

INDIVIDUAL REACTION TO PAST PERFORMANCE SEQUENCES: EVIDENCE FROM A REAL MARKETPLACE*

Angie Andrikogiannopoulou[†]
London School of Economics

Filippos Papakonstantinou[‡]
Imperial College London

Abstract

We use novel data on individual activity in a sports betting market, to study the effect of past performance sequences on individual behavior in a real market. The idiosyncratic nature of risk in this market and the revelation of assets' true terminal values enables us to disentangle whether behavior is caused by sentiment or by superior information about market mispricings, and to cleanly test two prominent theories of momentum and reversals — the regime-shifting model of Barberis, Shleifer, and Vishny (1998) and the gambler's/hot-hand fallacy model of Rabin (2002). Furthermore, our long panel enables us to study the prevalence across individuals of each type of behavior. We find that i) three quarters of individuals exhibit trend-chasing behavior; ii) seven times as many individuals exhibit behavior consistent with Barberis, Shleifer, and Vishny (1998) as exhibit behavior consistent with Rabin (2002); and iii) no individuals earn superior returns from momentum trading.

Keywords: Momentum, Individual Decision-making, Heterogeneity, Behavioral Biases, Sports Betting

JEL Classification: D12, D81, G02, G11, G14

*For helpful comments, the authors thank Nick Barberis, Tony Berrada, John Gandar, Zoran Ivkovic, Matti Keloharju, Tobias Moskowitz, Hyun Shin, Wei Xiong, seminar participants at the 2014 Research in Behavioral Finance conference in Rotterdam and the 2014 Swiss Finance Institute conference in Gerzensee, and the anonymous reviewers and the department editor, Lauren Cohen. The authors remain responsible for all errors, omissions, and typos. The authors also gratefully acknowledge financial support from the Swiss Finance Institute.

[†]London School of Economics, Department of Finance, Houghton Street, London WC2A 2AE, UK, e-mail: a.andrikogiannopoulou@lse.ac.uk.

[‡]Imperial College Business School, Tanaka Building, South Kensington Campus, London SW7 2AZ, UK, e-mail: fpapakon@imperial.ac.uk.

A large literature in finance has uncovered various predictable patterns in stock returns. Two of the most robust patterns are short-run momentum and long-run reversals: Assets with good recent performance tend to continue overperforming in the short term, and assets with good performance over a long period tend to underperform subsequently. Rational explanations of these patterns propose that they reflect patterns in systematic risk and are therefore consistent with market efficiency, while behavioral theories propose that they reflect mispricings which arise from individual biases and persist due to limits to arbitrage. However, direct tests of alternative theories using stock market data are very difficult, because tests of market efficiency rely on a specific asset pricing model and therefore face a joint hypothesis problem. Instead, it may be more fruitful to use clean empirical settings to test for the individual behavior predicted by the proposed theories. The literature has so far focused on testing for these behaviors using data from experiments, but these tests also face objections, because they study behavior in artificial settings.

In this paper, we test two theories that have been proposed to explain short-run momentum and long-run reversals in returns — the regime-shifting model of Barberis, Shleifer and Vishny (1998) and the gambler's/hot-hand fallacy model of Rabin (2002) — by utilizing a novel panel data set on individual activity in a sports betting market. Sports betting markets provide a useful middle ground between the stock market and controlled laboratories as an empirical setting. Contrary to the stock market, they feature assets (i.e., wagers) that bear no systematic risk, are short-lived, and have an observable terminal value which is exogenously determined and revealed by match outcomes. The fact that assets' risk is idiosyncratic implies that an asset pricing model is unnecessary, therefore tests of market efficiency and/or superior individual performance do not suffer from the joint hypothesis problem (Fama, 1970). The fact that assets' terminal values are observable and exogenous implies that there is no mispricing at termination, therefore performance can be measured accurately and cleanly. These features make it possible to disentangle whether individual behavior is driven by sentiment or by superior information about potential mispricings. Contrary to experiments, the sports betting market enables us to study individual behavior that arises naturally in a real-world market setting in which significant amounts of money are at stake. Finally, in contrast to both the

stock market and experimental settings, individuals in the sports betting market typically make over time a large number of choices with serially uncorrelated payoffs. This allows us to go beyond analyzing average behavior to analyzing the prevalence of these behaviors across individuals.

Barberis, Shleifer and Vishny (1998) develop a model in which a company's earnings follow a random walk process, but investors believe that at any point in time either a 'trending' or a 'mean-reverting' regime governs the earnings process. The psychological motivation for the belief in the trending regime is the representativeness bias (Tversky and Kahneman, 1974), according to which people expect small samples to be representative of the underlying distribution, and the motivation for the belief in the mean-reverting regime is the conservatism bias (Edwards, 1968), according to which people tend to underweight new evidence relative to prior beliefs. In this model, investors systematically rely on patterns in past outcomes to assess the likelihood of which regime governs the earnings process, and they expect that the probability of a streak continuing increases monotonically with the streak's length. Rabin (2002) also develops a model based on the representativeness bias, in which outcomes are generated by draws from an urn with replacement, but the decision-maker erroneously believes they are generated with replacement only every K draws, for some $K > 1$. When the decision-maker does not know, hence must infer, the urn's rate, he exhibits the gambler's fallacy after short streaks and the hot-hand fallacy after long streaks.¹ The interaction of these two effects implies a non-monotonic (first decreasing, then increasing) relation between streak length and the predicted probability of streak continuation. Thus, while both models predict that individuals underreact to news in the short-run and overreact in the long-run, they yield different predictions regarding the relation (monotonic or not) between streak length and individual belief in streak continuation.

Our analysis is based on the complete trading histories of a random sample of 500 individuals who placed wagers at an online sportsbook over a period of five years (2005–2010). Furthermore,

¹The gambler's fallacy refers to the mistaken belief that random sequences exhibit systematic reversals, hence future outcomes should balance the historical sequence toward the presumed rate. The hot-hand fallacy refers to the mistaken belief that random sequences exhibit systematic persistence, hence after observing a streak of similar outcomes, beliefs about the rate are biased toward the streak's outcome. Rabin and Vayanos (2010) adapt the gambler's/hot-hand fallacy model of Rabin (2002) to a setting with normally distributed outcomes and derive similar implications. Here, we test the Rabin (2002) model, which is more directly relevant to our setting in which outcomes are binary (win or loss).

we use data on teams' past performance to construct measures of momentum analogous to the ones used in stock market studies. We examine whether individuals' trading strategies are affected by momentum in team performance, and in particular, whether individuals' reaction to the length of performance streaks is monotonic as implied by Barberis, Shleifer and Vishny (1998) or non-monotonic as implied by Rabin (2002). Furthermore, the idiosyncratic and exogenous terminal value of assets in the sports betting market allows us to distinguish whether individual behavior is driven by an inherent cognitive bias, e.g., a misguided belief in momentum, or by superior information related to market mispricings. In our first set of tests, we use all wagers from all individuals to study the average behavior across individuals. In our second set of tests, we take advantage of our relatively long panel which contains, on average, 204 bets per individual, to repeat our analyses at the individual level. This enables us to uncover differences in trading behavior across individuals, that could not be uncovered in an analysis of average behavior. In particular, we use the wagers from each individual separately and we employ a multiple-testing methodology to obtain estimates of the proportions of individuals i) who exhibit trend-chasing or contrarian behavior, ii) whose behavior is more consistent with the model of Barberis, Shleifer and Vishny (1998) or the model of Rabin (2002), and iii) who possess an informational advantage about the teams toward which they exhibit a bias.

We find that individuals exhibit a pronounced preference (aversion) toward teams on long winning (losing) streaks. Comparing the composition of individual betting portfolios relative to price-matched 'market' portfolios that are contemporaneously available, we find that 16% (3%) of the former is allocated to teams on long winning (losing) streaks, while these teams make up 8% (7%) of the latter. Furthermore, conditional on the matches selected by each individual, the average individual is 36% more likely to back a team on a long winning streak relative to a team on a long losing streak. Furthermore, we find that individual reaction to streak length is monotonic, as predicted by Barberis, Shleifer and Vishny (1998): the ratio of individual-to-market portfolio weights and the probability of backing a team is the highest for teams that are on long winning streaks, and it decreases progressively for teams that are on short winning, short losing, and long losing streaks. Importantly, the average return individuals realize from wagering in favor of teams that

are on winning streaks is not different from the return they realize from wagering against these teams. This finding indicates that individual biased behavior is driven by sentiment rather than by an informational advantage, and that betting prices are, on average, efficient with respect to teams' performance streaks.²

To examine the prevalence of these behaviors across market participants, we repeat all aforementioned analyses at the individual-level, and we employ the powerful False Discovery Rate (FDR) methodology to correct for the problem of false discoveries in multiple testing. We find that an overwhelming majority of individuals (83%) are consistently affected by team performance streaks in their decisions: 78% (5%) exhibit momentum (contrarian) trading behavior as they systematically overweight teams on long winning (losing) streaks in their portfolios relative to the market and, conditional on their selected wagers, they are more likely to back a team that is on a long winning (losing) streak. Furthermore, we find that the majority of individuals (55%) show a monotonic response to streak length as implied by Barberis, Shleifer and Vishny (1998), while only a small minority (8%) exhibit the non-monotonic reaction that is implied by Rabin (2002). Since approximately seven times as many individuals exhibit a monotonic response to streak length as exhibit a non-monotonic response, this individual-level real-market evidence is strongly in favor of the model proposed by Barberis, Shleifer and Vishny (1998). Finally, we find that the proportion of individuals who generate significant excess returns from backing teams that are on long winning streaks is negligible. This implies that, not just for the average, but for all individuals who are affected by team performance streaks, this behavior is driven by a mistaken belief that historical performance predicts future performance rather than by superior information.

Our paper relates to the large finance literature that has tried to understand what motivates the trading behavior of individual investors. Using data from the stock market, a number of studies have

²Even though a thorough study of sportsbook pricing is beyond the scope of this paper, it seems that, contrary to the market maker in financial markets, the bookmaker in our sports betting market finds it optimal to set prices close to the efficient ones, allowing the book to become slightly unbalanced. This strategy saves on the costs associated with perfectly balancing the book on every single match and could lead to greater profits in the long-run as the bookmaker earns his commission on losing bets (see Paul and Weinbach, 2012). This feature of the sports betting market is attractive for our analysis since it implies that the stochastic process that underlies the observed sequences of performance outcomes is a random walk, as assumed in the models of Barberis, Shleifer and Vishny (1998) and Rabin (2002).

documented that investors systematically pay attention to stocks' past performance and follow momentum or contrarian trading strategies (e.g., Odean, 1998; Grinblatt and Keloharju, 2001; Barber, Odean and Zhu, 2009; Loh and Warachka, 2012), but the implications of this evidence are unclear since these trading strategies could be rational or not, depending on the true asset pricing model. Several experimental (e.g., Tversky and Kahneman, 1971; Burns and Corpus, 2004) and field (Clotfelter and Cook, 1993; Jorgensen, Suetens and Tyran, 2011) studies have demonstrated more clearly that individuals often expect outcomes of random sequences to exhibit systematic reversals or excessive persistence, but only a few studies have tested for the specific behaviors hypothesized in the models of Barberis, Shleifer and Vishny (1998) and Rabin (2002), and have yielded mixed evidence. Bloomfield and Hales (2002) and Asparouhova, Hertz and Lemmon (2009) test these theories in an experimental setting in which subjects are shown sequences of price movements generated by a random walk and are asked to assess the probability that the next movement will be upwards; while Bloomfield and Hales (2002) find evidence consistent with the model of Barberis, Shleifer and Vishny (1998), Asparouhova, Hertz and Lemmon (2009) find evidence consistent with the model of Rabin (2002). Contrary to us, these studies use experimental data rather than data from a real-world market setting in which behavior arises naturally. Using real-world data from the sports betting market, Avery and Chevalier (1999), Durham, Hertz and Martin (2005), and Moskowitz (2015) find that price movements over the betting period exhibit momentum patterns. Durham, Hertz and Martin (2005), in particular, find that the marginal bettor has a non-monotonic response to the length of teams' performance streaks, which is inconsistent with the model of Barberis, Shleifer and Vishny (1998). Contrary to us, these studies use market price data, which aggregate in an opaque manner the behavior of diverse market participants, including the market maker. Overall, the important difference between our study and the existing studies is that we use individual-level data from a marketplace which possesses experimental-like features; this enables us to perform clean tests of individual behavior from real-world decisions.

While extrapolating results from one domain to another should generally be undertaken with caution, it is likely that the behavioral biases we document are inherent and therefore they carry

over to other settings, perhaps particularly so to settings that are closer to the sports betting market. Indeed, the sports betting market shares several similarities with traditional financial markets, both in terms of operation and in terms of the mentality of its participants. First, it is populated by a large number of participants with different levels of sophistication (e.g., noise traders and arbitrageurs) who risk significant amounts of their money on the uncertain outcome of future events, while information about sports teams and events is widely available in the media.³ Second, recent survey evidence suggests that the key driver of gambling behavior is the monetary factor (“to make money” or to have “the chance of winning big money”), even though some market participants are also driven by non-pecuniary motives like team loyalty, fandom, or social factors.⁴ Conversely, it has been shown that many individual investors in the stock market also view trading as a gambling activity (Grinblatt and Keloharju, 2009; Dorn and Sengmueller, 2009; Gao Bakshi and Lin, 2015), prefer stocks with lottery characteristics (Kumar, 2009), and are largely motivated by loyalty in their decisions (Keloharju, Knapfer and Linnainmaa, 2012).⁵ Indeed, the literature has long identified the attractive features and the useful insights that sports betting markets can provide for traditional financial markets (e.g., Gandar et al., 1988; Brown and Sauer, 1993; Gray and Gray, 1997). These studies are all different from the current study in that they focus on aggregate analyses using price data rather than on an individual-level analysis of behavior.

The remainder of the paper is organized as follows. In Section 1, we describe the sports betting market, our data, and our variables of interest. In Section 2, we conduct a panel-level analysis to study whether individuals’ average behavior is affected by team performance streaks and how individuals react to streaks of different lengths. In Section 3, we investigate whether individuals’ behavior is driven by sentiment or superior information related to market mispricings, by testing

³According to the H2 Gambling Capital report from February 2013, close to \$1 trillion of wagers was placed in sports betting markets globally in 2012.

⁴See the 2010 British Gambling Prevalence Survey prepared for Great Britain’s Gambling Commission.

⁵Grinblatt and Keloharju (2009) show in a large sample of Finnish investors that sensation seekers (measured by the number of speeding tickets received) and those who exhibit overconfidence (measured by military psychological tests) trade more frequently. Dorn and Sengmueller (2009) find that German investors who reported in a survey that they enjoy investing or gambling turn over their stock portfolios at twice the rate of their peers. Gao Bakshi and Lin (2015) find that when there is a large jackpot lottery in Taiwan, some individuals substitute toward buying lottery tickets and away from trading stocks. Kumar, 2009 shows that many stock market investors prefer stocks with lottery characteristics.

whether individuals generate excess returns. In Section 4, we conduct an individual-level analysis to study the prevalence of biased behavior in the cross-section. In Section 5, we repeat our analyses using alternative measures of team performance streaks. In Section 6, we conclude.

1 Data

1.1 The sports betting market

Gambling in general, and sports betting in particular, is a hugely popular and economically significant activity. In terms of participation, recent surveys in the U.S. show that two thirds of the adult population participates in some form of gambling each year, and that about 50% of the adult population and 67% of college students bet on sports. In terms of monetary activity, it is estimated that close to \$1 trillion of wagers was placed in sports betting markets globally in 2012.⁶ Out of this, it is estimated that about 70% is wagered on soccer matches and 32% is placed on legal online sportsbooks.

A sports betting market is run by a bookmaker, who sets the price of a unit payout for each possible outcome of each event on which individuals can place wagers. For example, an outcome with price 0.8 (quoted as having odds 1.25) implies that an individual who backs this outcome to win will make a profit of €25 for each €100 staked if he wins.⁷ The traditional model of sportsbook pricing suggests that the bookmaker sets prices so that he minimizes his risk and earns a commission. For instance, if at the current prices one of the outcomes in an event is heavily bet, the bookmaker should increase its price to shift betting activity to the other outcomes so that the book would become balanced, in which case the total payout to winners, hence his commission, would be the same regardless of the realized outcome. However, in fixed-odds sports betting markets like the one we study here, large and/or frequent changes of the quoted prices are generally not observed, which calls into

⁶For survey results on participation, see the 2008 report by ESPN and data provided by the National Council on Problem Gambling. For statistics on the size of the sports betting market, see the February 2013 report from H2 Gambling Capital.

⁷It is important to note that, even though prices may shift over time, the payoff of each bet is determined by the prices prevailing at the time the bet was placed.

question the view that bookmakers simply set prices to balance the book. Some studies (e.g., Paul and Weinbach, 2008, 2009, 2012) suggest that this may be because bookmakers find it optimal to set prices that are efficient (to the best of their knowledge) even if this implies that the book may at times become slightly unbalanced, as this strategy saves on the costs associated with perfectly balancing the book at all times and for all matches hence could lead to greater long-run profits. Other studies (e.g., Pope and Peel, 1989; Levitt, 2004) suggest that bookmakers optimally exploit individuals' biases and maximize their profits by setting prices between the efficient ones and those at which the book is balanced. Even though a thorough study of the bookmaker's price-setting behavior is beyond the scope of and the data available for this paper, our findings (see Section 3) and the discussions we had with the bookmaker who provided our data give support to a pricing behavior that is close to the efficient pricing model. This implies that, conditional on the quoted prices, previous match outcomes should not predict future outcomes, which is consistent with the random-walk hypothesis.

1.2 A novel individual-level data set

We use a panel data set of individual betting activity obtained from a large European online sports betting company. Our data contain detailed information about the betting histories of 500 randomly selected individuals over a period of 5 years, from October 2005 to November 2010. We focus on bets placed by these individuals on the final outcome of soccer matches, i.e., on a home-team win, a draw, or an away-team win.⁸ Since our goal is to examine whether people are biased toward teams with specific characteristics, in our analysis we utilize only bets backing the home team or the away team; hence, we drop 15% of all bets because they back the draw outcome. In addition, we drop 10 individuals who have placed fewer than 5 bets. Thus, our final sample contains 99,969 bets from 490 individuals corresponding to soccer matches from over 100 competitions that include many national leagues worldwide, as well as international competitions such as the World Cup.

For each bet placed by each individual in our sample, we observe the following information: i)

⁸Even though we have information about bets placed on all sports, we restrict our attention to standard bets on soccer matches, because this is the most active market segment with the highest transaction volume by far. Also, historical data, which are necessary for our analysis, are significantly more readily available for soccer than for other sports.

bet date; ii) bet event (e.g., Premier League match between Chelsea and Liverpool); iii) outcome chosen (i.e., home or away win); iv) bet amount; and v) prices associated with all outcomes of the bet event at the time the bet was placed. In addition, we have information about individual characteristics, like the gender, age, and residence zip code.

Furthermore, we use several online sources to obtain a comprehensive list of 119,539 soccer matches that we consider to be the universe of matches that were available in the sportsbook under study during the years covered by our sample; from the same sources, we obtain the results of these matches.⁹ The list of matches is used to create the market portfolio, while the historical match results are used to create measures of team past performance and to calculate bet returns.

1.3 Variables of interest

To examine how individuals react to patterns in sequences of past team performance, we construct the following variables. First, for each match, we calculate the duration, i.e., the number of matches, of the active performance streak for each participating team. The team/match-specific *Streak* variable is positive (negative) if the team participating in the match is on a winning (losing) streak at the time of the match. Individuals are unlikely to mentally extend streaks across league seasons and competitions, because significant managerial and roster changes often occur across seasons, and because the different competitions (e.g., the domestic league and an international competition) in which a team may participate can vary greatly in strength. As a result, our performance measure does not span across league seasons and competitions and it is therefore not defined for teams participating in the first match of a league season or competition. Draw outcomes are assumed to maintain (but not increase) a team's streak. Furthermore, to allow for a separate effect of winning and los-

⁹The list of matches available through several bookmakers (including the bookmaker who has provided our data) and match results are obtained from three data sets: i) www.football-data.co.uk provides a data set that covers all major and many minor European national leagues for the whole sample period; ii) www.matchstatistics.com provides a data set that covers many major and minor national and international leagues and competitions worldwide for the period up to mid 2009; and iii) www.betfair.com provides a data set that covers many major and minor national and international leagues and competitions worldwide for the whole sample period, though we only use data from this source for the period starting from mid 2009. We use a different source for data on European national leagues because this source has excellent coverage of these leagues but no coverage of other competitions, and we use two different sources for data on the remaining competitions because each of them has great coverage for a different part of our sample period.

ing streaks and for a possible non-monotonic effect of streak length on behavior, we construct four dummy variables that indicate long winning, short winning, short losing, and long losing streaks. In particular, *ShortWin* (*LongWin*) is a dummy indicating that a team is on a winning streak of at most (greater than) 3 matches, where this cutoff point is chosen to maximize the explanatory power of our regressions. *ShortLose* and *LongLose* variables are defined analogously for losing streaks.¹⁰

We also use the following variables as controls in our analysis. First, we use the team/match-specific variable *Price*, which is the price (expressed in decimal odds) associated with each team at the time of each match in which it participates. Since teams on winning streaks tend to be favored to win, including this variable in our analysis mitigates concerns that our results are driven by a preference to back teams that are associated with, e.g., favorable odds. Second, we construct the team/match-specific dummy variable *Home*, which indicates if a team plays a match at its home stadium or away. This variable controls for the possibility that our results are driven by a preference toward home teams, which may systematically tend to be on a specific performance streak (e.g., because teams are more likely to win home games due to home-field advantage, and because teams may alternate between playing home and away matches due to scheduling issues). Third, we construct a proxy for team visibility, to account for the possibility that our results are driven by a preference toward visible teams, which systematically tend to be on (long) winning streaks. Our primary measure of team visibility is based on teams' rankings, on the basis that highly ranked teams tend to be more visible as they attract the media's attention; in particular, we construct a team/season-specific dummy variable, *VisibleTeam*, that indicates if a team was ranked among the top 20 soccer clubs (top 5 national teams) according to the UEFA club coefficients (FIFA World Rankings) of the preceding season.¹¹ Finally, we construct the individual-specific variable *MostBetTeam*, which indicates the team that has been backed most frequently by each individual; this controls for the pos-

¹⁰Alternative definitions of streak in which draws terminate or increase a team's streak yield similar results. Furthermore, as we discuss in more detail in Section 5, we have repeated our analysis using alternative measures of past team performance in which i) we exclude losses (wins) in which the team played against heavy favorites (longshots), and ii) we underweight losses (wins) in which the team played against heavy favorites (longshots), and overweight losses (wins) in which the team played against heavy longshots (favorites). In all cases, our qualitative results remain the same.

¹¹FIFA (Fédération Internationale de Football Association) is the international governing body for soccer, and UEFA (Union of European Football Associations) is the administrative body for soccer in Europe.

sibility that our results are driven by frequent bets, e.g., due to loyalty or fandom, by each individual on one specific team, which systematically happens to be on a specific performance streak.¹²

1.4 Summary statistics

In Table 1, we present summary statistics for our data. In Panel A, we present the characteristics of the individuals in our sample. The vast majority (93%) of individuals are men and the mean (median) age is 34.3 (33) years.¹³ Each individual, on average, has placed €2,590 on 204 wagers, and has participated in the sportsbook for a period of 7.3 months, with an average betting frequency of approximately once a week.

In Panel B, we present the characteristics of the bets we use in our analysis. 68% of the bets back the home team to win, while the remaining 32% back the away team to win. The odds of the selected outcome range from 1.01 to 251, with a mean (median) of 2.01 (1.75). 17% of the bets back a visible team, while only 3% back the most-frequently-backed team of the individual placing the bet. The performance streak of the team that is backed to win ranges from -20 to 25, with a mean (median) of 1.15 (1).

For the list of 119,539 soccer matches that we consider to be the universe of matches, in Panel C we present the characteristics for the corresponding wagers backing the home team or the away team to win. The closing odds range from 1.01 to 1001, with a mean (median) of 3.35 (2.60). 3% of the teams involved in these matches are visible. The length of team performance streaks ranges from -25 to 27, with almost half of the teams being on winning streaks and the other half on losing

¹²In unreported results, we have also used two slightly different definitions of team visibility and an alternative measure to control for loyalty/fandom, and the results remain qualitatively unchanged. For visibility, the first alternative measure is a team-specific dummy variable that indicates if a team belongs to the 20 most-widely supported soccer clubs in the world based on the number of their fans according to a 2010 survey by the research company SPORT+MARKT; and the second is a team/season-specific dummy variable that indicates teams that were included in the Forbes list of the 20 highest-net-worth soccer clubs for the preceding season. For loyalty/fandom, the alternative measure uses each individual's residence zip code together with the location of each team's stadium and identifies the teams that are local to each individual.

¹³Interestingly, these average characteristics are not very different from the average characteristics for samples of individuals who invest in the stock market through online brokers. In a sample of 1,607 U.S. individuals who switched from a phone-based service and made online trades between 1991 and 1996, Barber and Odean (2002) find that 86% of investors are men and that the mean (median) age is 49.6 (48) years. In a more recent sample of 3,079 German individuals holding online brokerage accounts between 1997 and 2001, Glaser (2003) finds that 95% of investors are men and that the mean (median) age is 40.8 (39) years.

streaks.

2 Individual trading behavior

In this section, we study whether individuals' behavior is affected by trends in teams' past performance, and, in particular, whether individual reaction to streaks of different length is more consistent with the behavior predicted by the model of Barberis, Shleifer and Vishny (1998) or that predicted by the model of Rabin (2002). For this purpose, we first use the whole panel data set consisting of all wagers from all individuals in our sample to examine whether individuals' portfolios systematically differ from a market portfolio, and whether team past performance affects individual bet-selection decisions conditional on the observed wagers. Then, in Section 4, we revisit these questions at the individual-level and estimate the proportion of individuals who exhibit biased behavior consistent with one or the other model.

Barberis, Shleifer and Vishny (1998) propose a model in which the decision-maker erroneously believes that a random-walk earnings process switches between a 'reversal' and a 'continuation' regime. As a result of this belief, investors rely on patterns in past performance sequences to assess the likelihood of which regime governs the earnings process. In particular, the model of Barberis, Shleifer and Vishny (1998) predicts that people believe that the probability of streak continuation increases monotonically with its length. On the other hand, Rabin (2002) proposes a model in which outcomes are generated by draws from an urn with replacement, but the decision-maker erroneously believes they are generated with replacement only every $K > 1$ draws. When the individual does not know, hence must infer, the urn's rate, he exhibits the gambler's fallacy after short streaks and the hot-hand fallacy after long streaks; the interaction of these two effects implies a non-monotonic (first decreasing, then increasing) relation between streak length and the predicted probability of streak continuation.

2.1 Composition of individual versus market portfolio

Here, we test whether individuals overweight/underweight teams that are on long/short winning/losing performance streaks in their weekly betting portfolios relative to a contemporaneous ‘market’ portfolio, to be defined below. Teams are first divided into groups based on their streak length. Then, based on the price associated with them in the sportsbook at the time of each match in which they participate, teams are further subdivided into four price groups — strong favorites, favorites, longshots, and strong longshots — where the price cutoffs correspond to the 5th, 50th, and 95th percentile (1.36, 2.60, and 7.50, respectively, expressed in decimal odds) of the distribution of closing prices in the universe of bets.¹⁴ For each team group g , we compute the portfolio weight that individual n allocates to this group in week t as

$$Indiv_{ngt} := \frac{B_{ngt}}{\sum_g B_{ngt}}, \quad (1)$$

where B_{ngt} is the amount of money staked by individual n on team group g in week t .¹⁵ In addition, we compute the weight that corresponds to team group g in the market portfolio in week t as

$$Market_{gt} := \frac{N_{gt}}{\sum_g N_{gt}}, \quad (2)$$

where N_{gt} is the number of teams that belong to group g and participate in a match that an individual could have placed a wager on in week t . Essentially, N_{gt} is the number of available wagers (assets) that back team group g in week t , and $Market_{gt}$ is the weight of team group g in week t in the equal-weighted market portfolio that buys all available assets, or the expected weight of team group g in week t in a portfolio constructed by picking bets at random.

Since the number of matches available in the sportsbook every week is very large and individuals are unlikely to consider all of them when selecting their wagers, we exclude from the market portfolio all matches from obscure leagues. That is, instead of using the whole universe of 119,539 matches to determine N_{gt} in the calculation of $Market_{gt}$, we use a more restricted ‘market’ con-

¹⁴We note that these cutoffs are very stable over time, so redoing the analysis using year-specific cutoffs yields very similar results.

¹⁵The results we present in this section are very similar if instead of value-weighted we use equal-weighted portfolios for individuals, i.e., we define $Indiv_{ngt} := \frac{N_{ngt}}{\sum_g N_{ngt}}$, where N_{ngt} is the number of teams that belong to group g and are backed by individual n in week t . Using monthly instead of weekly portfolios also yields very similar results.

sisting of 60,822 matches; these cover more than 90% of all wagers placed by individuals in our sample. It is important to note that using the whole universe of available matches instead of this reduced subset does not affect our results.

In Table 2, we present results for the two-way split of teams into groups by streak length and price. In particular, we present the mean portfolio weight $Indiv_{ngt}$ that individuals allocate to each team group (in columns labeled ‘Indiv.’), the mean weight $Market_{gt}$ of each team group in the contemporaneous market portfolio (in columns labeled ‘Market’), and the ratio of the individual-to-market portfolio weights for each team group (in columns labeled ‘Ratio’). We find that, on average, individuals allocate significantly higher portfolio weights to teams that are on winning rather than on losing streaks (65% versus 35%), even though the two groups are available in roughly equal proportions in the market (51% versus 49%). In particular, 16% of the average individual portfolio is allocated to teams on long winning streaks, while these teams make up only 8% of the market portfolio; and approximately 3% of the average individual portfolio is allocated to teams that are on long losing streaks, while these teams constitute on average 7% of the market portfolio. Furthermore, we observe that the ratio of individual-to-market portfolio weights is increasing with teams’ streak, from 0.38 to 0.77 to 1.11 to 2.12 for teams on long losing, short losing, short winning, and long winning streaks, respectively. This finding provides a first indication that individuals’ expectation of streak continuation increases with streak length, as predicted by the model of Barberis, Shleifer and Vishny (1998). But since teams on winning streaks tend to be favored to win, it is possible that individuals simply prefer to back teams that are associated with favorable odds. We examine this possibility by comparing individual with market portfolio weights within the aforementioned four price groups: strong favorites, favorites, longshots, and strong longshots. In Figure 1, we plot how the ratio of individual-to-market portfolio weights (reported in Table 2) varies within each price group. Even though individuals show a preference for teams with smaller odds (i.e., favorites), it is clear that, within all price groups, individuals show a preference for teams on winning streaks. Hence, we conclude that individual preference for winning teams is only partially driven by a preference for favorite odds groups.

2.2 Conditional holding decision

The analysis of decisions in a real-world market setting is complicated by the fact that the alternatives that individuals consider when making their selections are unobserved. As a result, it is important to supplement the preceding unconditional analysis with an analysis that conditions on the observed transactions. Similar to studies of stock-buying decisions that examine how stock past performance affects whether a transaction is a buy or a sell conditional on the stocks selected by each individual (e.g., Grinblatt and Keloharju, 2001), in this section we examine how team past performance affects which of the two teams participating in each match is backed to win, conditional on the matches selected by each individual.

To examine the conditional team selection decision, we use the linear probability model

$$Bet_{nim} = x'_{im}\beta + z'_{im}\gamma + \varepsilon_{nim}, \quad (3)$$

where Bet_{nim} takes value one if team i is backed by individual n in match m and zero if its opponent is backed; x_{im} is a vector of past performance measures for team i and its opponent at the time of match m ; z_{im} is a vector that contains the constant and characteristics of team i and its opponent at the time of match m , which are used as controls; and β , γ and ε_{nim} are the regression coefficients and residuals, respectively.¹⁶ The past performance measures vary by specification, as described below, and the control variables contain the price associated with team i in match m and dummies indicating whether team i is the home team in match m , whether it is visible during the year in which match m takes place, and whether it is the most-frequently-backed team for individual n . As mentioned in Section 1.3, including these controls in our analysis accounts for the possibility that individuals' preference toward performance streaks is driven by a preference to back teams that, e.g., have favorable odds, play at home, are visible, or are frequently backed due to loyalty or fandom. The analysis includes two observations, one for the backed team and one for its opponent, for each bet placed by each individual. Since we include in the analysis multiple observations for the same match, we cluster standard errors at the match level to account for correlations in the residuals.

¹⁶All results we report in this section are qualitatively identical to those from a logit model.

First, we assume a linear effect of performance streak on individuals' bet choices, in which case x_{im} in Equation 3 is the difference between the *Streak* measure for team i and its opponent in match m , which we denote $\Delta Streak$. Estimating a coefficient on $\Delta Streak$ that is significantly different from zero would indicate that individuals exhibit, on average, a preference toward performance streaks. Subsequently, to test whether behavior responds to streak length monotonically (consistent with the theory of Barberis, Shleifer and Vishny, 1998) or not (consistent with the theory of Rabin, 2002), we estimate alternative specifications that allow for a non-monotonic effect of streak on behavior.¹⁷ In particular, we additively include in the model the ordered categorical streak dummies *LongLose*, *ShortLose*, *ShortWin*, and *LongWin*, that is, we include dummies *LongLose*, *LongLose + ShortLose*, and *LongLose + ShortLose + ShortWin* for each team, so their coefficients capture the mean differential effect between adjacent categories. For example, the coefficient on a team's *LongLose + ShortLose + ShortWin* dummy is the mean change in the dependent variable when that team switches from a *LongWin* to a *ShortWin* streak. As with the linear measure of streak, we express these variables as differences between teams. In particular, letting β_1 , β_2 , and β_3 be, respectively, the coefficients of $\Delta (LongLose_{im} + ShortLose_{im} + ShortWin_{im})$, $\Delta (LongLose_{im} + ShortLose_{im})$, and $\Delta LongLose_{im}$ and collecting terms, we can rewrite the model as

$$Bet_{nim} = \beta_1 \Delta ShortWin_{im} + \sum_{k=1}^2 \beta_k \Delta ShortLose_{im} + \sum_{k=1}^3 \beta_k \Delta LongLose_{im} + z'_{im} \gamma + \varepsilon_{nim}. \quad (4)$$

From this equation, we see that an increasing reaction to streak length, as predicted by Barberis, Shleifer and Vishny (1998), implies $\beta_1 < 0$, $\beta_2 < 0$, and $\beta_3 < 0$, while an expectation that short streaks will reverse and long streaks will continue, as predicted by Rabin (2002), implies $\beta_1 < 0$, $\beta_2 > 0$, and $\beta_3 < 0$.

In Table 3, we present the Ordinary Least Squares (OLS) regression results for the effect of past performance streaks on individuals' bet choices. Specifications 1 and 2 assume a linear effect of performance streaks on team choice, while specifications 3 and 4 allow for a non-monotonic effect,

¹⁷We note that, as we suggest in Section 1.1 and show in Section 3, information in team performance streaks is efficiently incorporated in the quoted prices and does not predict future outcomes, hence our analysis can be used to test the theories of Barberis, Shleifer and Vishny (1998) and Rabin (2002), which assume that the observed sequences of performance outcomes follow a random walk.

as discussed above. Specifications 1 and 3 include price and the home/away dummy as controls, while specifications 2 and 4 include additional controls for team visibility and the team that is most frequently backed by each individual. The results from specifications 1 and 2 show that, even under the assumption of linearity, performance streaks have a strong and statistically significant effect on individuals' behavior. In particular, we find that a unit increase in a team's streak relative to that of its opponent yields a 3% increase in the probability of backing the team. The results from specifications 3 and 4 show that individual reaction to streak length is monotonic, as predicted by Barberis, Shleifer and Vishny (1998): the probability of backing a team is the highest for teams that are on long winning streaks and it decreases progressively for teams that are on short winning, short losing, and long losing streaks. In particular, the probability of backing a team on a long winning streak is 15% higher than the probability of backing a team on a short winning streak, which in turn is 8% higher than the probability of backing a team on a short losing streak, which is 13% higher than the probability of backing a team on a long losing streak.

Thus, we conclude that individuals systematically base their choices on patterns in past performance sequences, and, as predicted by the model of Barberis, Shleifer and Vishny (1998), the longer the length of the performance streak, the more likely they are to expect that the streak will continue. Using data from experiments, Bloomfield and Hales (2002) and Asparouhova, Hertzels and Lemmon (2009) also find that patterns in past performance affect individuals' choices, though the former find evidence in favor of the Barberis, Shleifer and Vishny (1998) model while the latter find evidence in favor of the Rabin (2002) model. Using aggregate data from a real-world marketplace (the U.S. college football betting market), Durham, Hertzels and Martin (2005) also observe price changes that reveal that market participants' behavior is affected by patterns in past performance, though in a manner that is not consistent with the Barberis, Shleifer and Vishny (1998) model.¹⁸ Thus, the existing evidence is only partially consistent with our finding that, as predicted by the model of Bar-

¹⁸In detail, Durham, Hertzels and Martin (2005) find that the length of teams' performance streaks has a non-monotonic effect on point spread movements during the betting period in the U.S. college football betting market, which is inconsistent with the Barberis, Shleifer and Vishny (1998) model. It is unclear whether the change in the point spread reflects a bias or its correction, hence it is inconclusive whether the observed pattern corresponds to the same non-monotonic relation between streak length and the expectation of its continuation as implied by the model of Rabin (2002), or to the opposite pattern.

beris, Shleifer and Vishny (1998), the longer the length of the performance streak, the more likely are individuals to expect that the streak will continue. The important difference between our study and these existing studies is that we use individual-level data from a marketplace, which allow us to directly examine individual behavior from real-world decisions. On the one hand, experimental studies may present choice settings and incentives that are significantly different from those in a real market. On the other hand, though studies of market prices are informative about market participants' behavior, it is not clear whose behavior they reflect, since they aggregate the behavior of diverse participants including the market maker.

3 Sentiment-driven versus information-driven behavior

In this section, we examine whether the bias toward performance streaks we document above is driven by sentiment or information related to exploitable market mispricings. The fact that in the sports betting market all assets (wagers) are idiosyncratic and their terminal value is exogenously revealed at the conclusion of the relevant events allows us to carry out direct tests of superior performance. This stands in stark contrast to other asset markets in which i) asset returns are exposed to systematic risk, therefore tests of superior performance are joint tests of an assumed asset pricing model, and ii) true fundamentals are usually not revealed, therefore performance cannot be measured accurately since assets may, at any point in time, be mispriced. In what follows, we first conduct a test of average individual superior performance at the panel level, and in Section 4 we perform tests of superior performance at the individual level, which enable us to obtain a precise estimate of the proportion of individuals who exhibit sentiment- versus information-driven trend-chasing behavior.

If individual behavior is driven by superior information related to market mispricings, then we should find that individuals earn superior returns from wagers backing teams that are on long winning streaks relative to their other wagers. If individual behavior is driven by sentiment, then we should find that individuals do not earn superior returns from wagers on their preferred team

groups.¹⁹ We directly measure individuals' performance using their realized returns, and we estimate the model

$$Return_{nim} = \alpha_t + x'_{nim}\beta + \varepsilon_{nim}, \quad (5)$$

where $Return_{nim}$ is the rate of return realized by individual n on the wager backing team i in match m ;²⁰ α_t is a monthly time fixed effect, where t is the month during which match m takes place; x_{nim} contains the wager's characteristics (including the price and the characteristics of the teams involved), individual-specific measures of the bias toward teams that are on long winning streaks, and interaction terms between team streaks and the individual-specific measure of the corresponding bias; and β and ε_{nim} are the regression coefficients and residuals, respectively. For each individual, his bias toward teams on long winning streaks is measured by the mean ratio of individual-to-market portfolio weights allocated to these teams.²¹ Our panel analysis includes one observation for each bet placed by each individual; since it is possible that multiple individuals have placed bets on the same match, we cluster standard errors at the match level to account for correlations in the residuals.

In Table 4, we present results from the panel analysis of this test of individuals' superior performance. Specifications 1 to 3 assume a linear effect of performance streaks on realized returns, while the remaining specifications allow for a non-monotonic effect. We find that, in all specifications, all variables of interest are statistically insignificant, suggesting that individual bias toward trends in team past performance is not related to superior betting performance. In particular, the

¹⁹Berk and Green (2004) propose a model in which investors rationally chase past mutual fund performance even though it does not predict future performance. They show that if i) past performance conveys information about fund managers' stock-picking skill (e.g., due to superior information), ii) funds exhibit decreasing returns to scale in deploying their skill, and iii) there is efficient provision of capital by investors to mutual funds, then investors will move money toward funds with good past performance until their performance is pushed down to the same level as that of competing funds. Consequently, in equilibrium, investors will rationally chase past performance even though it does not predict future performance and — since the scarce resource is fund managers' ability to identify profitable opportunities — highly skilled fund managers will reap high fees while investors' expected returns will equal zero. This line of reasoning cannot rationalize the unprofitable trend-chasing behavior we observe in the sportsbook, because i) unlike a fund manager's performance, a team's performance is independent of the individuals' trading behavior (i.e., there are no decreasing returns to scale), and ii) like fund managers (and unlike fund investors), individual bettors who possess superior information directly identify profitable opportunities, hence reap the resulting benefits themselves.

²⁰Note that no commission is paid after this return is realized; it is implicitly paid by all individuals placing wagers since the return from placing a wager with unit payout on each of the possible outcomes of an event is smaller than one.

²¹Alternative measures of individual preference for performance streaks — such as i) the difference between the individual-to-market portfolio weight allocated to teams on long winning streaks and the corresponding ratio for teams on long losing streaks, ii) the average streak of the teams backed by each individual, or iii) a time-varying measure based on each individual's bets during the weeks leading up to the match, instead of all his bets — yield very similar results.

insignificant coefficients on team streaks suggest that the average return individuals generate from wagering in favor of teams that are on winning streaks is not significantly different from the return they generate from wagering against these teams. This finding indicates that betting prices are, on average, efficient with respect to teams' performance streaks and do not predict match outcomes. The insignificant coefficients on the individual-specific bias measures suggest that individuals with a stronger bias toward teams that are on long winning streaks do not generate returns that are significantly different from the returns of individuals with a weaker bias toward these teams. Finally, the interaction of the two aforementioned effects is also insignificant, suggesting that individuals who exhibit a stronger bias do not earn higher returns from backing their preferred teams (presumably, because they do not possess superior information). *F*-tests of the null that all these variables' coefficients are jointly equal to zero are also not rejected at the 10% significance level. These findings imply that the average biases which we found in Section 2 above are likely rooted in behavioral biases rather than driven by superior information related to market mispricings.

While our finding that individuals do not earn significantly *higher* returns from backing their preferred team groups is informative about the sources of individual biases, our finding that they do not earn significantly *lower* returns from these wagers is informative about how prices are set in the market under study. In particular, under the traditional assumption that the bookmaker balances the book, it seems puzzling that individuals do not experience significantly negative mean returns from their wagers on teams toward which the majority of individuals exhibit a bias. While not crucial for the conclusions of this study, it is interesting to consider why market prices are, on average, efficient despite the presence of individual biases. One possible explanation for this puzzle is that informed arbitrageurs (e.g., gambling syndicates) continuously correct any mispricings that arise due to bettor biases, hence the prices at which non-arbitrageurs trade are on average efficient. This explanation would require not only that arbitrageurs persistently make profits at the expense of the bookmaker, but also that the bookmaker frequently shifts prices through the course of betting to reflect the actions of the marginal market participant. Since we generally do not observe large and/or frequent changes of betting prices in fixed-odds betting markets (see Levitt, 2004), it is unlikely that

this scenario accounts for the efficiency of bet prices. A more plausible explanation of the puzzle is that the bookmaker does not always balance the book perfectly, but rather he deliberately sets prices near the efficient ones, occasionally allowing the book to become slightly unbalanced in the presence of biased behavior. Consistent with this explanation, several studies in different sports have found that even though sportsbooks are not always perfectly balanced, their prices are efficient (e.g., Paul and Weinbach, 2008, 2009, 2012).²² This strategy saves on the costs associated with perfectly balancing the book at all times and for all matches and is not particularly risky for the bookmaker, since the commission he charges provides a cushion against the unbalanced liabilities implied by an unbalanced book, hence overall it could lead to greater long-run profits.

4 Individual heterogeneity

The panel structure of our data that contain, on average, 204 observations per individual allows us to go beyond analyzing the mean behavior of individuals to analyzing the extent of heterogeneity in behavior across individuals. Similar to our panel analysis above, here we perform an individual-level analysis of the unconditional and conditional bet-selection decisions. Importantly, we augment this analysis with the False Discovery Rate methodology, which allows us to estimate the proportions of individuals that exhibit trend-chasing behavior and a monotonic versus non-monotonic reaction to streak length. In Section 4.1 we present our estimation methodology, and in Section 4.2 we present our results on heterogeneity.

4.1 Estimation methodology

A seemingly reasonable way to estimate the proportion of individuals who exhibit a preference toward a specific team group would involve i) performing multiple individual-level hypothesis tests of no bias toward this team group, and ii) calculating the proportion of individuals for whom the

²²The exception is Paul and Weinbach (2007), who find in the NFL betting market evidence of an unbalanced book with inefficient prices consistent with the Levitt (2004) model of sportsbook pricing, according to which bookmakers optimally exploit individuals' biases and maximize their profits by setting prices between the efficient ones and those at which the book is balanced. However, such a strategy would still lead to negative expected returns for biased individuals, which we do not observe in our data.

tests are rejected at some significance level. However, this approach does not control for the fact that the probability of incorrectly rejecting the null hypothesis (Type I error) is inevitably increased when performing multiple comparisons simultaneously. In order to properly adjust for the problem of false discoveries in this multiple testing setting, we use the powerful False Discovery Rate (FDR) approach (Benjamini and Hochberg, 1995; Storey, 2002). The FDR methodology suggests estimating the proportion π_0 of true null hypotheses (in our case, individuals who do not exhibit a bias toward a specific team group) in the population as

$$\hat{\pi}_0(\lambda) := \frac{1}{1-\lambda} \frac{\#\{\hat{p}_n : \hat{p}_n > \lambda\}}{N}, \quad (6)$$

where \hat{p}_n are the estimated p -values from testing the N null hypotheses (one for each individual) and λ is some number in $(0, 1)$, and then using this $\hat{\pi}_0$ to adjust for false rejections and acceptances of the null at some significance level $\gamma \in (0, 1)$. The intuition behind this adjustment is that, for large enough λ , the p -values that lie above λ should correspond to true null hypotheses, and since p -values corresponding to the true null are uniformly distributed on the interval $[0, 1]$, then $\hat{\pi}_0$ is an estimate of π_0 , the true proportion of nulls in the population. Then, for a significance level γ high enough so that it contains all the alternatives, we can calculate the proportion of rejections and adjust (downward) for the proportion of true nulls that are expected to be rejected at significance level γ . In short, in our example, the estimated proportion $\hat{\pi}_+$ of individuals who suffer from a bias toward, e.g., teams on long winning streaks, is equal to the proportion of individuals in the sample with positive and significant (at level γ) test statistics in the two-sided hypothesis test of no bias, adjusted for the proportion of false discoveries, i.e.,

$$\hat{\pi}_+(\gamma, \lambda) := \frac{\#\left\{\left(\hat{\beta}_n, \hat{p}_n\right) : \hat{\beta}_n > 0, \hat{p}_n < \gamma\right\}}{N} - \hat{\pi}_0(\lambda) \times \frac{\gamma}{2}, \quad (7)$$

where $\hat{\beta}_n$ is the estimated parameter for individual n . As already noted, λ and γ need to be chosen judiciously; to choose them, we use a bootstrap procedure. We also note that because our individual samples are small-to-medium sized, we calculate p -values by inverting the bias-corrected accelerated (BC_a) bootstrap confidence intervals. The proportions of individuals exhibiting biased behavior that we report below are estimated using the FDR methodology.

4.2 Proportions of biased individuals

First, we estimate the proportions of individuals who overweight/underweight specific teams in their betting portfolios relative to the ‘market’ portfolio on the basis of these teams’ performance streaks. In particular, we conduct multiple-hypothesis tests in which we test, for each individual, the null hypothesis that the composition of the weekly portfolios of teams he has backed is not different from that of contemporaneous, price-matched market portfolios. To test this hypothesis for each individual, we conduct a bootstrap test, where we construct 1,000 matching weekly portfolios by randomly sampling from a reasonably broad set of wagers that are contemporaneously available in the sportsbook and in the same price group (i.e., strong favorites, favorites, longshots, and strong longshots) as the wagers contained in the individual’s portfolio.²³ For example, if an individual places two wagers on a strong favorite team and three wagers on a longshot team in a given week, we sample two wagers from the set of strong favorite teams and three wagers from the set of longshot teams that are available in the sportsbook that week. Then, we compare the mean streak length of the teams backed by each individual across all weeks with the corresponding statistic in the bootstrap samples. Finally, we calculate bootstrap p -values and use the FDR methodology to control for the problem of false discoveries in multiple testing.²⁴ Since we compare individuals’ bet selections to a reasonably broad set of alternative wagers, the proportions of individuals we present below should provide a reasonable upper bound on the true proportions of individuals who exhibit a bias toward specific team groups. In Panel A of Table 5, we report our results based on 403 individuals with at least 5 weeks of betting activity.²⁵ We find that an overwhelming majority of individuals (83%) are consistently affected by team performance streaks when forming their betting portfolios: out of those,

²³This reasonably broad set of alternatives is the weekly equivalent of the set of 60,822 matches (which cover 90% of all wagers placed by individuals in our sample) that we use in the corresponding panel-level analysis in Section 2.1. We also note that all portfolios — both individuals’ actual portfolios and the ones generated for the bootstrap samples — are equal-weighted. Our results are very similar if we use value-weighted portfolios instead.

²⁴Following Davidson and MacKinnon (2000), we compute the p -value for each individual as $\hat{p}^*(\hat{\tau}) := 2 \min \left\{ \frac{1}{B} \sum_{j=1}^B \mathbf{1}(\tau_j^* < \hat{\tau}), \frac{1}{B} \sum_{j=1}^B \mathbf{1}(\tau_j^* > \hat{\tau}) \right\}$, where B is the number of bootstrap samples and $\mathbf{1}(\tau_j^* < \hat{\tau})$ equals one if the realized value of the test statistic $\hat{\tau}$ is higher than that of the test statistic in bootstrap sample j , τ_j^* .

²⁵Using weekly portfolios as the unit of observation, we can use approximately 80% of our initial sample of 490 individuals in this analysis. Using monthly portfolios as the unit of observation and retaining the filter of keeping individuals with at least 5 observations, we can use approximately half of our initial sample, but the results are essentially unchanged. Also, using a filter of 10 observations per individual does not change the results.

78% exhibit momentum behavior by overweighting (underweighting) teams that are on longer winning (losing) performance streaks, while only 5% of the individuals exhibit contrarian behavior, by underweighting (overweighting) teams that are on longer winning (losing) performance streaks.

Subsequently, we perform individual-level OLS regressions to study how team characteristics affect each individual's decision to back one of the two teams participating in each match on which he has placed a bet (see Section 2.2 for the corresponding panel analysis).²⁶ Since this analysis relies on a narrowly-defined set of wagers among which choice is exercised, i.e., the pairs of teams that participate in the matches on which each individual has placed wagers, it is expected to yield a conservative estimate for the proportion of individuals who disproportionately bet on specific team groups. In Panel B of Table 5, we report the results of multiple-hypothesis tests, in which we simultaneously test for all individuals whether the regression coefficients on the streak variables are positive/negative. Similar to the results of our unconditional analysis, we find that 75% of the individuals are more likely to back a team on a long winning streak than a team on a long losing streak, while no individuals exhibit the opposite pattern. Overall, from the unconditional and conditional analyses we conclude that performance streaks affect the trading of the vast majority of individuals in our sample, with three quarters of the individuals exhibiting momentum behavior and less than 5% exhibiting contrarian behavior.

Furthermore, we perform individual-level OLS regressions to study whether the reaction of each individual to streak length is monotonic, as predicted by Barberis, Shleifer and Vishny (1998), or non-monotonic, as predicted by Rabin (2002). In particular, for each individual n , we estimate the model in Equation 4 and test whether the coefficients on the additive streak dummies are consistent with the model of Barberis, Shleifer and Vishny (1998) (i.e., $\beta_{1,n} < 0$, $\beta_{2,n} < 0$, and $\beta_{3,n} < 0$) or with the model of Rabin (2002) (i.e., $\beta_{1,n} < 0$, $\beta_{2,n} > 0$, and $\beta_{3,n} < 0$). Then, as above, we use the FDR methodology to compute the proportion of individuals that behave according to each model while correcting for the problem of false discoveries. In Panel B of Table 5, we report the

²⁶In this analysis, we use individuals with at least 30 observations, which means that we can use 89% of the 490 individuals in our sample. Using individuals with at least 60 observations reduces the proportion of individuals we can use, but yields similar results.

results of our multiple-hypothesis tests. We find that the majority of individuals (55%) exhibit a monotonic response to streak length as implied by Barberis, Shleifer and Vishny (1998), while only a small minority (8%) exhibits the non-monotonic reaction that is implied by Rabin (2002).²⁷ That is, approximately seven times as many individuals exhibit a monotonic response to streak length as exhibit a non-monotonic response, hence this individual-level real-market evidence is strongly in favor of the model proposed by Barberis, Shleifer and Vishny (1998).

Finally, we study whether there are any individuals in our sample who exhibit a preference toward teams on long winning streaks due to a rational profit motive to exploit a market mispricing rather than due to sentiment toward these teams. In this analysis, we estimate the model in Equation 5 individual-by-individual, where we exclude individual-specific regressors and we include individual-specific coefficients β_n . In Panel C of Table 5, we report the results of multiple-hypothesis tests, where we test for each individual whether the coefficient on the streak variable is positive/negative, and as above, we utilize the FDR methodology to estimate the proportion of individuals who generate higher/lower returns from backing teams on long winning streaks. The results in the table show that the proportion of individuals who generate significant excess returns from backing teams that are on long winning streaks is negligible. This implies that, not just for the average but for almost all individuals who exhibit trend-chasing behavior, this behavior is driven by a mistaken belief that historical performance predicts future performance rather than by superior information related to market mispricings.

²⁷We report results from a high-power but less conservative F -test of the null $(\beta_{1,n}, \beta_{2,n}, \beta_{3,n})' = \mathbf{0}$ versus the alternative $(\beta_{1,n}, \beta_{2,n}, \beta_{3,n})' \neq \mathbf{0}$. Among the individuals for whom the null is rejected, those with estimated coefficients $\hat{\beta}_{1,n} < 0, \hat{\beta}_{2,n} < 0, \hat{\beta}_{3,n} < 0$ are counted as rejections consistent with Barberis, Shleifer and Vishny (1998), while those with $\hat{\beta}_{1,n} < 0, \hat{\beta}_{2,n} > 0, \hat{\beta}_{3,n} < 0$ are counted as rejections consistent with Rabin (2002). A lower-power but more conservative test would be to test the null $\neg \left((\beta_{1,n}, \beta_{2,n}, \beta_{3,n})' < \mathbf{0} \right)$ versus the alternative $(\beta_{1,n}, \beta_{2,n}, \beta_{3,n})' < \mathbf{0}$ for Barberis, Shleifer and Vishny (1998) and the null $\neg \left((\beta_{1,n}, -\beta_{2,n}, \beta_{3,n})' < \mathbf{0} \right)$ versus the alternative $(\beta_{1,n}, -\beta_{2,n}, \beta_{3,n})' < \mathbf{0}$ for Rabin (2002). With the low-power test, we find that 17% of individuals exhibit the behavior predicted by Barberis, Shleifer and Vishny (1998), while 0% exhibit the behavior predicted by Rabin (2002). The power of this test is known to be very poor (Liu and Berger, 1995), so these proportions are likely to be severely understated, but importantly, our qualitative result with respect to the relative prevalence of the two theories remains the same as with the high-power test.

5 Alternative measures of past team performance

In our main analysis in Sections 2 through 4 above, we measure past team performance by counting the number of successive wins/losses for each team. While this measure of streak likely captures the most natural and intuitive momentum characteristic of assets in the sports betting market, it is possible that some individuals follow momentum strategies based on alternative streak measures and are able to generate excess returns from these strategies. To examine this possibility, we have repeated our analysis using some sensible alternative definitions of streak, which take into account information contained in match prices (i.e., quality of the opponents, whether games are played home or away, players' injuries, etc.). Specifically, for all matches in our data, we first convert the prices associated with the two teams into their implied win probabilities, and use this information to calculate several alternative measures of streak.²⁸ One set of streak measures is based on counting the number of successive losses (wins) of each team excluding matches in which the team played against heavy favorites (longshots). That is, we disregard all matches in which the ratio of the win probability of the team to that of its opponent is below (above) a cutoff k , where we use several different values of k from 0.5 to 1. Another streak measure we calculate is based on counting the number of successive losses (wins) of each team, where losses (wins) are multiplied by the (inverse) ratio of the win probability of the team to that of its opponent. As a result, we underweight losses (wins) in which the team played against heavy favorites (longshots), and overweight losses (wins) in which the team played against heavy longshots (favorites).

We find that the proportion of individuals who follow momentum strategies based on any of these alternative streak measures is small (between 5% and 30% depending on the measure considered) compared to the large majority (about 80%) of individuals who follow momentum strategies based on the standard streak measure we use in our main analysis above. But our qualitative results remain the same: in all cases, i) the proportion of individuals following momentum strategies is greater than the proportion following contrarian strategies, ii) more individuals exhibit a monotonic

²⁸The prices quoted by the bookmaker on all outcomes of an event imply a probability with which each outcome is expected to occur. An outcome with price p has an implied probability of $p \times (1 - c)$, where c is the bookmaker's commission.

response to streak length as implied by Barberis, Shleifer and Vishny (1998) than the non-monotonic reaction implied by Rabin (2002), and iii) no individuals earn significant excess returns from their momentum strategies. While it is admittedly impossible to exhaustively check all possible definitions of streak, the results on these sensible alternatives should strengthen our conclusion that trend-chasing behavior is sentiment-driven rather than information-driven.

6 Concluding discussion

In this paper, we have used a novel panel data set of individual activity in a sports betting market to study the extent to which individuals' behavior is affected by momentum in teams' performance, and to test for the differential behavior predicted by two prominent behavioral theories of momentum and reversals in stock returns — the regime-shifting model of Barberis, Shleifer and Vishny (1998) and the gambler's/hot-hand fallacy model of Rabin (2002). Sports betting markets attract a large number of participants who wager substantial amounts of their own money. Besides their rising popularity and significant economic value, these markets have some attractive experimental-like features that render them an ideal test bed for behavioral research: Individuals make a large number of repeated choices, assets are short-lived, bear no systematic risk, and their true terminal value is revealed by match outcomes which are exogenous to individuals' behavior. These features enable us to perform clean tests of the theories above and of individuals' superior performance, that are not affected by a potentially misspecified asset pricing model. The long panel enables us to study behavior at the level of each individual and to classify individuals to behavioral types.

Our findings indicate that the majority of individuals use momentum-based trading strategies, with about three quarters of the individuals exhibiting a preference toward long winning streaks. Furthermore, seven times as many individuals exhibit a monotonic response to streak length (consistent with Barberis, Shleifer and Vishny, 1998) as exhibit a non-monotonic response (consistent with Rabin, 2002). Individuals do not earn superior returns from momentum trading, which indicates that market prices efficiently incorporate all value-relevant information contained in teams' performance streaks, and that the observed behaviors are not driven by an informational advantage

related to market mispricings. We see no reason to believe that our findings on individuals' bias toward past performance streaks are strictly confined to the market setting that we study, rather we believe that they reveal an inherent bias toward past performance sequences that individuals are likely to also exhibit in other market settings, e.g., the stock market. Therefore, our results on the effect of streaks on behavior and the relative support that we find for the model of Barberis, Shleifer and Vishny (1998) versus that of Rabin (2002) should provide valuable insights about the short-run momentum and long-run reversal asset pricing anomalies.

References

- Asparouhova, E, M Hertz, and M Lemmon.** 2009. "Inference from streaks in random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers." *Management Science*, 55(11): 1766–1782.
- Avery, C, and J Chevalier.** 1999. "Identifying investor sentiment from price paths: The case of football betting." *Journal of Business*, 72(4): 493–521.
- Barber, B, and T Odean.** 2002. "Online investors: Do the slow die first?" *Review of Financial Studies*, 15(2): 455–487.
- Barber, B, T Odean, and N Zhu.** 2009. "Systematic noise." *Journal of Financial Markets*, 12(4): 547–569.
- Barberis, N, A Shleifer, and R Vishny.** 1998. "A model of investor sentiment." *Journal of Financial Economics*, 49(3): 307–343.
- Benjamini, Y, and Y Hochberg.** 1995. "Controlling the false discovery rate: A practical and powerful approach to multiple testing." *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1): 289–300.
- Berk, J. B, and R. C Green.** 2004. "Mutual fund flows and performance in rational markets." *Journal of Political Economy*, 112(6): 1269–1295.
- Bloomfield, R, and J Hales.** 2002. "Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs." *Journal of Financial Economics*, 65(3): 397–414.
- Brown, W, and R Sauer.** 1993. "Does the basketball market believe in the hot hand? Comment." *American Economic Review*, 83(5): 1377–86.
- Burns, B. D, and B Corpus.** 2004. "Randomness and inductions from streaks: "Gambler's fallacy" versus "hot hand"." *Psychonomic Bulletin & Review*, 11(1): 179–184.
- Clotfelter, C. T, and P. J Cook.** 1993. "Notes: The "gambler's fallacy" in lottery play." *Management Science*, 39(12): 1521–1525.
- Davidson, R, and J. G MacKinnon.** 2000. "Bootstrap tests: How many bootstraps?" *Econometric Reviews*, 19(1): 55–68.

- Dorn, D, and P Sengmueller.** 2009. "Trading as entertainment?" *Management Science*, 55(4): 591–603.
- Durham, G. R, M. G Hertz, and J. S Martin.** 2005. "The market impact of trends and sequences in performance: New evidence." *Journal of Finance*, 60(5): 2551–2569.
- Edwards, W.** 1968. "Conservatism in human information processing." *Formal representation of human judgment*, 17–52.
- Fama, E. F.** 1970. "Efficient capital markets: A review of theory and empirical work." *Journal of Finance*, 25(2): 383–417.
- Gandar, J, R Zuber, T O'Brien, and B Russo.** 1988. "Testing rationality in the point spread betting market." *Journal of Finance*, 43(4): 995–1008.
- Gao Bakshi, X, and T.-C Lin.** 2015. "Do individual investors treat trading as a fun and exciting gambling activity? Evidence from repeated natural experiments." *Review of Financial Studies*, 28(7): 2128–2166.
- Glaser, M.** 2003. "Online broker investors: Demographic information, investment strategy, portfolio positions, and trading activity." Unpublished Paper.
- Gray, P. K, and S. F Gray.** 1997. "Testing market efficiency: Evidence from the NFL sports betting market." *Journal of Finance*, 52(4): 1725–1737.
- Grinblatt, M, and M Keloharju.** 2001. "What makes investors trade?" *Journal of Finance*, 56(2): 589–616.
- Grinblatt, M, and M Keloharju.** 2009. "Sensation seeking, overconfidence, and trading activity." *Journal of Finance*, 64(2): 549–578.
- Jorgensen, C. B, S Suetens, and J.-R Tyran.** 2011. "Predicting lotto numbers." Unpublished Paper.
- Keloharju, M, S Knuifer, and J Linnainmaa.** 2012. "Do investors buy what they know? Product market choices and investment decisions." *Review of Financial Studies*, 25(10): 2921–2958.
- Kumar, A.** 2009. "Who gambles in the stock market?" *Journal of Finance*, 64(4): 1889–1933.
- Levitt, S. D.** 2004. "Why are gambling markets organised so differently from financial markets?" *Economic Journal*, 114(495): 223–246.
- Liu, H, and R. L Berger.** 1995. "Uniformly more powerful, one-sided tests for hypotheses about linear inequalities." *The Annals of Statistics*, 23(1): 55–72.
- Loh, R. K, and M Warachka.** 2012. "Streaks in earnings surprises and the cross-section of stock returns." *Management Science*, 58(7): 1305–1321.
- Moskowitz, T. J.** 2015. "Asset pricing and sports betting." Unpublished Paper.
- Odean, T.** 1998. "Are investors reluctant to realize their losses?" *Journal of Finance*, 53(5): 1775–1798.
- Paul, R. J, and A. P Weinbach.** 2007. "Does sportsbook.com set pointspreads to maximize profits? Tests of the Levitt Model of sportsbook behavior." *Journal of Prediction Markets*, 1(3): 209–218.
- Paul, R. J, and A. P Weinbach.** 2008. "Price setting in the NBA gambling market: tests of the Levitt model of sportsbook behavior." *International Journal of Sport Finance*, 3: 137–145.
- Paul, R. J, and A. P Weinbach.** 2009. "Sportsbook behavior in the NCAA football betting market: Tests of the traditional and Levitt models of sportsbook behavior." *Journal of Prediction Markets*, 3(2): 21–37.
- Paul, R. J, and A. P Weinbach.** 2012. "Sportsbook pricing and the behavioral biases of bettors in the NHL."

- Journal of Economics and Finance*, 36(1): 123–135.
- Pope, P. F, and D. A Peel.** 1989. “Information, prices and efficiency in a fixed-odds betting market.” *Economica*, 56(223): 323–41.
- Rabin, M.** 2002. “Inference by believers in the law of small numbers.” *Quarterly Journal of Economics*, 117(3): 775–816.
- Rabin, M, and D Vayanos.** 2010. “The gambler’s and hot-hand fallacies: Theory and applications.” *Review of Economic Studies*, 77(2): 730–778.
- Storey, J.** 2002. “A direct approach to false discovery rates.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3): 479–498.
- Tversky, A, and D Kahneman.** 1971. “Belief in the law of small numbers.” *Psychological bulletin*, 76(2): 105–110.
- Tversky, A, and D Kahneman.** 1974. “Judgment under uncertainty: Heuristics and biases.” *Science*, 185: 1124–1131.

Table 1: Summary Statistics

This table presents summary statistics for the data we use in our analysis. Panel A presents the characteristics of the 490 individuals in the sample. *Female* is a dummy indicating gender. *Age* is in years. *Number of bets (Value of bets)* is the number (value) of bets placed per individual. *Number of bet months* is the number of months during which an individual places at least one bet. Panel B presents the characteristics of the 99,969 bets placed by the individuals in our sample on the home team or the away team. *Home* and *Away* are dummies indicating the selected outcome. *Price* is the price — expressed as decimal odds — associated with the selected outcome. *Streak* is the duration — the number of matches — of the active winning or losing streak of the backed team at the time of the match; positive (negative) values indicate winning (losing) streaks, and the draw outcome is assumed to maintain a team’s current streak. *VisibleTeam* is a dummy indicating bets that back teams that are highly ranked according to the previous season’s annual FIFA (for national teams) or UEFA (for clubs) rankings. For each individual, *MostBetTeam* indicates his bets that back his most-frequently backed team. Panel C presents the characteristics of the bets (two for each match, one backing the home team and one backing the away team) available in the sportsbook during our sample period. *Price*, *Streak*, and *VisibleTeam* are defined as above.

Panel A: Characteristics of individuals

	N	Mean	Median	Std. Dev.	Min	Max
Female	490	0.07	0	0.25	0	1
Age	490	34.27	33	9.53	18	67
Number of bets	490	204.02	112	286.80	6	2,642
Value of bets	490	2,590	570	8,152	7.50	128,860
Number of bet months	490	7.27	5	6.94	1	56

Panel B: Characteristics of bets placed

	N	Mean	Median	Std. Dev.	Min	Max
Home	99,969	0.68	1	0.47	0	1
Away	99,969	0.32	0	0.47	0	1
Price	99,969	2.01	1.75	1.47	1.01	251
Streak	78,792	1.15	1	2.99	-20	25
VisibleTeam	99,969	0.17	0	0.37	0	1
MostBetTeam	99,969	0.03	0	0.16	0	1

Panel C: Characteristics of bets available in the sportsbook

	N	Mean	Median	Std. Dev.	Min	Max
Price	239,078	3.35	2.60	5.14	1.01	1001
Streak	212,847	0.14	1	2.83	-25	27
VisibleTeam	239,078	0.03	0	0.16	0	1

Table 2: Composition of Individual and Market Portfolio by Odds Group and Team Past Performance Streaks

This table shows the weights that individuals allocate to various team groups in their betting portfolios and the weight of the respective groups in the market portfolio. Teams are divided into groups based on their past performance, and further subdivided into groups based on the price associated with them in the sportsbook. Columns labeled ‘Indiv.’ report the mean across individuals of the time-series mean of the shares of weekly portfolio value wagered by each individual on each team group. Columns labeled ‘Market’ report the cross-sectional mean of the time-series mean of the proportion of all bets from a broad set of wagers that are available in the sportsbook each week that involve this team group. Columns labeled ‘Ratio’ report the ratio of the corresponding values in the columns labeled ‘Indiv.’ and ‘Market’. Teams are on a short (long) winning/losing streak if they are on a winning/losing streak that is at most (more than) 3 matches long. *Strong Favorite*, *Favorite*, *Longshot*, and *Strong Longshot* indicate prices shorter than 1.36 (the 5th percentile), between 1.36 and 2.6 (the 50th percentile), between 2.6 and 7.5 (the 95th percentile), and longer than 7.5, respectively, where prices are expressed in decimal odds.

	All			Strong Favorite			Favorite			Longshot			Strong Longshot		
	Indiv.	Market	Ratio	Indiv.	Market	Ratio	Indiv.	Market	Ratio	Indiv.	Market	Ratio	Indiv.	Market	Ratio
LongLose	2.5%	6.5%	0.38	0.06%	0.02%	2.54	1.85%	2.03%	0.91	0.55%	3.72%	0.15	0.06%	0.78%	0.07
ShortLose	32.4%	42.0%	0.77	3.52%	0.97%	3.63	24.16%	18.55%	1.30	4.36%	19.99%	0.22	0.37%	2.44%	0.15
ShortWin	48.7%	43.7%	1.11	6.93%	1.81%	3.82	35.55%	18.88%	1.88	5.67%	20.99%	0.27	0.53%	2.06%	0.26
LongWin	16.4%	7.7%	2.12	3.55%	0.89%	3.99	11.55%	4.50%	2.56	1.28%	2.20%	0.58	0.03%	0.13%	0.22

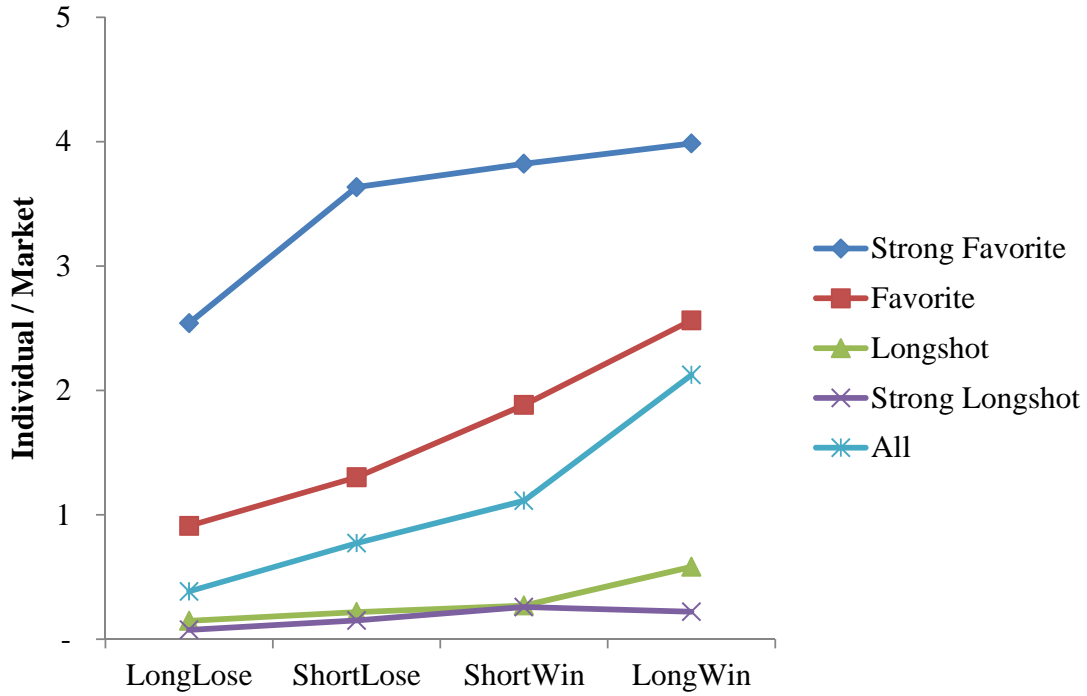


Figure 1: Plot of the ratio of individual-to-market portfolio weights of wagers grouped by the backed team's performance streak (i.e., short/long winning/losing). The ratio of individual-to-market portfolio weight for all wagers in each group is plotted using a star-marked line. In addition, wagers are further sub-divided into groups by the backed team's price, so for each sub-group of wagers the ratio is also plotted for strong favorites (diamond-marked line), favorites (square-marked line), longshots (triangle-marked line), and strong longshots (x-marked line).

Table 3: Conditional Bet-selection Decision

This table presents OLS estimates from a linear probability model in which the dependent variable is a dummy that takes the value one if a team is backed, and zero if a team is not backed by the individual placing a bet on a given match. The number of observations for each individual corresponds to the number of teams participating in the matches he has wagered on. *Price* is the price — expressed as decimal odds — associated with the team. *Home* is a dummy variable indicating whether the team is playing at home. *VisibleTeam* is a dummy indicating teams that are highly ranked according to the previous season’s annual FIFA (for national teams) or UEFA (for clubs) rankings. *MostBetTeam* indicates the team that is most-frequently backed by the individual placing the bet. *Streak* is the duration — the number of matches — of the active winning or losing streak of the team at the time of the match; positive (negative) values indicate winning (losing) streaks, and the draw outcome is assumed to maintain a team’s active streak. *ShortWin/ShortLose* (*LongWin/LongLose*) is a dummy indicating teams that are on a winning/losing streak of at most (more than) 3 matches long; because we are interested in testing for a monotonic response to the ordered streak categories, these dummies are included additively in the regression, i.e., we include *LongLose*, *LongLose+ShortLose*, and *LongLose+ShortLose+ShortWin*, so their coefficients capture the mean differential effect between adjacent categories. Δ denotes the difference between measures for the team and its opponent in the match. *t*-statistics using standard errors clustered at the match level are reported below the coefficients. */**/** indicate significance at the 10%/5%/1% levels.

	(1)	(2)	(3)	(4)
Constant	0.575 *** 31.710	0.536 *** 32.150	0.585 *** 31.847	0.545 *** 32.107
Price	-0.055 *** -14.436	-0.046 *** -13.121	-0.057 *** -14.813	-0.048 *** -13.437
Home	0.246 *** 24.293	0.262 *** 28.307	0.242 *** 23.679	0.258 *** 27.551
Δ VisibleTeam		0.171 *** 20.720		0.172 *** 20.314
Δ MostBetTeam		0.220 *** 25.276		0.223 *** 25.399
Δ Streak	0.030 *** 28.702	0.028 *** 31.886		
Δ (LongLose+ShortLose+ShortWin)			-0.155 *** -21.094	-0.142 *** -21.430
Δ (LongLose+ShortLose)			-0.079 *** -18.000	-0.074 *** -18.389
Δ LongLose			-0.134 *** -18.161	-0.133 *** -18.950
Number of Obs	157,584	157,584	157,584	157,584
Adjusted R^2	35.19%	37.83%	34.65%	37.33%

Table 4: Realized Returns

This table presents results from OLS regressions in which the dependent variable is the rate of return realized by each bet placed by individuals in our sample. The explanatory variables include: the price associated with the selected team; a dummy variable indicating if the selected team plays at home; the difference (denoted by Δ) between measures of team visibility, indicators for the most frequently backed team by each individual, and measures of team past performance for the selected team and its opponent; individual-specific measures of the preference for teams that are on a long winning streak; and interaction terms. *Price* is the price — expressed as decimal odds — associated with the team. *Home* is a dummy variable indicating whether the team is playing at home. *VisibleTeam* is a dummy indicating bets that back teams that are highly ranked according to the previous season’s annual FIFA (for national teams) or UEFA (for clubs) rankings. *MostBetTeam* indicates bets that back the team that is most-frequently backed by the individual placing the bet. *Streak* is the duration — the number of matches — of the active winning or losing streak of the backed team; positive (negative) values indicate winning (losing) streaks. *ShortWin/ShortLose (LongWin/LongLose)* is a dummy indicating bets on teams that are on a winning/losing streak of at most (more than) 3 matches long. The individual-specific preference for teams on long winning streaks (*LongWinPref*) is proxied by the mean ratio of individual-to-market portfolio weights allocated to these teams across all weeks in Specifications 2 and 5, and across the most recent 4 weeks in Specification 3. All specifications include monthly time fixed effects. *t*-statistics using standard errors clustered at the match level are reported below the coefficients. */**/* indicate significance at the 10%/5%/1% levels.

	(1)	(2)	(3)	(4)	(5)
Price	-0.003	-0.003	-0.000	-0.003	-0.003
	-0.237	-0.223	-0.025	-0.209	-0.194
Home	0.003	0.003	0.008	0.007	0.006
	0.185	0.164	0.420	0.355	0.342
Δ VisibleTeam	0.017	0.017	0.016	0.016	0.016
	0.798	0.784	0.720	0.738	0.726
Δ MostBetTeam	0.048	0.047	0.050	0.046	0.046
	1.500	1.483	1.430	1.462	1.457
Δ Streak	0.000	-0.000	0.001		
	0.057	-0.183	0.675		
Δ ShortWin				-0.009	-0.012
				-0.459	-0.472
Δ ShortLose				-0.024	-0.028
				-1.235	-1.050
Δ LongLose				0.017	0.021
				0.624	0.559
LongWinPref		0.042	0.046		0.044
		0.512	1.278		0.552
Δ Streak*LongWinPref		0.005	-0.011		
		0.325	-1.365		
Δ ShortWin*LongWinPref					0.040
					0.251
Δ ShortLose*LongWinPref					0.044
					0.268
Δ LongLose*LongWinPref					-0.034
					-0.138
Number of Obs	78,792	78,792	70,394	78,792	78,792
Adjusted R^2	0.14%	0.14%	0.12%	0.16%	0.16%

Table 5: Classification of Individuals

This table presents the results of the individual-level analysis of the unconditional and conditional bet-selection decisions. In Panel A, we present the estimated proportions of individuals who exhibit momentum (contrarian) behavior by overweighting long winning (losing) teams in their betting portfolios relative to a ‘market’ portfolio. We perform multiple-hypothesis tests, in which we simultaneously compare for all individuals the mean streak length of backed teams across all weeks with the corresponding statistics in 1,000 bootstrap samples of price-matched portfolios. Results are based on 403 individuals with at least 5 weeks of betting activity. In Panel B, we condition on the matches wagered by each individual and we present the estimated proportions of individuals i) that exhibit momentum (contrarian) behavior by being more likely to back teams on longer winning (losing) streaks, and ii) for whom the probability of backing a team increases monotonically with streak length, as predicted by the model of Barberis, Shleifer and Vishny (1998) (BSV), or first decreases and then increases, as predicted by the model of Rabin (2002). The estimated proportions are derived from testing for all individuals whether streak coefficients from estimating a series of individual-level linear probability models (Equation 3) are positive/negative. Results are based on 445 individuals with at least 30 observations. The number of observations for each individual corresponds to the number of teams participating in the matches he has wagered on. In Panel C, we present the estimated proportions of individuals who earn higher/lower returns from backing teams that are on longer winning streaks. The estimated proportions are derived from testing for all individuals whether streak coefficients from estimating a series of individual-level OLS regressions (Equation 5) are positive/negative. Results are based on 384 individuals. The number of observations for each individual corresponds to the number of teams he has backed. In all Panels, the False Discovery Rate methodology (Storey, 2002) is used to control for the problem of false discoveries in multiple testing.

Panel A: Unconditional Holding Decision

Momentum	77.83%
Contrarian	5.21%

Panel B: Conditional Holding Decision

Momentum	75.48%
Contrarian	0.00%
BSV	55.00%
Rabin	8.00%

Panel C: Realized Returns

Higher returns	0.03%
Lower returns	0.00%