

Discussion Paper No. 2014-17

Michael Bleaney, Spiros
Bougheas and Zhiyong Li
November 2014

Do Psychological Fallacies
Influence Trading in
Financial Markets?
Evidence from the Foreign
Exchange Market



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
Fax: +44 (0) 115 95 14159
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

Do Psychological Fallacies Influence Trading in Financial Markets?

Evidence from the Foreign Exchange Market

Michael Bleaney, Spiros Bougheas*, Zhiyong Li
University of Nottingham

November 21, 2014

Abstract

Research in both economics and psychology suggests that, when agents predict the next value of a random series, they frequently exhibit two types of biases, which are called the gambler's fallacy (GF) and the hot hand fallacy (HHF). The gambler's fallacy is to expect a negative correlation in a process which is in fact random. The hot hands fallacy is more or less the opposite of this – to believe that another heads is more likely after a run of heads. The evidence for these fallacies comes largely from situations where they are not punished (lotteries, casinos and laboratory experiments with random returns). In many real-world situations, such as in financial markets, succumbing to fallacies is costly, which gives an incentive to overcome them. The present study is based on high-frequency data from a market-maker in the foreign exchange market. Trading behaviour is only partly explained by the rational exploitation of past patterns in the data, but there is also evidence of the gambler's fallacy: a tendency to sell the dollar after it has risen persistently or strongly.

Keywords: Gambler's Fallacy, Hot hand Fallacy, Foreign Exchange Market

JEL: G020

*Corresponding Author: School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom; e-mail: spiros.bougheas@nottingham.ac.uk ; tel.no. ++44-115-8466108

1. Introduction

Research in both economics and psychology suggests that, when agents predict the next value of a random series, they frequently exhibit one of two types of bias, called respectively the gambler's fallacy (GF) and the hot hand fallacy (HHF) (for a survey, see Ayton and Fischer, 2004). The gambler's fallacy is to expect a negative correlation in a process which is in fact random (for example, to believe that the next toss of a fair coin is more likely to be tails after a run of heads). The hot hand fallacy, which works in the opposite direction, is to expect that another toss of heads is more likely after a run of heads. These fallacies may seem mutually contradictory, but the empirical evidence suggests that they apply in slightly different contexts, with the HHF more likely to occur where human performance is thought to be involved. For example, Croson and Sundali (2005) find that, amongst gamblers at a roulette wheel in a casino, a streak of one colour is more likely to induce bets on the opposite colour (GF), but that bets are made on more numbers after winning than losing, as if the gambler's own performance is positively serially correlated (HHF). Though there is plenty of evidence both from laboratory experiments and from the field (lotteries and casinos) on the GF and the HHF, most of it is limited to purely random environments with no variation in signal strength (e.g. the toss of a coin), and with winning probabilities that are not affected by psychological biases, so we know remarkably little about whether these biases can be overcome when they are punished.

Financial markets are especially interesting because the wrong decisions are potentially very costly, so they offer a strong incentive for traders not to succumb to psychological biases in decision-making. Unlike at the roulette wheel or in the choice of lottery numbers, where players' choices make no difference to the chance of success because the underlying process is random, in financial markets such biases can prove to be very costly if they do not match the underlying dynamics of prices. If people learn to overcome psychological biases in environments where such biases are punished, these biases should not be observed in financial markets. Moreover financial markets offer a particularly rich field of investigation in that the behaviour of participants might be affected by characteristics of the data other than the direction of the signal (whether the price is moving up or down), such as, for example, the signal strength (how fast the price moves). On the other hand, financial markets present an extra challenge because the underlying data generating process is unknown, and the research design needs to take account of the fact that some behaviour may represent the rational exploitation of patterns in the past data.

In this paper, we use high-frequency data from the foreign exchange market (the same data as used in Evans and Lyons, 2002). The dataset gives the time, direction of trade (buy or sell the US dollar) and the price of the trade. The price varies with the direction of trade because of the bid-ask spread, and standard methods are used to take out the spread in order to estimate the mid-price (the average of buy and sell prices) at the time of each transaction. The price trend immediately before any given trade can be deduced from the estimated mid-prices for the preceding trades. Although trades are not evenly spaced in time, our results do not appear to be distorted by this. We interpret the psychological fallacies discussed above as hypotheses about the effect of exchange rate movements up to trade T on the direction of trade T (we do not have data on the size of the transaction). We define a streak as a sequence of movements of the estimated mid-price of the dollar in the same direction over successive transactions, and we investigate how the characteristics of a streak affect the buy/sell decision. The characteristics of the streak are its length (the number of transactions for which the streak has lasted) and its width (the average movement in price per transaction over the streak). The length of the streak is a measure of the persistence of exchange rate movements, and the width is a measure of the speed of these movements. If the behaviour of traders in the market is influenced by GF and/or HHF, we expect the orders placed to be affected by the properties of the streak. For example, the gambler's fallacy would suggest that persistent upward movement of the dollar would make dollar sales more likely (which we term "trading against the trend"). In other words, we expect the buy/sell decision to be predictable to some extent from the streak characteristics.

Our analysis does not rely on the assumption that the trades in the dataset are observed by all market participants (they are not). Traders can easily obtain high-frequency information about how exchange rates are evolving from dealers' quotes and computer screens, so the exchange rate information in our dataset is effectively public. Thus in our view the dataset may be interpreted as traders' reactions to high-frequency price movements.

Our results show that there is a significant tendency to trade against the streak (i.e. to bet that the streak will end), and that this effect increases with the length and width of the streak. This is apparently consistent with traders suffering from the gambler's fallacy. However, it is possible that streaks in exchange rates are indeed followed by movements in the opposite direction, and that traders are simply exhibiting rational behaviour, trying to exploit a pattern that they have observed in past data. We implement various tests to investigate whether this rational

explanation of our results stands up to scrutiny. Our results are mixed. It is not clear that streak variables explain future exchange rate movements, as the rational explanation requires. On the other hand, we find that the predictions from an autoregressive model of exchange rate returns help to explain trading behaviour. Nevertheless the streak variables retain their statistical significance, which suggests that psychological factors also play a role.

The rest of the paper is organised as follows. In Section 2, we review related studies from both the economic and the psychology literatures. In section 3, we describe the raw data and formally define the streak length and the streak width. In Section 4, we analyse the relationship among traders' behaviour (the order flow), streak length and width using a probit model. Our results in Sections 3 and 4 uncover a statistically significant relationship between the trend and the behaviour of traders. In Section 5, we control for irregular time spaces and past order flows to examine the robustness of our results. In Section 6, we examine whether the behaviour is consistent with rationality. We conclude in Section 7.

2. Related Literature

Both the gambler's fallacy and the hot hand fallacy are cognitive biases arising when individuals predict the future outcomes of a random series.

Agents who suffer from GF expect a positive (negative) shock after a negative (positive) shock. Clotfelter and Cook (1993) find that lottery players exhibit GF in a numbers game: numbers which have already been drawn become very quickly less popular among the players.

In contrast, agents who suffer from HHF expect a positive (negative) shock after a positive (negative) shock. Gilovich et al. (1985) have demonstrated the presence of HHF among basketball players and fans, who believe (erroneously) that a player who is having a streak of successes also has a high probability of scoring the next time. Camerer (1989) and Brown and Sauer (1993) also find evidence on HHF in the basketball betting market. Guryan and Kearney (2008) provide evidence for the "lucky store effect", which refers to the tendency of gamblers to believe that there is a higher probability of winning if they buy a ticket from a store which has recently sold a winning ticket.

Though GF and HHF make contradictory predictions, researchers find that both can be observed in the behaviour of the same individual. Sundali and Croson (2006) and Croson and Sundali (2005) identify both effects in the behaviour of the same single individual by using data

collected from a casino. Individuals tend to shun recently winning outcomes (GF), but bet more frequently after successful bets (HHF). Psychologists and economists offer several explanations for this phenomenon.

2.1. Psychological Theories

The psychology literature provides two explanations for the relationship between GF and HHF. One explanation suggests that both fallacies are cognitive biases caused by lack of knowledge of probability theory. The other explanation distinguishes two different types of expectations with respect to human performance and natural events (Ayton and Fischer, 2004). In human performance, individuals normally exhibit positive recency (HHF). In natural events, individuals normally exhibit negative recency (GF).

Tversky (1974) suggests that GF, one of the representativeness heuristic biases, is caused by poor understanding of probability theory. Individuals tend to expect a small sample to have the statistical characteristics of a large sample. Consequently, when individuals toss a fair coin they attach a probability of more than 50% of getting a tail after observing a streak of heads. On the other hand, when the streak of heads becomes very long or when the frequency of heads is consistently larger than that of tails, individuals may reject the belief that the series is random and start to believe that the outcome of any toss is positively related to the previous one. Then GF becomes HHF. Rao (2009) has designed an experiment to study the relation between GF and HHF. The experiment suggests that GF normally happens after short streaks and HHF happens after long streaks. The experiment also identifies the transition from GF to HHF. Guryan and Kearney (2008) argue that the “lucky store” effect challenges this interpretation, but in fact it seems consistent with HHF, since stores with long losing streaks are avoided.

Ayton and Fischer (2004) give an alternative explanation. They argue that the representativeness explanation cannot show why HHF is observed in basketball games but the GF is not, and why the GF is observed in lottery games but HHF is not. In their experiments individuals suffer from GF when predicting a series of natural events such as the toss of a coin or the spin of a roulette wheel, and suffer from HHF when predicting a series of human activities such as basketball players’ shooting success or whether the first serve of a tennis player is returned.

2.2. Economic Explanations

Economists are interested in GF and HHF because the two biases may help to explain puzzles or anomalies in financial markets. The model of Rabin and Vayanos (2010) (hereafter RV) follows the suggestion of the psychological literature and tries to link the two fallacies. The RV model assumes that agents believe that random shocks are negatively serially correlated (in other words, they suffer from GF). GF becomes HHF after a long streak because agents update their beliefs about the state of the world in an erroneous manner.

In the RV model, agents can observe a signal (x) which is generated by a mean-reverting process and an i.i.d shock: $x_t = \alpha x_{t-1} + u_t$. The parameter α is unknown to the agents. Agents erroneously believe that u_t is negatively autocorrelated, which generates a bias that is akin to GF. As agents learn from observing outcomes, it is assumed that they update their beliefs about α , but also (crucially) that they fail to correct their erroneous belief about the shocks. After a succession of positively correlated shocks, agents may revise their beliefs about α upwards so much that they believe that α is greater than one. When this happens their expectations become extrapolative, as if agents suffered from the HHF. In effect agents correct for their erroneous beliefs about shocks by overestimating α , and the prediction process is a combination of a GF bias and this compensating HHF bias. The model predicts GF behaviour for short streaks and HHF behaviour for long streaks.

Barberis et al. (1998) (hereafter BSV) assume that agents believe that random movements of an asset price have two regimes: reversing (which can be interpreted as GF) and trending (which can be interpreted as HHF). Agents continuously update their subjective probabilities (relative strength) of the two regimes. GF becomes HHF after a long streak, as in the RV model, because the relative strength of GF and HHF changes as the streak increases. As in the case of the RV model, traders have erroneous beliefs that they cannot entirely correct through learning. Conservatism means that individuals update their beliefs slowly when facing new information. Traders believe that the series, which is actually random, follows either one of two regimes with given parameters: mean-reverting or trending (continuation in the same direction). The prediction is a weighted average of the predictions of the two regimes, with the weights being adjusted in an optimal fashion in response to past data. Thus, as in the RV model, agents adjust their beliefs in the light of experience, but never arrive at the correct data-generating process. The model suggests that traders under-react to short term shocks and over-react to long term trends.

The mean-reverting regime can be considered as the GF effect and the trending regime can be explained as the HHF effect.

Both models can be criticised theoretically on the grounds that what agents can or cannot learn is somewhat arbitrary, but they are supported by experimental evidence when compared to either the random walk model or the Bayesian learning model (Bloomfield and Hales, 2002; Asparouhova et al., 2009). However, overall, the RV model does a better job than the BSV model.

3. Data and Series Construction

3.1. Data

In this section, we describe the dataset and the methodology we employed for generating the series which we used in our analysis.

The data include all the interbank tick-by-tick prices and order flows of nine currencies against the US dollar from May 1 to August 31 1996 on the Reuters Dealing 2000-1 (hereafter D2000-1). We report results for the Deutsche Mark (DEM) only; our findings for the Japanese yen were similar. The data set format is shown in Table 1. Trades on D2000-1 happen between two anonymous dealers: a calling dealer, who requires quotes, and a quoting dealer. The quoting dealer offers bid and ask prices to the calling dealer. The calling dealer has to make a quick decision to buy dollars (make a positive order flow) or sell dollars (make a negative order flow) or reject the quote. If a transaction is made, the time and the direction will be recorded by the system. Two things need to be mentioned. First, traders can only observe their own trading records. Second, though both bid and ask prices (two series of exchange rates) were quoted by the calling dealers, only the price that reflects the direction of actual trade is in the dataset (and the price may be slightly more favourable to the trader than the quote).

Although the Reuters D2000-1 system does not publish transaction prices, this does not affect our analysis. Traders in the market are aware of prices at all times for several reasons. First, traders in the market frequently ask for quotes to obtain the latest prices and to seek arbitrage opportunities. The no-arbitrage condition, on which the Evans and Lyons (2002) model relies, requires that quotes of different market makers are identical at a given time. Second, users of the Reuters D2000-1 system have indicative FXFX data (real time indicative prices) for reference. Although FXFX data are not transactions data, they do reflect the basic tendencies of prices.

Third, users of the Reuters D2000-2 system have data from Reuters D2000-1. Therefore, traders have many indirect methods of obtaining real-time price information, and it is reasonable to assume that transaction prices are known by everyone in the market. Moreover, if the trades recorded in our data set differ significantly from quoted prices at the time, one side of the trade is knowingly making a significant loss, which would be irrational, so one can assume that such trades are unlikely to occur.

The average time between trades is 26 seconds, with a standard deviation of 76 seconds, but the distribution is heavily skewed to the right, with a maximum of 10560 seconds (in a period when all the major financial markets are closed), and a minimum of zero. The median time between trades is nine seconds, and the interval exceeds one minute in only 9% of the sample.

3.2. Estimation of the Spread

Let x_t be the logarithm of the nominal exchange rate of the US dollar against the German mark for transaction t (an increase representing an appreciation of the US dollar), and let OF_t be the order flow which is equal to +1 when an agent buys dollars and equal to 0 when an agent sells dollars. Then

$$x_t = \begin{cases} A_t, & \text{Buy order } (OF_t = 1) \\ B_t, & \text{Sell order } (OF_t = 0) \end{cases} \quad (1)$$

where B_t denotes the logarithm of the bid price and A_t denotes the logarithm of the ask price and $A_t \geq B_t$. The unobserved bid-ask spread SP_t is the difference between the ask and the bid prices,

$$SP_t = A_t - B_t.$$

The spread distorts the trend in prices: if a buy order is followed by a sell order, or *vice versa*, the change in x will not reflect the true price trend. We therefore need to estimate and eliminate this effect.

We define the mid-price for transaction t , M_t , as the average of the bid and the ask prices. More formally,

$$M_t = \frac{1}{2}(A_t + B_t) = \begin{cases} A_t - \frac{SP_t}{2}, & (OF_t = 1) \\ B_t + \frac{SP_t}{2}, & (OF_t = 0) \end{cases}. \quad (2)$$

Under the supposition that the bid-ask spread is fixed, the change in the mid-price is given by:

Table 1: Data Format (USD/DEM)

Month	Day	Hour	Min	Sec	B/S	Ask	Bid
4	30	18	45	40	1	1.5326	-
4	30	18	46	23	0	-	1.5326
4	30	18	47	56	1	1.5328	-
4	30	18	48	22	1	1.533	-
4	30	18	49	53	1	1.5332	-
4	30	18	51	0	0	-	1.5327
4	30	18	52	34	1	1.5327	-
4	30	18	53	8	1	1.533	-
4	30	18	53	35	0	-	1.5329
4	30	18	54	21	1	1.5329	-
4	30	18	54	27	0	-	1.5333
4	30	18	55	10	0	-	1.533

This table shows the format of the data on the Reuters D2000-1 system. If B/S is equal to 1 there is a buy order and if B/S is equal to 0 there is a sell order.

$$\Delta M_t = \frac{1}{2}[(A_t - A_{t-1}) + (B_t - B_{t-1})] = A_t - A_{t-1} = B_t - B_{t-1}. \quad (3)$$

where Δ denotes the first-difference operator. Since in this case the spread is unknown, it has to be estimated from the data, which we do by using the model of Huang and Stoll (1997).¹ Using the estimate of the spread (\widetilde{SP}) we can adjust the observed changes in the transactions price returns using the following formula:

$$\Delta s_t = \Delta x_t - \widetilde{SP}_t(OF_t - OF_{t-1}), \quad (4)$$

where s is the estimate of the mid-price.

3.3. Streak Length and Width

If agents believe that a given series is autocorrelated, then the properties of past observations should have predictive power over agents' actions (in this case to buy or sell the US dollar). In our analysis we will focus on two properties of a streak, which is defined as a sequence of movements of the estimated mid-price of the dollar in the same direction over successive transactions.

¹This is a standard model when the direction of trade is known, as here. For details of the estimation see Bleaney and Li (2014).

The first property is the streak length, k_t , which denotes the number of transactions up to transaction t that the series continues to move in the same direction. The streak length is a non-negative integer.

We also need a measure of the strength of the streak, and a candidate would be the distance between the current exchange rate and the one at the beginning of the streak. However, clearly the streak length and the distance are positive correlated. In order to avoid this problem we use the streak width, w_t , defined as the streak distance divided by the streak length.

Figure 1 shows graphically the streak length, streak distance and streak width. Focusing on the upward trend section of the series, the horizontal right-angle side of the triangle represents the streak length, and the vertical right-angle side of the triangle represents the streak distance. The streak width is equal to $\tan\theta$, that is the slope of the hypotenuse of the triangle, and provides a measure of the strength of the shock. The steeper the hypotenuse is, the stronger the shocks are. The length of the segment AB is equal to the streak width of the trend before period t .

We can formally express the streak length as:

$$k_t = \begin{cases} k_{t-1}, & \text{sign}[\Delta s_t] = \text{sign}[\Delta s_{t-1}] \\ 0, & \text{sign}[\Delta s_t] \neq \text{sign}[\Delta s_{t-1}] \end{cases} \quad (5)$$

where

$$\text{sign}_t = \begin{cases} 1, & \Delta s_t > 0 \\ \text{sign}_{t-1}, & \text{if } \Delta s_t = 0 \\ -1 & \Delta s_t < 0 \end{cases}$$

The formal expression for the streak width is given by:

$$w_t = \frac{100|s_t - s_{t-k_t-1}|}{k_t+1}. \quad (6)$$

Figure 1: Streak Length, Streak Distance and Streak Width



If traders suffer from GF and/or HHF, their choices should be influenced by the streak length and the streak width. For example, if the dollar has been rising for some time (a long streak) or particularly fast (a wide streak), the HHF would predict that the next transaction will tend to be a dollar purchase. Most of the evidence relating to these fallacies has been gathered from coin-tossing experiments where, by definition, the streak width is constant. Consequently, the literature has focused on the effects of streak length on the behaviour of agents and has ignored streak width. That is fine when the evidence is related to scoring in a basketball game or gambling in a numbers game, but is not necessarily appropriate for the foreign exchange market where the shocks vary in size.

Table 2 presents the summary statistics of the data. The streak length (k) and streak width (w) are calculated from the series of estimated mid-price exchange rate returns (Δs).

Table 2: Summary Statistics

	Obs.	Mean	Std. Dev.	Min.	Max
Δs	256838	-1.12×10^{-7}	3.85×10^{-4}	-0.00663	0.00271
Streak length (k)	256838	0.700	1.0437	0	19
Streak width (w)	256838	0.0235	0.0292	2.98×10^{-6}	0.663

Source: USD/DEM transaction data from the Reuters D2000-1 system.

Figure 2 shows the histogram of streak length. The average streak length is less than one, and the fraction of instances where $k_t = 0$ is greater than 50%, which suggests that the sign of exchange rate returns changes very frequently. Figure 3 shows the histogram of streak width, which is also strongly skewed to the right.

4. Streak Characteristics and Trading Behaviour


In this section we report some empirical tests of whether the characteristics of a streak are correlated with trading decisions in a way that is consistent with the literature on psychological fallacies. Specifically, we test whether streak length and width are correlated with order flows as the GF or HHF would predict.

We wish to test whether trading behaviour is consistent with the GF, with the HHF or with neither. To do this we use a modified indicator of trade direction:

$$TF_t = \begin{cases} 1 & \text{if } \frac{2(OF_t - 1/2)}{sign_t} = 1 \\ 0 & \text{if } \frac{2(OF_t - 1/2)}{sign_t} = -1 \end{cases}$$

When $TF_t = 1$ the trader trades with the trend (buys the \$ when it has just risen or sells the \$ when it has just fallen) and when $TF_t = 0$ the trader trades against the trend. Formally, $TF = 1$ if $OF = 1$ and $sign = +1$ or if $OF = 0$ and $sign = -1$. Conversely, $TF = 0$ if $OF = 1$ and $sign = -1$ or if $OF = 0$ and $sign = +1$. The HHF would tend to predict $TF = 1$, particularly in the case of a long streak, whereas the GF would predict $TF = 0$.

Figure 2: Histogram of Streak Length (USD/DEM pooled)




Source: The histogram is obtained from (7) using USD/DEM transaction data on the Reuters D2000-1 system.

We can use the following regression to study the impact of the streak length on trend-following behaviour:

$$TF_t = \Phi(\beta_1 + \beta_2 k_t) + \varepsilon_t. \quad (7)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. If the coefficient β_2 is positive, then the probability of trend-following behaviour increases as the streak length becomes longer (as the HHF would predict) and vice versa. A negative β_2 would be consistent with the GF.

Figure 3: Histogram of Streak Width (USD/DEM pooled)



Source: The histogram is obtained from (8) using USD/DEM transaction data on the Reuters D2000-1 system.

The first column (7) of Table 3 reports the results, while the second column (7*) reports the average marginal effects. The coefficient β_2 is significant and negative, which indicates that trend-following becomes less probable as the streak length increases. The marginal effects suggest that for each increase of one in the streak length, the probability of following the trend on average decreases by 13.7%. These results are consistent with the hypothesis that trading decisions reflect the GF.

We can examine the relationship between streak length and trading decisions more closely by replacing streak length by a series of dummies (Dk) for different streak lengths such that

$$Dk_i = \begin{cases} 1 & \text{if } k_t = i \\ 0 & \text{if } k_t \neq i \end{cases}$$

The probit model is written as:

$$TF_t = \Phi(\beta_1 + \sum_{i=1}^n \delta_i Dk_{it}) + \varepsilon_t. \quad (8)$$

Compared with (7), (8) uses dummies for the streak length rather than the streak length itself, which allows the influence of the streak length on order flows to be non-linear. This non-linear relationship will be captured by the coefficients of the dummies. Since the proportion of observations with streak length over five is rather small, we amalgamate all the dummies for a streak length of five or more into one dummy for a streak length of five or more (Dk_{5+}).

The third and fourth columns of Table 3 report the results of estimating equation (8) and the marginal effects (denoted by 8*) respectively. The results are similar to those for equation (7): as the streak length increases, the probability of following the trend decreases, and this occurs at each step, which indicates that equation (7) is a good approximation to the underlying pattern. For example, from (8*), compared with the case when the streak length is zero, the probability of following the trend is 24.2% lower when the streak length is one and 31.9% lower when the streak length is two. Furthermore, the decline in the probability of trend-following falls as the streak length increases: when the streak length increases from one to two, the probability of trend-following falls by 7.7%, and when the streak length increase from two to three, the probability of trend-following falls by 5.8%. Figure 4 plots the predicted probability of trend following against the streak length implied by equation (8) in Table 3.

Trend-following behaviour when the streak length is very short (less than zero or equal to one) might arise from the sequential nature of order arrivals. Traders holding the same expectations with regard to exchange rate movements may place orders in the same direction almost simultaneously, which causes two adjacent orders to be positively correlated. For example, suppose that the streak length is zero and the exchange rate rises in period t , and that both traders, A and B, believe that the exchange rate will rise in the future. The buy order of trader A comes first. The buy order of trader B comes after A's order, when (we shall assume) the exchange rate has started to rise. As a result, B's order is recorded as a trend-following order with a streak length of one. Thus, almost simultaneous actions by traders reacting to similar information might be responsible for some apparent trend-following behaviour. This effect will tend to work against the GF.

Table 3: Trend-following behaviour and streak length

	(7)	(7*)	(8)	(8*)
k_t	-0.343*** (-105.10)	-0.137*** (-105.10)		
Dk_1			-0.605*** (-101.93)	-0.242*** (-101.93)
Dk_2			-0.800*** (-91.42)	-0.319*** (-91.42)
Dk_3			-0.945*** (-67.05)	-0.377*** (-67.05)
Dk_4			-1.054*** (-45.23)	-0.420*** (-45.23)
Dk_{5+}			-1.148*** (-37.15)	-0.458*** (-37.15)
Constant	0.171*** (54.47)		0.242*** (73.03)	
Pseudo-R²	0.050		0.059	
Probability		0.473		0.474
Obs.	256838		256838	

Dependent variable: TF_t . Estimation method: probit.
Columns numbered * are marginal effects. k = streak length.
 $Dk_i=1$ for $k=i$, =0 otherwise. $Dk_{5+}=1$ for $k \geq 5$, =0 otherwise.
Z-statistics in parentheses: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$


4.1. Streak Width

In this section, we investigate whether the strength as well as the direction of the signal appears to make a difference to trading decisions. We do this by adding the streak width to equation (7):

$$TF_t = \Phi(\beta_1 + \beta_2 w_t + \sum_{i=1}^5 \delta_i Dk_{it}) + \varepsilon_t, \quad (9)$$

where β_2 measures the effect of the streak width on trend-following behaviour.

Figure 4: Predicted probability from equation (8) in Table 3



Source: The graph is obtained from (8) using USD/DEM transaction data on the Reuters D2000-1 system

The results and the marginal effects are reported in Table 4. The streak-width coefficient is negative and significantly different from zero, which suggests that the streak width has an additional negative effect on trend-following behaviour, for a given streak length. On the other hand its Z-statistic is quite small, considering the sample size, and the estimated effect of a one-standard-deviation increase in streak width on the probability of trend following is only -1.4% ($= -0.479 \times 0.0292$). Thus streak width is of minor importance compared to streak length.

Table 4: Trend-following Behaviour, Streak Length and Streak Width

	(9)	(9*)
Dk_1	-0.617*** (-102.66)	-0.237*** (-109.06)
Dk_2	-0.816*** (-92.36)	-0.300*** (-110.27)
Dk_3	-0.964*** (-68.05)	-0.329*** (-92.38)
Dk_4	-1.075*** (-46.04)	-0.350*** (-69.59)
Dk_{5+}	-1.171*** (-37.82)	-0.368*** (-62.38)
<i>Streak width</i>	-1.204*** (-13.68)	-0.479*** (-13.68)
Constant	0.277*** (66.43)	
Pseudo-R²	0.0583	
Probability		0.474
Obs.	256838	

Dependent variable: TF_t . Estimation method: probit.
Columns numbered * are marginal effects. k = streak length.
 $Dk_i=1$ for $k=i$, =0 otherwise. $Dk_{5+}=1$ for $k \geq 5$, =0 otherwise.
Z-statistics in parentheses: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5. Robustness checks

5.1. Time intervals

One of the important properties of the data is that time is measured in trades, which means that the time interval between consecutive observations is not identical. To some extent this is an advantage, because periods when there is little trading, which may be atypical, are represented by few observations and should therefore have less influence on the results. Nevertheless it is important to check that our results are not sensitive to the time interval between trades.

Let T_t be the number of seconds between orders t and $t-1$. For example, the first trade in Table 1 occurred at 18:45:40, and the second trade occurred at 18:46:23. There are thus 43 seconds between the two trades, and accordingly $T_2 = 43$. As a first robustness test, we allow all the coefficients in equation (9) to vary with T . Thus we estimate the following probit regression:

$$TF_t = \Phi(\beta_1 + \beta_2 w_t + \sum_{i=1}^5 \delta_i Dk_{it} + \beta_3 T_t + \sum_{i=1}^5 \gamma_i T_i Dk_{it}) + \varepsilon_t. \quad (10)$$

As a second robustness test, we exclude the data with the shortest and the longest intervals between trades. To be specific, we re-estimate equation (9) using only the trades between the 25th and the 75th percentiles for T (respectively three seconds and 23 seconds). Table 5 reports the results of these two tests and the corresponding marginal effects, respectively.

In the first test, (columns 10a and 10a*) the dummy variable interacted with the time interval is significant at the 5% level in only one out of the five cases ($k = 1$). Thus the estimated coefficients do not seem to vary much with the time interval between trades. In the second case (the truncated sample – columns 10b and 10b*), the estimated marginal effects are very similar to those shown in Table 4.

5.2. Microstructure

According to Evans and Lyons (2002), there is “hot-potato” trading on the Reuters D2000-1 system; traders pass their inventory imbalance onto others, and thus current trading behaviour is influenced by past trading behaviour. To allow for this, we add lagged trade direction to the model:

$$TF_t = \Phi(\beta_1 + \beta_2 w_t + \sum_{i=1}^5 \delta_i Dk_{it} + \beta_4 TF_{t-1}) + \varepsilon_t \quad (11)$$

Table 6 reports the results and the corresponding marginal effects. The coefficient of TF_{t-1} is significant and positive, as the hypothesis of hot-potato trading would predict. More important for our purposes is the fact that including a lag of the dependent variable produces estimates of the other coefficients that are similar to those shown in Table 4, so that our previous results seem to be robust to controlling for hot-potato trading.

6. Do Traders Behave Rationally?

In the last two sections, we have found strong evidence that traders tend to trade against the trend. This behaviour is consistent with the influence of the gambler’s fallacy: traders observe that a currency is moving in one direction, and they tend to trade as if they were betting on that movement being reversed. This may, however, be entirely rational behaviour, and not the result of psychological fallacies at all, if this is a pattern that is observable in past data. To test this, we need to examine the dynamics of the exchange rate.

Table 5: Time Interval Effects

	(10a)	(10a*)	(10b)	(10b*)
Dk_1	-0.617*** (-102.67)	-0.237*** (-109.08)	-0.666*** (-80.68)	-0.255*** (-86.37)
Dk_2	-0.816*** (-92.35)	-0.295*** (-110.26)	-0.863*** (-71.10)	-0.309*** (-86.39)
Dk_3	-0.964*** (-68.03)	-0.329*** (-92.35)	-1.013*** (-52.51)	-0.341*** (-73.29)
Dk_4	-1.076*** (-46.03)	-0.350*** (-69.60)	-1.107*** (-35.20)	-0.357*** (-54.40)
Dk_{5+}	-1.166*** (-37.35)	-0.367*** (-61.35)	-1.224*** (-29.39)	-0.378*** (-50.64)
w_t	-1.231*** (-13.91)	-0.490*** (-13.91)	-1.351*** (-10.73)	-0.538*** (-10.73)
T_t	0.000041 (0.92)	0.000016 (0.92)	0.00061 (1.00)	0.00025 (1.00)
$Dk_1 * T$	0.00024*** (2.76)	0.000094*** (2.76)		
$Dk_2 * T$	0.00025* (1.93)	0.000098* (1.93)		
$Dk_3 * T$	0.00028 (1.34)	0.00111 (1.34)		
$Dk_4 * T$	-0.00005 (-0.26)	-0.000023 (-0.26)		
$Dk_{5+} * T$	0.00107* (1.66)	0.000425* (1.66)		
TF_{t-1}				
Constant	0.276*** (64.60)		0.274*** (64.70)	
Pseudo- R^2	0.0584		0.0659	
Probability		0.474		0.475
Obs.	256838	256838	136118	136118

Dependent variable: TF_t . Estimation method: probit.
Columns numbered * are marginal effects. k = streak length.
 $Dk_i=1$ for $k=i$, =0 otherwise. $Dk_{5+}=1$ for $k \geq 5$, =0 otherwise.
 w = streak width. Regression (12a) uses observations
between the 25th and 75th percentiles of T only.
Z-statistics in parentheses: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Microstructure Effects

	(11)	(11*)
Dk_1	-0.670*** (-107.14)	-0.256*** (-114.86)
Dk_2	-0.833*** (-94.78)	-0.300*** (-113.85)
Dk_3	-0.968*** (-68.86)	-0.329*** (-93.68)
Dk_4	-1.070*** (-46.23)	-0.349*** (-69.61)
Dk_{5+}	-1.155*** (-37.64)	-0.365*** (-61.27)
w_t	-1.282*** (-14.49)	-0.510*** (-14.49)
TF_{t-1}	0.174*** (32.86)	0.0692*** (32.93)
Constant	0.210*** (45.48)	
Pseudo-R^2	0.0613	
Probability		0.474
Obs.	256822	256822

Dependent variable: TF_t . Estimation method: probit.
Columns numbered * are marginal effects. k = streak length.
 $Dk_i=1$ for $k=i$, =0 otherwise. $Dk_{5+}=1$ for $k \geq 5$, =0 otherwise.
 w = streak width.
Z-statistics in parentheses: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We begin by estimating the relationship between short-run exchange rate movements (i.e. the movement in the price between two successive trades) and the streak variables previously considered. If traders are rational, the streak variables should have the same signs in the exchange rate regression as in the previous regression for trade direction. The rational explanation for our previous results would then be that traders bought (sold) the dollar at time t because they predicted from past data that the dollar was likely to rise (fall) between times t and $t+1$. Thus we estimate an equation of the form:

$$\Delta s_{t+1} = \Phi(\beta_1 + \beta_2 w_t + \sum_{i=1}^5 \delta_i Dk_{it}) + \varepsilon_t, \quad (12)$$

The snag is that, even if the results of estimating (12) are consistent with the rational explanation, the finding could be driven by reverse causality. This is because a purchase of the dollar by a customer tends to cause market-makers to adjust the price upward, partly to re-equilibrate their inventories, and partly to guard against the possibility that the equilibrium price has shifted permanently upwards (Huang and Stoll, 1997). This effect means that, whatever the true reason for trading decisions at time t , the impact on the price will make it look as though the decisions were driven by correct anticipation of price movements. In this sample, a regression of Δs_{t+1} on OF_t yields the following result (z -statistics in parentheses):

$$1000\Delta s_{t+1} = -0.000133 + 0.0340OF_t$$

(-0.18)
(44.9)

To get around this difficulty, we also examine whether streak length helps to predict price movements over a longer period starting at time $t+1$. To be precise, we look at the correlation between streak length and width and the movement in the exchange rate between trades $t+1$ and $t+6$. If streak variables predict exchange rate movements at this horizon with the same sign as in Table 4, that would be strong evidence of rational exploitation of observable patterns in the data.

An alternative approach is to include the predictions from an autoregressive model of exchange rate returns in our model of trading behaviour; if rational exploitation of past patterns in the data is at play, this variable should have a significant positive coefficient. If the streak variables then become insignificant, that would suggest that trading behaviour is driven by rational profit-seeking rather than by psychological fallacies. If the streak variables remain significant, it would seem that psychological fallacies are also playing a role in trading behaviour.

Let the variable *prediction* be the forecast of a time series model of exchange rate returns divided by the sign of the current return.² Thus *prediction* measures the size of future returns and is positive if the model suggests that the trend continues in the future. Letting $E(\Delta s_{t+1}|\Delta s_t, \Delta s_{t-1}, \dots)$ be the prediction from an AR(15) time series model, the variable *prediction* is formally given by:

$$prediction = \frac{E(\Delta s_{t+1}|\Delta s_t, \Delta s_{t-1}, \dots)}{sign_t} \times 100,$$

² Division by the sign of the current return is necessary because of the specification of the dependent variable.

The probit model is given by:

$$TF_t = \Phi(\beta_1 + \beta_2 w_t + \sum_{i=1}^5 \delta_i Dk_{it} + \beta_5 prediction) + \varepsilon_t \quad (13)$$

The first column of Table 7 shows that streak length and width do help to predict the next movement of the exchange rate (the R-squared is 0.183). The streak length dummies and streak width all have the same negative signs as in our trade direction regression (Table 4), which supports the rational model. On the other hand, the second column of Table 7 demonstrates that they have virtually no capacity to explain exchange rate movement over the subsequent five trades ($t+1$ to $t+6$); although all the coefficients are all negative and some of them are statistically significant, t -statistics below ten have virtually no economic significance in such a large sample, as is indicated by the R-squared of 0.000. Thus it is far from clear that traders should rationally use information about streak length and width to determine their trades.

Table 8 includes the predictions from an AR(15) model together with streak variables in a probit model of trade direction. The predictions from the autoregressive model are significant with the expected positive sign, which suggests that trades are partly explained by rational exploitation of past patterns in the data. Nevertheless the streak variables remain significantly negative with coefficients only slightly smaller in absolute value than in Table 4. This suggests that, although there is a rational element to trading, psychological fallacies also play a part.

7. Conclusion

There is evidence from previous research that people suffer from the gambler's fallacy and/or the hot hand fallacy when they predict future values of uncertain series. However, the evidence comes from laboratory and field experiments from casinos and lotteries where the underlying series is the product of a random process. Since the outcome is then entirely random, there is no cost attached to succumbing to these fallacies: the player's chances of winning remain the same. A natural question to ask is whether people learn to overcome their tendency to succumb to psychological fallacies in environments where mistakes are punished. To this end, we have studied the behaviour of professional traders in the foreign exchange market. Moreover, many studies focus only on the direction of the signal with no variation in signal strength; in financial markets, as in many other real-world situations, the strength of the signal (the speed of price movements) is potentially as important as its direction.

Table 7: Streaks and Exchange Rate Returns

Dependent variable:	$1000\Delta s_{t+1}$	$1000(s_{t+6} - s_{t+1})$
Dk_1 * sign	-0.045*** (-32.4)	-0.00544*** (-3.29)
Dk_2 * sign	-0.077*** (-35.4)	-0.00835*** (-3.36)
Dk_3 * sign	-0.091*** (-25.2)	-0.00677 (-1.64)
Dk_4 * sign	-0.102*** (-17.0)	-0.00509 (-0.77)
Dk_{5+} * sign	-0.104*** (-13.9)	-0.00317 (-0.37)
w_t * sign	-3.97*** (-127.5)	-0.108*** (-4.08)
Constant	-0.00015 (-0.21)	-0.00061 (-0.74)
R^2	0.183	0.000
RMSE	0.350	0.410
Obs.	256822	256742

The variable “sign” =1 if $\Delta s_t > 0$ and = -1 if $\Delta s_t < 0$.
Z-statistics in parentheses: * p<0.05, ** p<0.01, *** p<0.001

We have shown that the actions of traders depend on the two properties of the trend in the series of price returns, namely the streak length and the streak width, in a manner that is consistent with the gambler’s fallacy (but not the hot hand fallacy). Persistent (and to a lesser extent fast) upward (downward) movements in the dollar tend to be followed by dollar sales (purchases). We then asked whether our results could be explained by traders’ rational exploitation of non-random patterns in exchange rate returns, by examining whether (a) streak variables predicted exchange rate returns; and (b) by including forecasts from an autoregressive model of exchange rate returns in our model of trading. It was not clear that streak variables predict exchange rate movements over anything longer than the shortest possible horizon (where there are issues of reverse causality). We found that the predictions from an autoregressive model of exchange rate returns help to explain trading behaviour, but that the streak variables retain their statistical significance, which suggests that psychological factors also play a role.

Table 8: Trading Strategies

	(14)	(14*)
Dk_1	-0.534*** (-82.26)	-0.207*** (-86.21)
Dk_2	-0.717*** (-76.79)	-0.264*** (-88.35)
Dk_3	-0.857*** (-58.90)	-0.300*** (-75.25)
Dk_4	-0.964*** (-40.80)	-0.325*** (-57.00)
Dk_{5+}	-1.056*** (-33.80)	-0.345*** (-50.87)
w_t	-3.007*** (-28.44)	-1.197*** (-28.44)
Prediction	5.318*** (33.34)	2.117*** (33.34)
Constant	0.242*** (56.48)	
Pseudo-R^2	0.0618	
Probability		0.474
Obs.	256822	256822

Dependent variable: TF_t . Estimation method: probit.
Columns numbered * are marginal effects. k = streak length.
 $Dk_i=1$ for $k=i$, =0 otherwise. $Dk_{5+}=1$ for $k \geq 5$, =0 otherwise.
 w = streak width.

Z-statistics in parentheses: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

- Asparouhova, E. N., M. G. Hertzel, and M. L. Lemmon (2009). Inference from streaks of random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers. *Management Science* 55, 1776-1782
- Ayton, P. and I. Fischer (2004). The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory and Cognition* 32, 1369–78.
- Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of Financial Economics* 49, 307–343
- Bleaney, M. and Z. Li (2014). Decomposing the bid-ask spread in multi-dealer markets. *University of Nottingham School of Economics Discussion Papers* 14/03.
- Brown, W. O. and R. D. Sauer (1993). Does the basketball market believe in the hot hand? comment. *American Economic Review* 83, 1377–1386
- Camerer, C. F. (1989). Does the basketball market believe in the 'hot hand?' *American Economic Review* 79, 1257–1261
- Clotfelter, C. T. and P. J. Cook (1993). The “gambler's fallacy” in lottery play. *Management Science* 39, 1521–1525
- Croson, R. and J. Sundali (2005). The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty* 30, 195–209
- Evans, M. D. D. and R. K. Lyons (2002). Order flow and exchange rate dynamics. *Journal of Political Economy* 110, 170–180
- Gilovich, T., R. Vallone, and A. Tversky (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology* 17, 295–314
- Guryan, J. and M. S. Kearney (2008). Gambling at lucky stores: Empirical evidence from state lottery sales. *American Economic Review* 98, 458–473
- Huang, R. D. and H. R. Stoll (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies* 10, 995–1034

Rabin, M. and D. Vayanos (2010). The gambler's and hot-hand fallacies: Theory and applications. *Review of Economic Studies* 77, 730–77

Rao, J. M. (2009). Inference of Subjective Sequences: When the gambler's fallacy becomes the hot hand fallacy. Yahoo! Research Labs, Mimeo

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

Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases." *Science* 185(4157): 1124-1131.