j in week t : a player who has already moved away from a number in week $t - 1$ cannot move away from the same number in week t . Yet, it can also refer to player i *staying with* the same number j in week t : a player who has already moved toward a number in week $t - 1$ cannot move toward the same number in week t .

A majority of players seems to pick the same numbers week after week. In fact, out of the 25'807 Systemlotto players in our dataset, 17'318 have at least two consecutive observations such that Reaction_{ijt} can be measured at least once. Of these, 10'434 players do not change their weight on numbers at all, meaning that for these players $\text{Reaction}_{ijt} = 0$ for all j and t . About a quarter of players in our data set (6'884) pick different numbers in consecutive weeks, i.e. have $\text{Reaction}_{ijt} \neq 0$ for at least one j in one t , and our analysis below is therefore based on these players.

Table 1 provides summary statistics on the dependent variable Reaction_{ijt} . By construction,

Reaction_{ijt} varies between -0.125 and 0.125 with a mean of 0. Intuitively, the symmetric range results from the normalization and fact that for every move toward a number there is a move away from another number or other numbers of the same absolute size. The reason why the maximum is equal to 0.125 and the minimum to -0.125 is that at least 8 tickets are generated from each set and the maximum weight of a particular number is therefore $1/8$ ($= 0.125$). Hence, the maximum absolute change in weight — changing from a weight of zero to one of $1/8$ or the other way around — is 0.125. The median of Reaction_{ijt} is close to zero (-0.00069) and the interquartile range is [-0.03704, 0.03704]. These statistics show that unconditional on drawing history, there is no clear tendency for players to move away from ($\text{Reaction}_{ijt} < 0$) or toward ($\text{Reaction}_{ijt} > 0$) lotto numbers. In other words, there are no pure time trends in number picking. However, as shown next, the extent to which players avoid or prefer numbers is significantly related to outcomes of recent drawings.

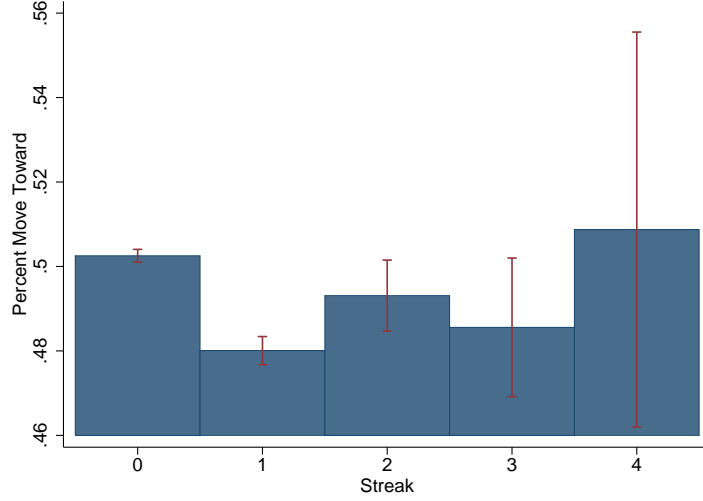
3 Aggregate reaction of lotto players to recent drawings

This section analyzes how players react to recent drawings by reporting the results from pooled regressions. Figure 2 shows the percentage of moves toward numbers across all data points where $\text{Reaction}_{ijt} \neq 0$ as a function of the streak length. Had all players picked numbers randomly, the proportion of moves toward numbers would be 50% for each streak length. In contrast, Figure 2 shows that, on average, players avoid numbers that have been drawn recently as long as the streak is not “too” long (lower than 4), and prefer numbers that are on a “long streak” (i.e. equal to 4). Thus, on average, players act in line with the gambler’s fallacy as long as streaks are not “too” long and switch to hot hand fallacy as streak get “long” which is in line with Rabin (2002) and Rabin and Vayanos (2010).

The figure shows that players tend to most strongly avoid numbers that have been drawn *only* in the previous week. In particular, the effect is most pronounced at $\text{Streak}_{jt} = 1$ compared to $\text{Streak}_{jt} = 0$. From Figure 2 we see that at $\text{Streak}_{jt} = 1$ there are 2% fewer moves toward numbers than at $\text{Streak}_{jt} = 0$, meaning that there are 2% fewer bets on numbers drawn only in the previous week (and not in consecutive weeks) than on numbers not drawn. Our finding is in the same ballpark as the one found in the lab by Asparouhova et al. (2009). They find that, on average, players reduce their probability estimate of continuation by approximately 0.9% for each unit increase in streak length for streaks of length up to three.

When comparing $\text{Streak}_{jt} > 1$ versus $\text{Streak}_{jt} = 1$ we see that at $\text{Streak}_{jt} > 1$ players decrease

Figure 2: Proportion of moves toward numbers as a function of the recent drawing history



Notes: The chart shows for each value of Streak_{jt} the proportion of moves toward numbers, i.e. the proportion of cases where $\text{Reaction}_{ijt} > 0$. Included are 95% confidence intervals.

their weight on average less than at $\text{Streak}_{jt} = 1$. Thus, as numbers are drawn in consecutive weeks, there is an increase in the move toward numbers compared to $\text{Streak}_{jt} = 1$.⁸

To test whether the effects illustrated in Figure 2 are statistically significant we adopt the following approach. We run two separate regressions. In the first regression, we estimate how players react to a lotto number being drawn in the previous week, irrespective of whether the number is on streak. We define the variable Drawn_{jt} as follows:

$$\text{Drawn}_{jt} = \begin{cases} 1 & \text{if } \text{Streak}_{jt} > 0, \\ 0 & \text{if } \text{Streak}_{jt} = 0. \end{cases} \quad (4)$$

The first regression is:

$$\text{Reaction}_{ijt} = \beta_0 + \beta_1 \text{Drawn}_{jt} + \epsilon_{ijt}, \quad i = 1, \dots, N, \quad j = 1, \dots, 36, \quad t = 1, \dots, T, \quad (5)$$

In the second regression, eq. 6, we estimate how players react to an increase in the streak of a number, *given that the number was drawn in the previous week*.⁹ Hence, the variable included in

⁸The confidence intervals at streak length 4 are much larger than those at other streak lengths because the number of observations at streak length 4 is much smaller than at other streak lengths.

⁹Because the independent variables are the same for each lotto player, it is not possible to include individual player effects in the pooled regressions.

Table 2: Pooled regression results: separate simple regressions

	Effect of Drawn _{jt}	Effect of Streak _{jt}
	Est. (s.e.)	Est. (s.e.)
Intercept	0.00044 (0.00006)***	-0.00199 (0.00026)***
Drawn _{jt}	-0.00223 (0.00028)***	
Streak _{jt}		0.00095 (0.00032)***

Notes: The table reports results for equations 5 and 6. The estimations of equations 5 are based on 549'071 data points coming from 6'884 players. The estimations of equations 6 are based on 108'318 data points (where Streak_{jt} > 0) from 6'709 players. Standard errors are robust to within-player dependency. The stars *** indicate that the effect of the variable is statistically significant at the 1% level.

the second regression is Streak_{jt} given that Streak_{jt} > 0.¹⁰

$$\text{Reaction}_{ijt} = \gamma_0 + \gamma_1 \text{Streak}_{jt} + \epsilon_{ijt}, \quad \text{Streak}_{jt} > 0, \quad i = 1, \dots, N, \quad j = 1, \dots, 36, \quad t = 1, \dots, T. \quad (6)$$

Table 2 shows that in both regressions the variables are statistically significant and the signs are in line with the predictions from Rabin and Vayanos (2010). Table 2 shows in the first column that, on average, the tendency to avoid numbers drawn in the previous week is stronger than the tendency to avoid numbers not drawn.¹¹ In addition, the second column shows that the tendency to prefer (i.e. move towards) numbers that have been drawn in the previous week becomes more pronounced as the streak gets longer.

4 Fallacies at the individual level

We now estimate the reaction of individual players to test the prediction by Rabin and Vayanos (2010) that the two fallacies — gambler's and hot hand fallacy — are not simply coexisting but rather that the same players who are prone to gambler's fallacy are also the ones who tend to

¹⁰Alternatives to our approach are running a piecewise linear regression as in Asparouhova et al. (2009) or a quadratic regression. All approaches yield similar results. We prefer our approach because in a piecewise linear or quadratic regression the independent variables refer to different streak lengths (short and long streaks, or linear and quadratic term) and these are correlated by definition. Such correlation potentially biases the econometric estimates. Note that imposing a switch point at 1, as we do, is plausible because a switch point of 1 gives the best fit. Also, a quadratic regression gives a switch point of 1.6.

¹¹Probit regressions give qualitatively similar results.

be prone to the hot hand fallacy as streaks get longer. Our empirical approach in this section involves estimating reactions in number picking to recent draws and streaks for each individual. That is, we run two sets of N separate regressions. Obviously, depending on the confidence level α , some of the estimated coefficients will turn out to be significant by pure chance (in the limit, α times N). In a first step, our empirical strategy therefore involves testing if the distribution of the estimated reactions (particularly, the p -values that correspond to tests of significance) is different from the distribution that would be observed if all players pick numbers randomly. We find that they are different. For example, players are about 4 times more likely to react to numbers drawn in the previous week than predicted by pure chance. This result clearly indicates that biased players are indeed present and that the significance of our estimates is not a statistical fluke. In a second step, we show that players who react by avoiding numbers drawn in the previous week also tend to move their bets to numbers on streaks as streaks get long. For example, among those avoiding numbers drawn last week, 57 percent are also prone to hot hand fallacy, i.e. they move their bets to numbers on streaks as streaks get long.

As in Section 3, we estimate two separate regression models to capture the gambler’s fallacy in the short run and the hot hand fallacy in the long run. And as before, we exclude players who do not react to drawings because non-reaction is difficult to interpret. We estimate the following two regressions for each player i which explain Reaction_{ijt} (defined in eq. 3) as a function of the recent drawing history:

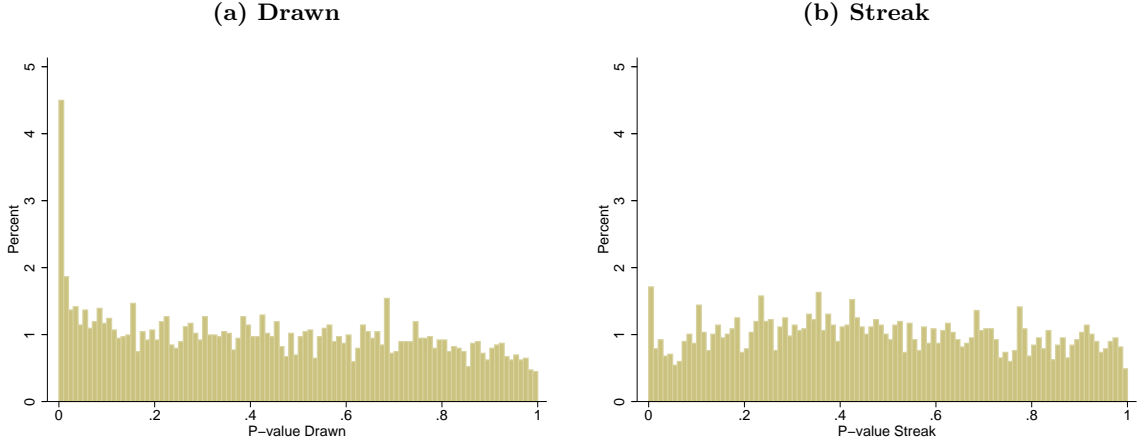
$$\text{Reaction}_{ijt} = \beta_{0j} + \beta_{1i}\text{Drawn}_{jt} + \epsilon_{ijt}, \quad i = 1, \dots, N, \quad j = 1, \dots, 36, \quad t = 1, \dots, T_i, \quad (7)$$

and

$$\text{Reaction}_{ijt} = \gamma_{0j} + \gamma_{1i}\text{Streak}_{jt} + \epsilon_{ijt}, \quad \text{Streak}_{jt} > 0, \quad i = 1, \dots, N, \quad j = 1, \dots, 36, \quad t = 1, \dots, T_i. \quad (8)$$

We define being prone to the gambler’s fallacy as avoiding numbers drawn in the previous week compared to numbers not drawn. Hence, for players prone to the gambler’s fallacy the effect of Drawn_{jt} in eq. 7 is negative. “Hot hand” players increasingly prefer numbers on a streak the longer the streak. For these players the effect of Streak_{jt} in eq. 8 is thus positive. Note that we do not expect any effect of Streak_{jt} in eq. 8 for “pure” gambler’s fallacy players, i.e. those who are prone to the gambler’s fallacy but do not switch to the hot hand fallacy after observing long streaks. To illustrate, consider a “pure” gambler’s fallacy player who strictly avoids numbers drawn in the previous week. Such a player has a negative coefficient on Drawn_{jt} . But since it is possible to avoid a number *only once* even if it is drawn in several consecutive weeks, this player

Figure 3: Distribution of p-values



Notes: The figure shows the distributions of p -values across players of a t -test of $H_0 : \beta_{1i} = 0$ in eq. 7 and of $H_0 : \gamma_{1i} = 0$ in eq. 8. Only players with more than 30 data points in the regressions are included. In panel (a) $j = 1, \dots, 4'024$ and in panel (b) $j = 1, \dots, 997$. The intervals have a size of 1%.

simply *stays away* from these numbers. The effect is that $\text{Reaction}_{ijt} = 0$ when $\text{Streak}_{jt} > 1$ such that eq. 8 cannot be estimated. As explained earlier, data points where $\text{Reaction}_{ijt} = 0$ are omitted from the regression. In addition, eq. 8 cannot be estimated for players who never experience a number on a streak because they do not play frequently (in fact, in 6 out of the 28 weeks studied, no number is on streak).

In the remainder of this section we show in 4.1 that players are significantly more likely to react to numbers drawn in the previous week, and more likely to react to the length of a streak, than predicted by pure randomness. In 4.2 we provide evidence for the gambler's fallacy by showing that players tend to avoid numbers drawn in the previous week. In 4.3 we argue that the majority of players who are prone to gambler's fallacy are also the ones who tend to be prone to the hot hand fallacy as streaks get longer.

4.1 Fallacy or statistical artifact?

Figure 3 shows the distributions of p -values on a t -test that the coefficients $\beta_{1i} = 0$ in eq. 7 (left panel, $n = 4'024$ players) and $\gamma_{1i} = 0$ in eq. 8 (right panel, $n = 997$ players) are different from zero. The figure only includes players with at least 30 data points (see Appendix A.4 for the respective figures when using all players and Appendix A.5 for descriptive statistics of how

data points are distributed across players).¹² If players were picking numbers randomly, the distribution of p -values would be uniform. However, the figure shows that this is not the case. For example, in Figure 3a, the lowest 1% of p -values contains more than 4% of the observations, i.e. more than four times as many as would be observed with random number picking. Also, for 9.2% of the players (412 out of 4475) the effect of Drawn_{jt} in eq. 7 is significantly different from zero at the 5% level.¹³ To test more formally, we find that Kolmogorov-Smirnov tests reject the null hypothesis that the distributions of p -values are uniform ($p = 0.000$ in both cases). Moreover, we find that the average p -value is significantly lower than 0.5 in both cases ($p = 0.000$ in two-tailed one-sample t -tests) while the average would be 0.5 if players picked numbers randomly. We conclude that fallacies are present among the lotto players.

4.2 The gambler's fallacy

To test the players' reaction to numbers drawn in the previous week, we classify the players according to whether they bet more or less on numbers drawn in the previous week ($\hat{\beta}_{1i} > 0$ vs. $\hat{\beta}_{1i} < 0$ in regressions of eq. 7). Table 3 provides the distribution of three sets of players depending on whether they increase or decrease their bets (all, those with at least 30 data points, and those who react significantly at the 5% level). For the distribution over all players we find that more players decrease their bets on numbers drawn in the previous week than players who increase their bets (52% vs. 48%). This distribution is significantly different from the uniform distribution at the 5% level. The finding is qualitatively the same for the group of players with more than 30 observations (52 vs. 48%, significant at the 10% level). If one focuses on those players whose reaction to the previous drawing is statistically significant, then the distribution gets more pronounced: 64% decrease their bets on numbers previously drawn — and thus act in line with the gambler's fallacy — and 36% increase their bets ($p = 0.000$).

4.3 Are players prone to gamblers fallacy also prone to the hot hand fallacy?

To test whether the two types of fallacies are related in the way predicted by Rabin and Vayanos (2010), we classify players according to whether their reaction to Drawn_{jt} in eq. 7 and their

¹²There are 6'288 players for whom eq. 7 can be estimated and 4'612 players for whom eq. 8 can be estimated. For 4'475 players, both regression equations can be estimated.

¹³For 171 players the effect of Drawn_{jt} in eq. 7 is significantly different from zero at the 1% level. This is 3.8% of 4475, which is 3.8 times higher than in the case all players would choose numbers randomly.

Table 3: The gambler’s fallacy

	All players	$n \geq 30$	Significant
Effect of Draw $_{jt}$:			
Negative	3256 (52%)	2092 (52%)	360 (64%)
Positive	3032 (48%)	1932 (48%)	205 (36%)
Total	6288	4024	565
Fisher exact p -value	0.045	0.078	0.000

Notes: The table reports numbers of players where $\hat{\beta}_{1i} < 0$ and $\hat{\beta}_{1i} > 0$ in eq. 7. The numbers are based on all players, players with at least 30 data points, and players for whom $H_0 : \beta_{1i} = 0$ is rejected at the 5% level. The Fisher exact p -values test the observed distributions against the uniform distribution.

reaction to Streak $_{jt}$ in eq. 8 is positive or negative. This allows us to study whether the combination of gambler’s fallacy (i.e. a negative reaction to Draw $_{jt}$ in eq. 7) and hot hand fallacy as streaks get long (i.e. a positive reaction to Streak $_{jt}$ in eq. 8) becomes more frequent than other combinations.

Below we provide supportive evidence for the theory of Rabin and Vayanos (2010) by showing that players prone to the gambler’s fallacy are significantly more likely to be prone to hot hand fallacy than other players. We proceed in two steps. Table 4 uses many but noisy observations, Table 5 few but highly informative ones. More specifically, Table 4 uses all players with more than 30 data points for whom eqs. 7 and 8 can be estimated, Table 5 uses only those who are significantly biased in eq. 7.

We first classify players for whom both regressions (eqs. 7 and 8) can be estimated and who have at least 30 data points in each of the regressions, irrespective of whether the estimated reactions are statistically significant or not ($n = 997$).¹⁴ Table 4 presents a crosstable for these players. The row variable refers to the effect of Draw $_{jt}$ in eq. 7 and the column variable to the effect of Streak $_{jt}$ in eq. 8.

Table 4 shows that the two types of fallacies are significantly related as predicted by Rabin and Vayanos (2010). In particular, among the gambler’s fallacy players (522 players who avoid numbers drawn in the previous week, see row labeled “Negative”), a majority of 57% ($= 298/522$) increasingly bet on numbers on streak as the streak gets longer (different from uniform with $p = 0.022$, Fisher exact test). Conversely, among players prone to the hot hand fallacy (521

¹⁴See Appendix A.6 for a classification of all players.

Table 4: Classification of fallacies ($n \geq 30$)

Effect of Drawn $_{jt}$	Effect of Streak $_{jt}$			
	Negative	Positive	Total	Fisher p -value
Negative	224	298	522	0.022
Positive	252	223	475	0.364
Total	476	521	997	0.325
Fisher p -value	0.399	0.022	0.303	0.002

Notes: The table reports numbers of players where $\hat{\beta}_{1i} < 0$ and $\hat{\beta}_{1i} > 0$ in eq. 7 and $\hat{\gamma}_{1i} < 0$ and $\hat{\gamma}_{1i} > 0$ in eq. 8. It is based on players with at least 30 data points in eqs. 7 and 8 each. The p -values come from Fisher exact tests that compare the observed distributions with the uniform distribution.

players who increasingly bet on numbers on streak the longer the streak, see column labeled “Positive”), a majority of 57% (= 298/521) avoids the numbers drawn in the previous week ($p = 0.022$). Overall, 30% (= 298/997) of all players bet less on numbers drawn in the previous week and increase their bets on numbers on streak, the longer the streak. This is 20% (5 percentage points) more than what would be observed under random selection of lotto numbers.

Next, we classify those players whose estimated reactions in regressions (eq. 7 and eq. 8) are *statistically significant* at the 5% level.¹⁵ The classification in Table 5 follows the same logic as Table 4 and provides the same overall picture. Those players prone to the gambler’s fallacy are also prone to the hot hand fallacy as streaks get longer. It also shows that the “overlap” between both fallacies gets more important as noise is reduced. To illustrate, about three quarters of players (42 out of 58) who are significantly biased in the short run (with respect to Drawn $_{jt}$) and in the long run (with respect to Streak $_{jt}$), are characterized by both fallacies.¹⁶

We should note that although in Table 5 the number of players for whom both effects (Drawn $_{jt}$ in eq. 7 and Streak $_{jt}$ in eq. 8) are significant at the 5% level looks rather low at first sight, it is much higher than the number that would be observed had all players chosen numbers randomly. Indeed, 58 is 1.30% of 4’475 (cf. footnote 12) and more than 5 times higher than the number that would be observed had all players chosen numbers randomly (0.25% of 4’475 is about 11.1).

We conclude that the majority of gambler’s fallacy players switch to the hot hand fallacy as streaks get longer.

¹⁵The same conclusions hold for $\alpha = 0.01$.

¹⁶If we further restrict the analysis to significantly biased players with $n \geq 30$ the support for Rabin and Vayanos (2010) becomes even stronger. Now, 85% of the players are consistent with their theory.

Table 5: Classification of fallacies (significant)

Effect of Drawn _{jt}	Effect of Streak _{jt}			Fisher <i>p</i> -value
	Negative	Positive	Total	
Negative	0	42	42	0.000
Positive	14	2	16	0.054
Total	14	44	58	0.006
Fisher <i>p</i> -value	0.006	0.000	0.021	0.000

Notes: The table reports numbers of players where $\hat{\beta}_{1i} < 0$ and $\hat{\beta}_{1i} > 0$ in eq. 7 and $\hat{\gamma}_{1i} < 0$ and $\hat{\gamma}_{1i} > 0$ in eq. 8. It is based on players i for whom $H_0 : \beta_{1i} = 0$ in eq. 7 and $H_0 : \gamma_{1i} = 0$ in eq. 8 are rejected at the 5% level. The *p*-values come from Fisher exact tests that compare the observed distributions with the uniform distribution.

5 Costly fallacies

We now show that holding the belief that lotto numbers can be predicted is costly. The cost comes in two guises. First, Section 5.1 shows that players who are significantly prone to the gambler’s fallacy tend to buy more tickets than others, and playing the lotto means losing money on average (because the overall payout rate is only 45%). Second, Section 5.2 shows that biased players win smaller amounts given that they win. Note that biased players are just as likely as non-biased players to win, i.e. to guess x numbers correctly. Their problem is that they tend to win in the “wrong weeks” because their individual prize money for guessing x numbers correctly is lower: due to the pari-mutuel nature of lotto they have to share the prize money with many other (biased) players.

5.1 Biased players lose more

We now show that players significantly prone to the gambler’s fallacy buy more tickets than other players. To do so, we regress the average number of tickets bought in Systemlotto on the variables indicating a bias defined in Section 4. In particular, we measure whether players who move away from numbers buy more tickets than those who move toward numbers, by including “Negative” in the regression, which is a binary variable equal to one for player i if $\hat{\beta}_{1i} < 0$ in eq. 7 and equal to zero if $\hat{\beta}_{1i} > 0$. We also include the variable “Significant Reaction”, which indicates whether $H_0 : \beta_{1i} = 0$ in eq. 7 is rejected at the 5% level for player i or not (0/1), and the interaction between Negative and Significant Reaction. Including the interaction allows us to test whether players prone to the gambler’s fallacy — those who react negatively and significantly

Table 6: Regression of number of tickets bought

Dep. var.: Avg. # tickets bought	Est. (s.e.)
Intercept	11.82 (1.26) ^{***}
Negative	0.81 (0.51)
Significant Reaction	-2.58 (1.38) [*]
Negative x Significant Reaction	4.74 (1.75) ^{***}
Male	4.08 (0.69) ^{***}
Age	-0.13 (0.02) ^{***}
Avg. # numbers picked	1.27 (0.02) ^{***}
<i>N</i>	6287

Notes: The table reports results from robust regressions where the dependent variable is the average number of tickets bought by player i through the Systemlotto device. Negative = 1 for player i if $\hat{\beta}_{1i} < 0$ in eq. 7 and 0 otherwise. Significant Reaction = 1 if $H_0 : \beta_{1i} = 0$ in eq. 7 is rejected at the 5% level for player i and 0 otherwise. The stars ^{***}, ^{**} or ^{*} indicate that the effect of the variable is statistically significant at the 1%, 5% or 10% level, respectively.

to the outcome of the previous week’s drawing — buy more or less tickets than other players.¹⁷ We also include a dummy for gender, age, and the average number of lotto numbers picked as control variables.¹⁸ Including the latter is important to avoid upwardly biased estimates for two reasons. First, because players who pick more lotto numbers are more likely to move away from a number than players who pick fewer lotto numbers — simply because there are more numbers to potentially move away from — and thus also more likely to have $\hat{\beta}_{1i} < 0$ in eq. 7. Second, because the number of lotto numbers picked is positively correlated with the number of tickets bought. Leaving out this control variable thus induces upward bias on the regression coefficient of Negative.

¹⁷Out of the 6’287 players for whom eq. 7 can be estimated, for 565 players i $H_0 : \beta_{1i} = 0$ is rejected at the 5% level (see also Table 3).

¹⁸The individual characteristics we have are gender, age and postal code. We experimented with regressions that include variables related to the estimation results of eq. 8, where the effect of Streak_{jt} on Reaction_{ijt} is estimated. For example, we included whether or not $\hat{\gamma}_{1i} > 0$ as a dummy regressor in a regression where the average number of tickets bought is the dependent variable. However, none of the variables related to eq. 8 turned out to be significantly correlated with the average number of tickets bought by Systemlotto players.

Table 6 shows that the interaction between Negative and Significant Reaction is statistically significant, whereas Negative is not. This implies that players who are (significantly) prone to gambler’s fallacy, i.e. for whom $\hat{\beta}_{1i} < 0$ and $H_0 : \beta_{1i} = 0$ is rejected at the 5% level, buy significantly more tickets than players who react significantly positive to the previous week’s drawing ($\hat{\beta}_{1i} > 0$), or who do not react significantly at all (for whom $H_0 : \beta_{1i} = 0$ is not rejected at the 5% level). In particular, gambler’s fallacy players buy almost 3 tickets more ($0.81 - 2.58 + 4.74$) or spend about 1.2 EUR more in an average week (this is significantly different from zero with $p = 0.006$). In contrast, players whose move toward numbers drawn in the previous week is significant ($\hat{\beta}_{1i} > 0$ and $H_0 : \beta_{1i} = 0$ rejected at the 5% level), buy about 2.5 tickets less than players whose move toward numbers drawn in the previous week is not significant ($\hat{\beta}_{1i} > 0$ and $H_0 : \beta_{1i} = 0$ cannot be rejected at the 5% level). However, this effect is only weakly significant.

Table 6 further shows that male players buy more tickets than female players, while older players buy significantly fewer tickets. For example, men buy on average 4 tickets (i.e. spend 1.6 EUR) more per week than women. The regression also shows that, as expected, the relation between the average number of numbers picked and the average number of tickets bought is positive and significant.

One hypothesis is that overconfidence might drive our result that players who react significantly in line with the gambler’s fallacy buy significantly more tickets. This resonates well with the fact that the gambler’s fallacy results are driven by male players, and males have been found to be particularly prone to overconfidence (see, e.g., Niederle and Vesterlund, 2007). We find that the average female player does not react significantly to recent drawings.¹⁹

5.2 Biased players win less

Given that it is not possible for lotto players to increase the probability of winning per ticket purchased, the best they can do is to maximize the amount won if they happen to win. Winnings are maximized by avoiding popular numbers and combinations. The reason is that because of the pari-mutuel structure of lotto, picking popular numbers entails the risk of having to share the winning prize with many other players. Thus, being biased is costly if the bias is common.

¹⁹Suetens and Tyran (2011) show that, in the aggregate, men are prone to the gambler’s fallacy whereas women are not. However, the switching result obtained in the current paper (from gambler’s fallacy to hot hand fallacy) is also obtained for female players who are significantly prone to the gambler’s fallacy with respect to the previous week’s drawing.

Picking the same numbers can be due to a bias as we study in this paper, or due to using visual patterns to pick numbers. For example, Simon (1999) shows that picking the least popular combinations yields a gain that is more than four times higher than picking the most popular combinations in UK lotto.

We now show that biased players win lower amounts than players who pick numbers randomly, given that they happen to guess x numbers right. In particular, we compare the prize that biased players receive when getting four or five lotto numbers correct with the prize random-picking players would receive in this case. Note that by comparing payoffs conditional on winning, we effectively isolate the effect of picking popular numbers from the effect of pure luck, provided that the number of observations is sufficiently large. Yet, winning the jackpot is very unlikely (about 1 in 8 million) and the chances are rather low even for 6 correct numbers. As a consequence, there are not many observations in these prize categories (24 in our dataset) which implies that the effect of pure luck on the payoff of lotto players is quite high in these categories. Hence, we refrain from investigating the difference for 6 and 7 correct numbers.²⁰

We establish the result that being biased is costly for the subset of players for whom we have information on which combinations were actually chosen by the lotto agency for these players (i.e. players using system M, see Appendix A.2). As a benchmark, we calculate the expected prize of players picking numbers randomly by means of Monte Carlo simulations. In the simulations we use the fact that a random player has the same probability of winning in all weeks. Our result suggests that biased players tend to win in the “wrong” weeks, where many winners share the prize.

Table 7 shows that biased players indeed realize lower winning prizes than they would if they would pick numbers randomly.²¹ The first column shows the average winning of a random player (DKK 42.6 for 4 correct and DKK 149.6 for 5 correct). The second column shows that players with $\hat{\beta}_{1i} < 0$ (whether significant or not), i.e. those that tend to avoid numbers drawn in the previous week win DKK 42.5 and DKK 148.2 in the two prize categories, respectively. These averages are significantly lower (at the 1% level or better) in both cases. The same conclusion holds for players with $\hat{\beta}_{1i} > 0$ (whether significant or not), but the differences are even more

²⁰The prizes in the lotto vary quite considerably over time. For example, the prize for guessing 6 numbers correctly was between 18'000 DKK and 113'000 DKK with a mean of 66'000 DKK in our dataset.

²¹Note that the numbers of observations reported in Table 7 do not refer to numbers of players, but to numbers of winnings from predicting four or five times correctly.

Table 7: Comparison of average prize conditional on winning

Prize	Simulated		Observed		
	(Random)	Neg. Reaction	Pos. Reaction		
		All	Signif.	All	Signif.
Four numbers correct					
Nr. of observations	-	9926	1276	7475	526
Average prize, DKK	42.57	42.49	42.32	42.42	41.95
Pr(Same or higher than random player)	-	< 0.01	< 0.01	< 0.01	< 0.01
Five numbers correct					
Nr. of observations	-	735	-	433	-
Average prize, DKK	149.6	148.2	-	143.8	-
Pr(Same or higher than random player)	-	< 0.01	-	< 0.01	-

Notes: The players are categorized based on the estimation of eq. 7. “Neg. Reaction” refers to players for whom $\hat{\beta}_{1i} < 0$, and “Pos. Reaction” to those for whom $\hat{\beta}_{1i} > 0$. “Signif.” refers to those players for whom $H_0 : \beta_{1i} = 0$ in eq. 7 is rejected at the 5% level. The random numbers used for the simulations of randomly picking players are based on atmospherical noise. They are acquired from the website random.org. P-values are derived using Monte Carlo simulations. The simulations are based on sample sizes of 9926 and 735 for the four-correct and five-correct prizes, respectively. The simulations are run 100 times and if the average prize of the Systemlotto players is strictly smaller than the smallest of the 100 simulated averages, the p -value is < 0.01.

pronounced (DKK 42.4 and 143.8, respectively, see fourth column). The result that the latter do even worse is not surprising because a coordinated move toward the seven numbers drawn in the previous week (when $\hat{\beta}_{1i} > 0$) makes it more likely to share the prize when winning than a coordinated move away from the seven numbers drawn in the previous week (when $\hat{\beta}_{1i} < 0$, which is in fact a move toward the other twenty-nine numbers not drawn in the previous week). The third and fourth columns show that the effects are more pronounced for having 4 numbers correct when considering only observations of significantly biased players (i.e. players for whom $H_0 : \beta_{1i} = 0$ in eq. 7 is rejected at the 5% level).²²

While the losses compared to random play are highly statistically significant, the absolute magnitudes are not that large. The winnings for having four numbers correct are between 0.6% and 1.5% lower for the significantly biased players compared to random players. The effect is larger, up to 3.9 percent, for winning with five correct numbers. We speculate that the effects grow exponentially for the two highest prize categories simply because sharing a large prize with

²²The number of observations for the case of winning with five correct numbers is too low to run meaningful regressions.

an additional person is more costly in absolute terms if the prize is large. Recall that the average prize money in our sample is about 1500 times larger for having 6 correct numbers than 4 correct (DKK 66'000 vs. DKK 43) We therefore believe that the calculated losses are a conservative estimate of the cost.

6 Conclusion

Given that lotto drawings are truly random, it seems absurd to believe that anyone can predict next week's numbers. Yet, our data suggests that such a belief is persistent in the aggregate for the group of lotto players we study, and curiously enough, the lotto agency itself describes the lotto as: "a number game which is about predicting the correct numbers drawn" (translated from danskespil.dk, see "rules of the game"). In line with recent behavioral theory, we find that players tend to avoid numbers that have been drawn in the previous week but also tend to favor numbers that are "on streak", i.e. have been drawn in several consecutive weeks. Our unique panel data set allows us to track individual players over time and we find that the "gambler's fallacy" and the "hot hand fallacy" are systematically related as suggested by Rabin and Vayanos (2010).

Our aggregate results are remarkable given the parimutuel nature of lotto. If anything, lotto is not about predicting the numbers drawn — because all numbers are equally likely — but about predicting numbers picked by others — because avoiding numbers picked by others increases the prize amount given that a player happens to win. In this sense having a bias is costly: we show that biased players win lower amounts than players who pick numbers randomly, given that they happen to win a prize.

Although our data come from a state lottery that has a particularly transparent stochastic nature, the same biases may also appear in contexts when the random process is less transparent. Think, for example, of financial markets, where aggregate outcomes are influenced by random shocks *and* human hands. For example, Kumar (2009) shows that it is the same type of people who participate in state lotteries and who buy lottery-type stocks (such as low-priced stocks with high idiosyncratic volatility and skewness).

References

- Asparouhova, E., Hertzfel, M., and Lemmon, M. (2009). Inference from streaks in random outcomes: experimental evidence on beliefs in regime shifting and the law of small numbers. *Management Science*, 55:1766–1782.
- Ayton, P. and Fischer, I. (2004). The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness? *Memory and Cognition*, 32:1369–1378.
- Bar-Hillel, M. and Wagenaar, W. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12:428–454.
- Burns, B. and Corpus, B. (2004). Randomness and inductions from streaks: “gambler’s fallacy” versus “hot hand”. *Psychonomic Bulletin and Review*, 11:179–184.
- Camerer, C. (1989). Does the basketball market believe in the hot hand? *American Economic Review*, 79:1257–1261.
- Clotfelder, C. and Cook, P. (1993). The “gambler’s fallacy” in lottery play. *Management Science*, 39:1521–1525.
- Croson, R. and Sundali, J. (2005). The gambler’s fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty*, 30:195–209.
- Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17:295–314.
- Kumar, A. (2009). Who gambles in the stock market? *Journal of Finance*, 64:1889–1933.
- Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122:1067–1101.
- Offerman, T. and Sonnemans, J. (2004). What’s causing overreaction? An experimental investigation of recency and the hot-hand effect. *Scandinavian Journal of Economics*, 106:533–553.
- Rabin, M. (2002). Inference by believers in the law of small numbers. *Quarterly Journal of Economics*, 117:775–816.
- Rabin, M. and Vayanos, D. (2010). The gambler’s and hot-hand fallacies: Theory and applications. *Review of Economic Studies*, 77:730–778.
- Rapoport, A. and Budescu, D. (1997). Randomization in individual choice behavior. *Psychological Review*, 104:603–617.
- Simon, J. (1999). An analysis of the distribution of combinations chosen by UK national lottery players. *Journal of Risk and Uncertainty*, 17:243–276.
- Suetens, S. and Tyran, J.-R. (2011). The gambler’s fallacy and gender. *Journal of Economic*

Behavior & Organization, Forthcoming.

Sundali, J. and Croson, R. (2006). Biases in casino betting: The hot hand and the gambler's fallacy. *Judgment and Decision Making*, 1:1–12.

Terrell, D. (1994). A test of the gambler's fallacy: Evidence from pari-mutuel games. *Journal of Risk and Uncertainty*, 8:309–317.

Tversky, A. and Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76:105–110.

A Appendix

A.1 Application

This section applies the model Rabin and Vayanos (2010) to the context of lottery play in Denmark. We assume that an individual observes a sequence of lotto drawings that depends on the probability that a number is drawn and an *i.i.d.* normal shock. The signal s_t in week $t = 1, 2, \dots$ is

$$s_t = \theta_t + \epsilon_t, \quad (9)$$

where θ_t is the probability that a number is drawn and ϵ_t the normal shock with mean zero and variance $\sigma_\epsilon^2 > 0$. The probability that a number is drawn is assumed to evolve according to the following auto-regressive process:

$$\theta_t = \mu + \rho(\theta_{t-1} - \mu) + \eta_t, \quad (10)$$

where μ is the long-run mean, $\rho \in [0, 1]$ the persistence parameter, and η_t an *i.i.d.* normal shock with mean zero, variance σ_η^2 , and independent of ϵ_t . Given that the drawing machines and the sets of lotto balls used are replaced from time to time, the parameter ρ can be seen as measuring the persistence of the drawing machines and lotto balls: a ρ close to one implies a high probability of using the same machine and the same set of balls and a ρ close to zero implies a high switching or replacing probability.²³ The variance σ_η^2 measures the variability in drawing outcomes between different drawing machines and sets of balls: a high σ_η^2 would imply that different machines or sets of balls generate different drawing probabilities. Since the lotto agency does everything it can to generate fair drawings, we assume that $\sigma_\eta^2 = 0$. Therefore, θ_t is constant and equal to the probability that a number is drawn.

The gambler's fallacy is modeled as the mistaken belief that ϵ_t is not *i.i.d.* but exhibits reversals in the following sense:

$$\epsilon_t = \omega_t - \alpha\rho \sum_{k=0}^{\infty} (\delta\rho)^k \epsilon_{t-1-k}, \quad (11)$$

where ω_t is an *i.i.d.* normal shock with mean zero and variance σ_ω^2 , and $\alpha, \delta \in [0, 1)$. Whether a lotto player who falls prey to the gambler's fallacy starts to develop hot hand beliefs after

²³The Danish lotto agency operates with three sets of lotto balls and two identical drawing machines. The machines are switched every six months and the balls are replaced when they have been used fifty-two times. The public is not informed about the timing of switching the machine nor replacing the balls.

observing long streaks depends on whether he is certain about the constancy of the probability that a number is drawn (equal to $7/36$).

Consider first the case where a lotto player is absolutely certain about the drawing probability: he believes that the probability that a number is drawn exhibits no variability ($\tilde{\sigma}_\eta^2 = 0$) and is equal to $7/36$. In this case $\theta_t = \mu$ and the player observes signals according to $s_t = \mu + \epsilon_t$, where ϵ_t refers to the gambler's fallacy beliefs modeled in eq. 11.

When a lotto player is uncertain about the drawing probability, for example, because he does not fully trust the drawing mechanism and believes in variability between different sets of lotto balls, it can be shown that the player will develop a belief in the hot hand. In particular, Proposition 5 in Rabin and Vayanos (2010) shows that under mild assumptions²⁴ an uncertain player will develop the wrong belief that the drawing probability varies over time: $\tilde{\sigma}_\eta^2 > 0$. The intuition is that in order to explain the absence of reversals, an uncertain lotto player will overestimate the turnover in the drawing machines or lotto balls (i.e. underestimate ρ) and overestimate the variability in drawing outcomes between these machines or sets of balls. The consequence is that an uncertain lotto player who believes in the law of small numbers will expect reversals after short streaks and continuation after long streaks (see Proposition 6 in Rabin and Vayanos, 2010).

²⁴The player should be confident that there is some persistence of the drawing mechanism: $\tilde{\rho} \in [\underline{\rho}, 1]$ with $\underline{\rho} > 0$. In other words, the player should believe that the drawing machine and the lotto balls are not replaced every week.

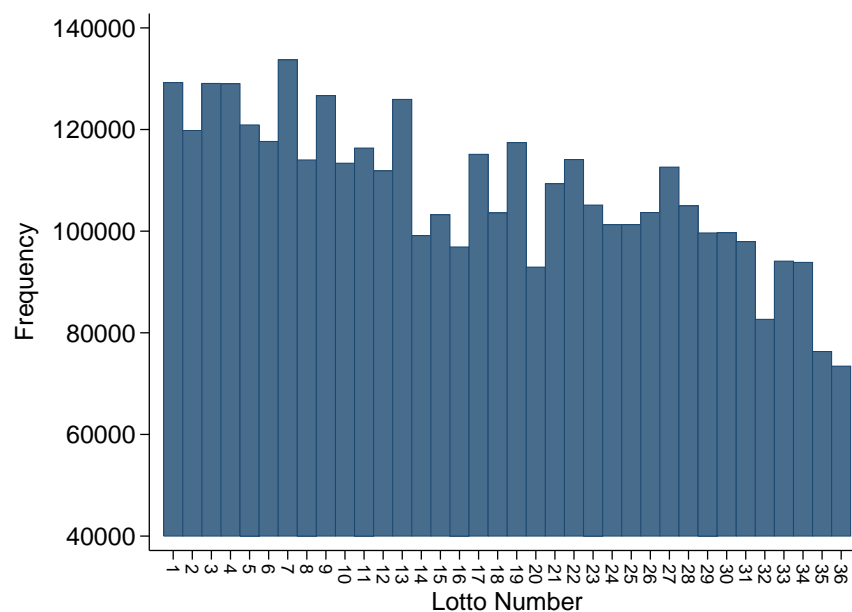
A.2 Overview of the different systems in Systemlotto

Option	Type of system	# chosen numbers in a set	# tickets/combinations generated
1	M	8	8
2	M	9	36
3	M	10	120
4	M	11	330
5	M	12	792
6	R	10	8
7	R	10	30
8	R	11	20
9	R	11	34
10	R	12	12
11	R	12	24
12	R	12	48
13	R	13	18
14	R	13	66
15	R	14	48
16	R	14	132
17	R	15	24
18	R	15	69
19	R	16	32
20	R	16	109
21	R	16	240
22	R	17	272
23	R	18	82
24	R	19	338
25	R	20	450
26	R	20	1040
27	R	21	198
28	R	23	345
29	R	24	455
30	R	25	600
31	C	17	17
32	C	18	33
33	C	19	52
34	C	20	20
35	C	20	80
36	C	22	60
37	C	24	24
38	C	24	120
39	C	25	100
40	C	25	200
41	C	28	194
42	C	30	268
43	C	31	155

Notes: Under system “M” all potential combinations of the chosen numbers in a set are generated. Systems “R” and “C” are both reduced systems, in which only some of the potential combinations are generated. Systems “C” are more heavily reduced than systems “R”. That is, a smaller share of the potential combinations are generated in systems “C” than in systems “R”.

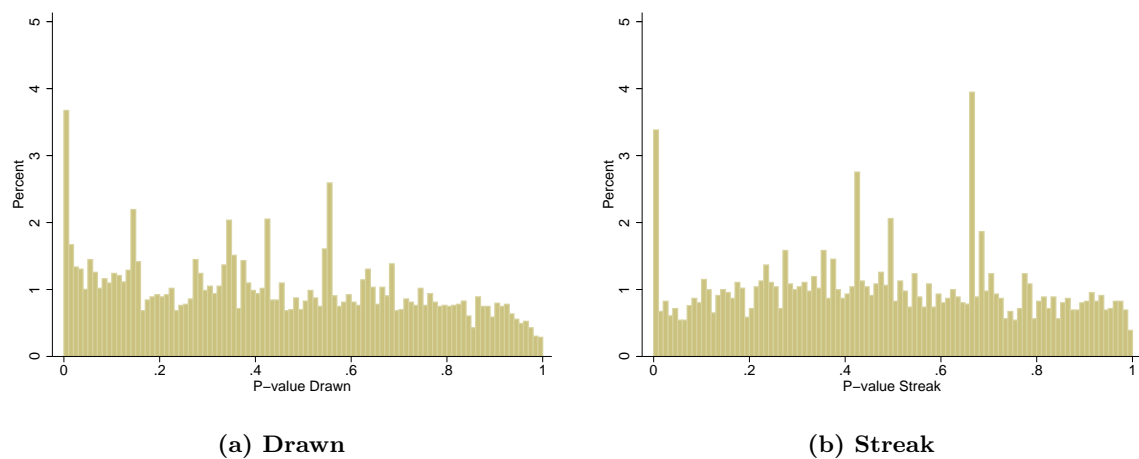
A.3 Popularity of lotto numbers

Figure A1: Popularity of lotto numbers



Notes: The chart shows for each lotto number the total number of times the number is picked in the period of study (weeks 25 to 52 of 2005) using the Systemlotto selection device over the Internet. Note that a player is not counted more than once in a week.

A.4 Distribution of p -values based on all players

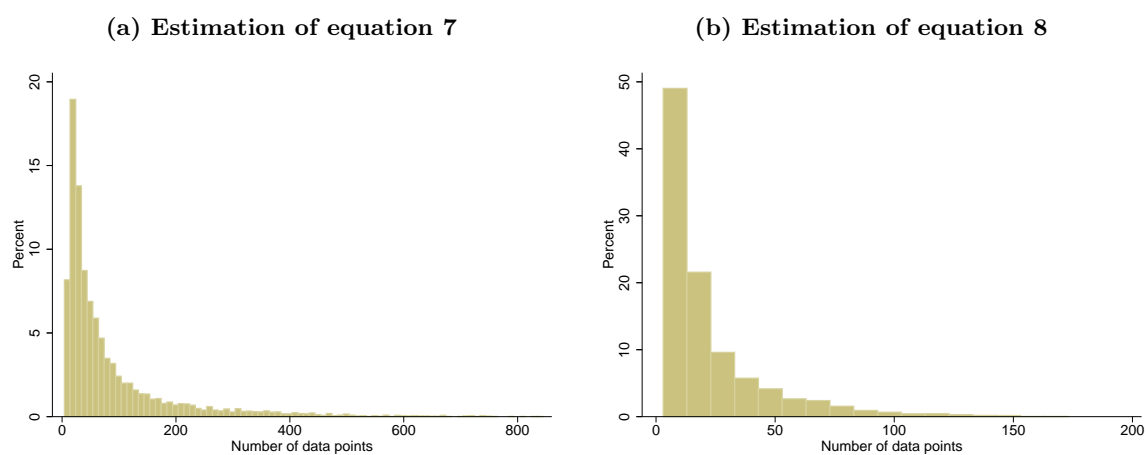


(a) Drawn **(b) Streak**

Notes: The figure shows the distributions of p -values across all players of a t -test of $H_0 : \beta_{1i} = 0$ in eq. 7 and of $H_0 : \gamma_{1i} = 0$ in eq. 8. In panel (a) $j = 1, \dots, 6288$, and in panel (b) $j = 1, \dots, 4612$. The intervals have a size of 1%.

A.5 Distributions of number of data points

Figure A2: Distribution of number of datapoints



Notes: The figure shows the distributions of the number of data points across players in regressions of eqs. 7 and 8. The intervals have a width of 10. The minimum number of data points is 4 in regressions of eq. 7 and 3 in regressions of eq. 8.

A.6 Crosstable based on all players

Effect of Draw n_{jt}	Effect of Streak $_{jt}$			Fisher p -value
	Negative	Positive	Total	
Negative	1144 49% 52% 26%	1181 51% 57% 26%	2325 100% 52% 52%	0.598
Positive	1057 53% 48% 24%	1093 47% 43% 24%	2150 100% 48% 48%	0.604
Total	2201 48% 100% 48%	2274 52% 100% 52%	4475 100% 100% 100%	0.447
Fisher p -value	0.194	0.202	0.066	1.000

Notes: The tables report numbers of players where $\hat{\beta}_{1i} < 0$ and $\hat{\beta}_{1i} > 0$ in eq. 7 and $\hat{\gamma}_{1i} < 0$ and $\hat{\gamma}_{1i} > 0$ in eq. 8. The p -values come from Fisher exact tests that compare the observed distributions with the uniform distribution.