What Drives Peer Effects in Financial Decision-Making? Neural and Behavioral Evidence

CARY FRYDMAN *

October 2014

ABSTRACT

We use neural data collected from an experimental asset market to test the underlying mechanisms that generate peer effects. Despite the fact that subjects are randomly assigned to pairs and endowed with the same information set, we find substantial evidence of peer effects in investment decisions. We then use the neural data to construct empirical tests that can separately identify a social learning mechanism and a relative wealth concern mechanism. We observe neural activity that indicates both mechanisms contribute to peer effects. The neural activity that we use in the empirical tests is measured at the time when subjects receive information about a peer's investment decision. Our results therefore demonstrate how neural data generated in the absence of choice can be used to distinguish between competing theories of financial decision-making.

*contact: cfrydman@marshall.usc.edu, Marshall School of Business, University of Southern California.
Financial support from a Zumberge Individual Award Grant is gratefully acknowledged. I thank Lawrence Jin, Alec Smith, David Solomon and seminar participants from USC for helpful comments. I also thank Read Montague and Terry Lohrenz for sharing their neuroimaging data set.

Most models of trading in financial markets typically do not allow investors to socially interact with each other. Indeed, the primary mechanism in most models through which one investor can affect the beliefs and decisions of another investor is through the market price. In reality, however, individuals observe each others' behavior directly or learn about each others' decisions and beliefs through conversation (Hirshleifer and Teoh (2009)). Over the last fifteen years, empiricists have shown that these social interactions can have a significant impact on financial decision-making in a wide variety of contexts. For example, peers can affect stock market participation (Hong et al. (2004); Brown et al. (2008)), retirement saving decisions (Duflo and Saez (2003); Beshears et al. (2014)), mutual fund manager stock selection (Hong et al. (2005); Pool et al. (2014)), and individual investor trading decisions (Ivković and Weisbenner (2007); Bursztyn et al. (2014)).

Despite the growing body of evidence that documents peer effects in financial decision-making, our understanding of the mechanism that generates peer effects is still limited. Broadly speaking, there are two competing channels through which peers can affect decisions. First, an agent can learn about the underlying fundamentals of a financial asset from a peer's decision (the information channel). Second, a peer's decision may directly enter another agent's utility function because of relative wealth concerns or a direct preference for conformity (the preference channel). Distinguishing between these two channels can be challenging because in many environments, an information-based mechanism will make similar predictions to a preference-based mechanism[1]. However, separating the two mechanisms is important because it can provide useful guidance for theorists as they build newer models of financial markets with social interactions, a new subfield that Hirshleifer (2014) calls "social finance".

[1] For example, suppose an investor observes his neighbor experience a large positive return in the stock market, and subsequently increases his investment in the stock market (Kaustia and Knüpfer (2012)). This can occur through a social information channel if the investor learns about underlying fundamentals from his neighbor. However, it could also occur through a preference channel if the investor increases his investment in order to "Keep up with the Jones."

In this paper, we take up the challenge of identifying the underlying mechanism of peer effects using a newly available source of data: neural activity measured with functional magnetic resonance imaging (fMRI) in a controlled experimental setting. While almost all empirical studies of peer effects in finance have used data from the field, we use a laboratory experiment with randomly assigned peers; this setting gives us complete control over the information that is provided to subjects, and even more importantly, over the information that is transmitted between subjects. We can therefore rule out, by construction, the well-known endogeneity problem where similar behavior among two agents may be driven by common preferences or exposure to common information (Manski (1993)). We then use the neural data to construct novel tests of the competing mechanisms that generate peer effects. These tests rely on neural data collected in the absence of choice[2], and hence showcase the marginal value of the neural data in testing between competing theories of financial decisions.

In our experiment, subjects are presented with a dynamic investment problem where they are asked, in each of two hundred periods, to allocate a fraction of their wealth between a risk-free asset and a risky asset. Among the various studies of peer effects in finance, our setting is most closely related to the study conducted by Ivković and Weisbenner (2007), who examine time-series variation in stock purchases among individual investors within the same neighborhood. These authors find that an individual investor is more likely to purchase a stock from industry *I* if his neighbor also purchased a stock from industry *I* in the same quarter. The authors propose that neighbors directly communicate with each other, and this communication has a causal effect on subsequent investment decisions. While it is difficult to test this conjecture using field data, our neuroimaging experiment enables us to perfectly observe two key variables: (i) the information transmitted between peers and (ii) the neural activity that is generated upon

[2] By "absence of choice" we mean that the neural data used in our tests is collected at a particular point in the experiment when subjects are not actively deliberating about a decision, but instead are simply receiving information.

receipt of this information transmission. Together, these two complementary pieces of data allow us to construct tests of the underlying peer effects mechanism.

We describe the details of the neural predictions that are generated from each competing mechanism later in the paper, but we summarize the main idea here. All of our neural tests rely on the theory of prediction error, which is a signal that measures the change in expected net present value of lifetime utility generated by new information. Critically, a large body of evidence from cognitive neuroscience has shown that prediction errors can be measured in a specific area of the brain called the ventral Striatum (vSt)[3]. For economists, the potential value of this measurement technique should be clear as it implies that we can infer which factors – in addition to the standard consumption factor – affect a subject's utility[4]. To begin the analysis, we develop precise predictions about neural activity that is generated under each of the two peer effects mechanisms.

First, we consider a social learning mechanism where a subject learns about a risky asset's value from his peer's decision. In our experiment, this decision is simply the fraction of wealth allocated to the risky asset in each period. When a subject observes his peer increase his allocation to the risky asset, the social learning mechanism predicts that the subject should update upwards his valuation of the asset. This new information implies a better investment opportunity set than previously expected, and will therefore increase a subject's net expected utility. This will, in turn, produce a positive prediction error in the vSt. Conversely, when a subject observes his peer decrease his allocation to the risky asset, a subject will interpret this information as bad news about future investment opportunities, and this will generate a negative prediction error in the vSt.

Second, we examine a relative wealth mechanism where a subject has utility over the difference between his final portfolio value and his peer's final portfolio value. Because all

[3] See, for example, Schultz et al. (1997), Pessiglione et al. (2006), and Caplin et al. (2010).
[4] Indeed, a series of recent papers in neuroeconomics has highlighted the vast potential that this technique holds for answering fundamental questions about choice theory and belief updating (Caplin and Dean (2008); Caplin et al. (2010); Rutledge et al. (2010)).

subjects invest each period in the same risky asset, the revelation of a peer's investment allocation carries information about a subject's per-period change in his relative portfolio value. Under the relative wealth mechanism, these per-period changes in relative portfolio value will impact net expected utility, and hence should generate prediction errors. Specifically, we expect to see a positive prediction error when a subject invests more than his peer and the market experiences a positive return. Conversely, we expect to see a negative prediction error when a subject invests less than his peer and the market experiences a positive return.

Empirically testing for relative wealth concerns can be difficult because in many environments, changes in absolute wealth are correlated with changes in relative wealth. However, a key feature of the current experimental design is that information about the market return is revealed before information about the peer decision is revealed. As we explain further in the main text, this staggered information arrival allows us to construct neural tests that can separately identify whether a change in utility is driven by absolute changes in wealth or relative changes in wealth.

Overall, our results are consistent with previous field studies that have documented peer effects in financial decision-making. Behaviorally, we find that risky asset allocations are causally affected by peer decisions, and part of the remaining variation in risky asset allocations can be explained by the conditional Sharpe ratio. In the neural data, we find two main results. First, at the time when a market return is revealed, activity in the vSt positively correlates with absolute changes in wealth that are generated by market returns. Second, at the time when a peer decision is revealed, vSt activity correlates with *both* the prediction errors that are generated under the social learning mechanism and the relative wealth concern mechanism. We interpret the neural data as evidence in favor of both a social learning mechanism and a relative wealth mechanism. It is worth emphasizing that the neural data used to test these two channels is generated at a specific point in the experiment when a subject is not making a decision, but is instead receiving information about a peer's decision. Our tests therefore demonstrate, for the

first time, how neural data generated in the absence of choice can be used to test competing theories of financial decision-making[5].

The results from our study mainly contribute to two literatures: one on peer effects and one on the growing field of neurofinance. First, as recent empirical work on peer effects has achieved growing success in constructing clean identification strategies, the natural next step is to ask, what drives these peer effects? There has been little work on this question, with the exception of a recent paper that uses a field study in Brazil (Burszstyn et al. (2014)). These authors use a clever experimental design to separate an investor's revealed preference for a risky asset from the ownership of the risky asset. It is then possible to infer whether investors are purchasing an asset due to social learning (information) or social utility (preferences). Burszstyn et al. (2014) find that both mechanisms explain a portion of the observed peer effects, which is consistent with the results from our neural tests. However, while their paper relies upon a single binary investment decision, we use a dynamic investment task where subjects allocate a continuous fraction of their wealth to a risky asset. Moreover, our setting differs in that we provide subjects with the outcomes of the risky asset return, and critically, this enables us to directly test the relative wealth mechanism using neural data.

Second, we contribute to the young but growing literature on neurofinance, which uses tools from cognitive neuroscience to help understand the neural mechanisms that give rise to financial decisions. Many of the early contributions to this field provide important evidence showing the neural and physiological correlates of financial risk-taking (Lo and Repin (2002); Kuhnen and Knutson (2005); Preuschoff et al. (2006)). Building on this earlier work, subsequent studies have demonstrated how exogenous emotional cues or emotional regulation strategies can causally impact financial risk-taking (Knutson et al. (2008); Sokol-Hessner et al. (2009)).

[5] Recent work has examined how neural data collected in the absence of choice can help predict choice in other settings such as consumer purchasing decisions (Levy et al. (2011); Smith et al. (2014)).

While recent work in neurofinance is starting to investigate a wider array of topics,[6] the main contribution of our paper to this area is to show how neural data, in the absence of choice, can be used to test between competing theories of financial decision-making. As such, our paper is related to two recent studies that use neural data to test theories of investor behavior. One uses neural activity generated from selling stocks to test the "realization utility" theory of trading (Frydman et al. (2014a)); the other uses vSt activity generated upon viewing stock returns to test whether regret theory can explain stock purchasing behavior (Frydman et al. (2014b)).

The rest of this paper is organized as follows. Section I presents the experimental design and the neural predictions from the two peer effects mechanisms. Section II provides a primer on fMRI methods for economists. Section III describes the behavioral and neural results. Section IV discusses the results and concludes.

I. Experimental Design and Predictions

In this section, we first describe the experimental stock market that was used to generate both the behavioral and neural data. We then outline the neural predictions that are generated under the potential mechanisms that can give rise to peer effects.

A. Design

The design and data are taken from the experiment conducted in Lohrenz et al. (2013) and we now describe their experimental setting in detail.

48 subjects were endowed with $100 of experimental cash and given the opportunity to invest this wealth in two separate assets over the course of two hundred trials. One asset was risk-

6 Other recent work in neurofinance examines how investors learn from financial information (Bossaerts and Payzan-Le Nestour (2014); Kuhnen (2014)) and the formation of experimental asset price bubbles (De Martino et al. (2013); Smith et al. (2014)).

free and paid a zero interest rate, and the other asset was risky and generated a random return R_t in each period $t=1,2,\ldots,200$. The sequence of risky asset returns $\{R_t\}$ is taken from historical markets, and subjects are not told anything about the process $\{R_t\}$, except that it is taken from real historical markets. As such, subjects are faced with a large amount of parameter uncertainty upon entering the experiment. At the beginning of each period t, subject i allocates a fraction of his wealth, $x_{i,t}$, to the risky asset. The remaining fraction of wealth, $(1 - x_{i,t})$, is invested in the risk-free asset that earns a zero interest rate.

A sequence $\{R_t\}$ of length twenty is referred to as a "market", and each subject participates in ten separate markets throughout the course of the experiment. Subjects are instructed that each market is mutually independent of the other nine markets and are given a short break in between each market; at the beginning of a new market, a subject's wealth is carried over from the end of the previous market. Upon entering the experiment, half of the 48 subjects are randomly assigned to a *social treatment,* and the remaining 24 subjects are assigned to a *control condition.* Subjects within the *social treatment* are then randomly assigned into twelve pairs.

Each of the two hundred trials consists of three different screens (Figure 1). First, a subject sees an "allocation" screen at which time he is instructed to enter his investment allocation $x_{i,t}$. Second, a "market" screen displays both the realized return of the risky asset and the subject's updated portfolio value as a function of his investment allocation. Finally, if the subject is in the *social treatment,* a "peer decision" screen reveals the investment allocation of subject *i's* peer. If instead the subject is the *control condition,* the "peer decision" screen reveals a randomly drawn investment allocation that is uniformly drawn from [0,1]. Subjects in the *control condition* are explicitly told that the investment allocation they observe is indeed uniformly drawn from [0, 1]. At the end of the experiment, a subject's final portfolio was converted from experimental currency to actual US dollars using a 5:1 exchange rate. In addition to the earnings from the experiment, subjects were paid a fixed "show-up" fee of $20.

B. Behavioral Predictions

As described in the introduction, identifying peer effects is difficult because of the researcher's inability to perfectly observe communication between two agents. Hence, similar behavior among two peers may stem from common information shocks or common preferences (e.g., investors in a given city may read the same local newspaper or have similar levels of risk aversion). If instead, a researcher is given full control over the information structure that is made available to all agents, then it becomes straightforward to test for peer effects. While this condition is unlikely to be satisfied in the field, we can easily construct this type of environment in a controlled laboratory setting.

In order to test for the presence of peer effects, it is useful to provide a precise description of what constitutes a peer effect in the current experimental setting. We say that a peer effect arises in the social condition when $x_{i,t}$ is dependent on the history of peer investment allocations $\{x_{j,u}\}_{u=1,\ldots,t-1}$ controlling for any function of past returns $\{R_u\}_{u=1,\ldots,t-1}$. For simplicity, the empirical tests we construct will invoke a stronger definition where a peer effect requires $x_{i,t}$ to be dependent on the most recent peer investment allocation, $x_{j,t-1}$. The reason we must control for any function of past returns is because investment decisions may be correlated within pairs because both agents are exposed to identical information sets (i.e., the history of past returns.)

While we cannot explicitly control for every function of past returns, we can exploit the fact that, within each market, subject *k* in the *control condition* sees the same sequence of returns as subject *i* and *j* in the *social treatment* (where *i* and *j* belong to the same pair). Therefore, in the absence of peer effects we expect that on average, $corr(x_{i,t}, x_{j,t-1}) = corr(x_{i,t}, x_{k,t-1})$. In other words, there should no difference between the within-pair and across-pair correlations in investment decisions. In contrast, if there is a causal effect of *j's* investment allocation in period *t-1* on *i's* investment allocation in period *t,* we should detect a difference between the within-pair and across-pair correlations. This leads to our first prediction:

Prediction 1 (Behavioral): Let *i* and *j* be randomly assigned peers in the *social treatment*, and let *k* be any subject in the *control condition.* If there are no peer effects, then the within-pair correlation should on average, equal the across pair correlation: $corr(x_{i,t} , x_{j,t-1}) = corr(x_{i,t} , x_{k,t-1})$. In contrast, if there are peer effects, then on average, $corr(x_{i,t} , x_{j,t-1}) \neq corr(x_{i,t} , x_{k,t-1})$.

C. Neural Predictions

While Prediction 1 is concerned with identifying the presence of peer effects, in this section we develop predictions regarding the different mechanisms that can give rise to peer effects. In particular, we exploit the availability of neural data at the peer decision screen, in order to generate testable predictions of competing theories.

Broadly speaking, there are two main channels through which peer effects can arise. First, peer decisions may be correlated if *i* learns about the underlying return distribution from *j's* decision. This is often referred to as observational learning (Bikhchandani et al. (1992)) and acts through an information channel. Second, peer effects may arise if *i*'s utility depends directly on *j*'s decision. In what follows, we focus on a specific mechanism within the class of social utility mechanisms where *i* has utility over relative wealth compared to *j*.

We focus our predictions on neural activity that is generated at the moment when the peer decision screen appears. Recall that when this screen appears, a subject observes his peer's investment allocation but is not required to make any active decisions. Hence, this screen is exclusively involved with information processing (as opposed to decision deliberation), and this is useful for two reasons. First, because we are using neural data that is generated in the absence of choice, it provides an explicit example of how the neural data is useful above and beyond traditional choice data. Second, we can rely on a large literature in decision neuroscience that can generate predictions about neural activity when new information -- in this case, a peer's

investment decision -- is revealed. Specifically, we focus on the concept of a *prediction error*, which we briefly review now.

A prediction error is a signal that the brain computes in response to new information. This signal can be thought of as measuring the change in the expected net present value of utility generated by the news, taking into account all sources of utility. While this terminology may be unfamiliar to most financial economists, the general concept of a prediction error should be familiar as it is closely related to basic properties of asset pricing. To see this, consider a simple model where an asset's price is equal to the sum of its discounted expected cash flows. Prices should then fluctuate only when unexpected news is revealed about future cash flows, and this price fluctuation should occur at the moment when the unexpected news is revealed. Furthermore, the sign and size of the price change should reflect whether cash flows are expected to increase or decrease, and by how much. Similarly, a prediction error signal embodies the same two core properties that a price change exhibits in an efficient market: (i) it is different from zero only at the time when unexpected news is revealed and (ii) it carries information about the signed change in expected utility generated from the unexpected news.

Critically, a large body of evidence in decision neuroscience shows that the prediction error signal can be accurately measured in a specific area of the brain, the vSt (Schultz et al. (1997); Pessiglione et al. (2006); Hare et al. (2008); Lin et al. (2012)). This is useful because it allows us to empirically measure the prediction error signal generated in response to news, and therefore, we can infer which factors cause a change in discounted expected utility. For each competing mechanism of peer effects, we can then derive the theoretical prediction about the change in expected discounted utility generated by unexpected news, and then test whether vSt activity significantly correlates with these changes. We now develop these theoretical predictions, which are difficult to test with only behavioral data, but can be directly tested with neural activity in the vSt.

First, we outline the predictions regarding neural activity that are generated under the social learning mechanism. While all subjects are given the same information regarding the markets at the start of the experiment, we cannot rule out *a priori* the possibility that *i* may perceive his peer, *j,* to possess superior information processing ability. Hence, it is possible that even though *i* and *j* have identical information sets (i.e., the history of past returns), *i* may believe that *j*'s decision carries information about the underlying distribution of returns. In this case, we expect *i* to update upwards (downwards) his belief about the value of the risky asset after observing *j* increase (decrease) the fraction of his wealth allocated to the risky asset. All things equal, when subject *i* updates his beliefs upward about the value of the asset, this news should increase his discounted expected utility as it implies that the risky asset is a better investment opportunity than previously expected. Therefore, the change in *j*'s risky asset allocation, $(x_{j,t} - x_{j,t-1})$, is a prediction error, and we expect to see this quantity correlate with vSt activity at the moment when the peer decision screen is revealed[7]. This leads to our second prediction:

Prediction 2 (Neural): If peer effects are driven by a social learning mechanism, then the vSt should compute a prediction error signal at the revelation of the peer decision screen. This prediction error signal is defined by: $(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})$

We now turn to our final prediction, one that characterizes the neural activity that should be observed if peer effects are generated through a preference channel. We test a specific theory

[7] Before we state the precise definition of this neural prediction, a small adjustment is needed because the market return in each trial is revealed before the peer decision is revealed. To understand why this adjustment is necessary, consider the peer decision revealed in trial *t*. When subject *i* observes $x_{j,t}$, he is observing how *j* changed his investment allocation as a function of the market return on trial *t-1*. However, subject *i* has also updated his beliefs about the risky asset as a function of the market return on trial *t-1*, and this should be captured by the change in his investment allocation $(x_{i,t} - x_{i,t-1})$. Therefore, to control for the updating that subject *i* does through private learning, we subtract this quantity from the change in *j's* risky asset allocation. The prediction error attributed to social learning by subject *i* on trial *t* is therefore equal to: $(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})$. We note that this specification assumes that subject *i* is naïve in the sense that he believes his peer will change an investment allocation only as a function of market returns, and hence not as a function of *i's* decision (Eyster and Rabin (2010)).

within the class of social preferences, one based on the "keeping up with the Jones'" preferences, where subjects have utility over relative wealth (Abel (1990); Gali (1994)). Note that there is no incentive in our experimental design that induces such a concern for relative wealth, and as such, our test of this mechanism relies on "homegrown" preferences over relative wealth.

To fix ideas, we assume that in addition to standard sources of utility, subjects also derive utility from the difference between their peer's final portfolio value and their own final portfolio value[8]. This implies that the change in *j's* portfolio value on each trial contains news about *i's* expected discounted utility. Therefore, the revelation of the peer decision (along with the risky asset return) should generate a prediction error for subject *i.* Recall that subjects are not explicitly shown their peer's portfolio value on each peer decision screen, and so it may be difficult for them to compute the actual dollar value change in their peer's portfolio each period. However, subjects are shown the fraction of wealth their peer allocated to the risky asset, $x_{j,t}$, and the risky asset return R_t on each trial. It is then reasonable to expect that subjects can compute both the percentage change in their peer's portfolio value, and the percentage change in their own portfolio value on each trial.

We therefore proxy the dollar change in portfolio values with the percentage change in portfolio values. The prediction error that is generated under the relative consumption mechanism should then be equal to the difference between the percentage change in *i's* portfolio value and the percentage change in *j's* portfolio value. That is, neural activity in subject *i's* vSt at the peer decision screen should negatively correlate with the quantity, $R_t \times (x_{j,t} - x_{i,t})$. This leads to our third and final prediction:

[8] For example, a subject could have utility of the form $u(m, y) = m + g(m-y)$, where m is a subject's own wealth, y is the amount of peer wealth, an g is an increasing function.

Prediction 3 (Neural): If peer effects are driven by a relative consumption mechanism, then the vSt should compute a prediction error signal at the revelation of the peer decision screen. This prediction error signal is defined by: $-R_t \times (x_{j,t} - x_{i,t})$.

II. fMRI Data Collection and Analysis

In this section, we describe how the fMRI measures of neural activity were collected and analyzed. Our goal in this section, which is taken primarily from Frydman et al. (2014a), is to provide a brief primer on the basics of fMRI data analysis, so that the underlying econometrics are approachable to those not already familiar with the fMRI literature. For a more detailed discussion, see Huettel, Song, and McCarthy (2004), Ashby (2011), and Poldrack, Mumford, and Nichols (2011).

A. fMRI Data Collection and Measurement

We collected measures of neural activity over the entire brain using BOLD-fMRI, which stands for blood-oxygenated level dependent functional magnetic resonance imaging. BOLD-fMRI measures changes in local magnetic fields that result from the local inflows of oxygenated hemoglobin and outflows of de-oxygenated hemoglobin that occur when neurons fire. In particular, fMRI provides measures of the BOLD response in small "neighborhoods" of brain tissue called *voxels*, and is thought to measure the sum of the total amount of neuronal firing into that voxel and the total amount of neuronal firing within the voxel.

One important complication is that the hemoglobin responses measured by BOLD-fMRI are slower than the associated neuronal responses. Specifically, although the bulk of the neuronal response takes place quickly, BOLD measurements are affected for up to 24 seconds thereafter. Figure 2A provides a more detailed illustration of the nature of the BOLD response. The top

panel depicts the path of the BOLD signal in response to one (arbitrary) unit of neural activity of infinitesimal duration at time zero. The function plotted here is called the canonical hemodynamic response function (HRF). It is denoted by $h(\tau)$, where τ is the amount of time elapsed since the neural activity impulse, and has been shown to approximate well the pattern of BOLD responses for most subjects, brain areas, and tasks.

Fortunately, there is a standard way of dealing with the complication described in the previous paragraph. In particular, the BOLD response has been shown to combine linearly across multiple sources of neural activity (Boynton et al. (1996)). This property, along with knowledge of the specific functional form of the HRF, allows us to construct a mapping from predicted neural activity to predicted BOLD responses. Specifically, if the predicted level of neural activity at any particular time is given by $a(t)$, then the level of BOLD activity at any instant t is well approximated by

$$b(t)=\int_0^\infty h(u)a(t-u)\,du, \tag{1}$$

which is the convolution between the HRF and the neural inputs. This integral has a straightforward interpretation: it is a lagged sum of all the BOLD responses triggered by previous neural activity. The properties of the BOLD response are illustrated in Fig. 2B, which depicts a hypothetical path of neural activity (solid line), together with the associated BOLD response (dashed line).

During our experiment, we acquire two types of MRI data in a 3.0 Siemens Tesla Trio MRI. First, we acquire BOLD-fMRI data while the subjects perform the experimental task. We use a voxel size of 3.4 mm x 3.4 mm x 4 mm, and collect these data for the entire brain every 2 seconds. We also acquire high-resolution anatomical scans that we use mainly for realigning the brains across subjects and for localizing the brain activity identified by our analyses.

B. fMRI Main Data Analyses

The key goals of our analysis are to test if the region of the vSt that has been repeatedly shown to encode prediction errors at the time of utility-relevant news exhibits activity consistent with Predictions 2 and 3. To do this, we run statistical tests to see if there are areas within this region of the brain, given by collections of spatially contiguous voxels called *clusters,* where the BOLD response reflects neural activity that implements the computations of interest (prediction errors generated by social learning or relative wealth mechanisms). This is complicated by the fact that, since every voxel contains thousands of neurons, the BOLD responses in a voxel can be driven by multiple signals. Fortunately, the linear properties of the BOLD signal allow the neural signals of interest to be identified using standard linear regression methods.

The general statistical procedure is straightforward, and will be familiar to most economists. The analysis begins by specifying two types of variables that might affect the BOLD response: target computations and additional controls. The target computations reflect the signals we are looking for (e.g., a prediction error at the revelation of a peer decision). They are specified by a time series $s_i(t)$ describing each signal of interest. For each of these signals, let $S_i(t)$ denote the time series that results from convolving the signal $s_i(t)$ with the HRF, as described above. The additional controls, denoted by $c_j(t)$, are other variables that might affect the BOLD time series (e.g., residual head movement or time trends). These are introduced to further clean up the noise in the BOLD signal, but are not explicitly used in any of our tests. The control variables are not convolved with the HRF because, while they affect the measured BOLD responses, they do not reflect neural activity which triggers a hemodynamic response.

The linearity of the BOLD signal implies that the level of BOLD activity $b^v(t)$ in any voxel v at time t should be given by

$$b^v(t) = \text{constant} + \sum_i \beta_i^v S_i(t) + \sum_j \alpha_j^v c_j(t) + \varepsilon(t), \quad (2)$$

where $\varepsilon(t)$ denotes AR(1) noise. This model is estimated independently in each of the voxels that fall within vSt. Our hypotheses can then be restated as tests about the coefficients of this regression model: signal i is said to be associated with activity in voxel v only if β_i^v is significantly different from zero.

Two additional considerations apply to most fMRI studies, including this one. First, we are interested in testing hypotheses about the distribution of the signal coefficients in the population of subjects, not hypotheses about individual subject coefficients. This would normally require estimating a mixed effects version of the linear model specified above, which, given the size of a typical fMRI dataset, would be computationally intensive. Fortunately, there is a shortcut that provides a good approximation to the full mixed effects analysis (Penny et al. (2006)). It involves estimating the parameters separately for each individual subject, averaging them across subjects, and then performing *t*-tests. This is the approach we follow here.

Second, we conduct our tests in an area of the vSt that, in prior work, has been linked to the computation of prediction errors. Specifically, we construct two 8 mm radius spheres (a total of 56 voxels) around the coordinates (MNI-space, $[x = -10, y = 12, z = -8]$, $[x = 10, y = 12, z = -8]$) that were found to exhibit peak correlation with prediction errors in Pessiglione et al. (2006). Because our hypothesis tests are therefore carried out in each of the 56 voxels in the relevant vSt region of interest, there is a concern about false-positives. To address this problem, we correct for multiple comparisons within the relevant region of interest, a procedure known in the fMRI literature as a small volume correction (SVC). We report results as significant if they pass SVC correction at a level of $p<0.05$[9].

[9] Specifically, we report results as significant if voxels pass SVC with a family-wise error rate of less than 0.05 and if they pass $p<0.005$ uncorrected with a 10-voxel extent threshold.

III. Results

A. Test of Behavioral Prediction 1

We begin our test of Prediction 1 by computing, for each subject *i* in the *social treatment,* the correlation between his risky asset allocation in period *t,* $x_{i,t}$, and his peer's risky asset allocation in period *t-1,* $x_{j,t-1}$. We find that the average value of $corr(x_{i,t}, x_{j,t-1})$ across all twenty-four subjects in the *social treatment* is 0.300 (standard error: 0.042). Next, we compute the correlations generated by the risky asset allocations of subject *i* in the *social treatment* and subject *k* in the *control condition.* Specifically, we fix subject *i* in the *social treatment,* and then compute $corr(x_{i,t}, x_{k,t-1})$, for each of the $k=1,\ldots,24$ subjects in the *control condition.* We then average these twenty-four correlations, and take this to be our across-pair correlation measure for subject *i*. Finally, we take the average of this across all twenty-four subjects in the *social treatment,* which equals 0.094 (standard error: 0.019).

We can immediately reject the null hypothesis that the within-pair and across-pair correlations are equal with a *t*-statistic of 6.30. Figure 3 shows that for all but two of the twenty-four subjects in the *social treatment,* the within-pair correlation is greater than the across-pair correlation. Because the experimental design ensures that subjects in the *control condition* and *social treatment* have identical information sets (i.e, they observe the same history of returns) the result that the within pair correlation is significantly greater than the across-pair correlation provides evidence that peer decisions causally affect investment allocations.

To gain further insight into other factors that explain the time series variation in risky asset allocations, we can impose additional structure on the model that subjects use to form their allocations. Because subjects are not told anything about the actual data generating process that governs the risky asset returns, they are faced with a learning problem. In such a setting, a natural

hypothesis is that a subject computes the conditional Sharpe ratio after observing market returns and then uses this statistic to guide the investment decision.

To test this hypothesis, we estimate an OLS regression of subject *i's* current investment allocation on the conditional Sharpe ratio, his peer's previous investment allocation, and subject *i's* previous investment allocation:

$$x_{i,t} = \alpha + \gamma_i + \beta_1 x_{j,t-1} + \beta_2 sharpe\_ratio_t + \beta_3 x_{i,t-1} + \varepsilon_{i,t} \quad (3)$$

We estimate the model separately for the *social treatment* and for the *control condition* and results are displayed in Table 1. The first column of Table 1 confirms the previous univariate correlation results, as *j's* previous investment allocation is a significant predictor of *i's* current investment allocation. Furthermore, column 1 (using data only from *social treatment*) and column 2 (using data only from *control condition*) both show that the conditional Sharpe ratio is a significant predictor of risky asset allocations. Therefore, in the *social treatment,* variation in risky asset allocations can be decomposed into a social component (the peer investment decision) and a non-social component (the conditional Sharpe ratio)[10]. Interestingly, when pooling data across control and treatment (column 3), the interaction between the treatment and the conditional Sharpe ratio has a negative coefficient, suggesting that when peer decisions are observable, subjects rely less on past market data.

While the regression results presented above are consistent with peer effects, one alternative theory that can potentially explain part of the observed behavior is anchoring (Tversky and Kahneman (1974)). Anchoring occurs when individuals rely on arbitrary initial values or starting points when computing decisions in complex environments. Our subjects may perceive the experiment as a complex environment, and peer decisions can provide anchors for subjects in the *social treatment* when they are choosing their risky asset allocations. In this case, we would

[10] The fact that the conditional Sharpe ratio is a significant predictor of investment decisions is consistent with the positive *across-pair* correlations we document in Figure 3. In other words, the conditional Sharpe ratio may be the channel through which two subjects exhibit correlated decisions, despite their inability to observe each other's decisions.

expect to see the observed positive correlation between *j's* allocation on trial *t-1* and *i's* allocation on trial *t.*

However, we can rule out this alternative theory using the *control condition* data in Table 1. Recall that in the *control condition,* subjects do not have access to a peer's decision on each trial, but instead see a randomly drawn number over the interval [0,1] on the peer decision screen. Hence, if peer effects are driven by anchoring, we should still observe a correlation between the uniformly drawn number (the anchor) and the risky asset allocation. Instead, we find that the coefficient on the peer investment is not significantly different from zero in the *control condition*, and furthermore, the coefficient on the interaction between the social condition and the peer investment (column 3) is significantly positive. Together, these results cast doubt on the anchoring hypothesis.

B. Neural Response to Market Returns

Before turning to the results on neural predictions 2 and 3, we first describe a preliminary result on neural activity at the time market returns are revealed. This preliminary result will act to validate the methodology we use in the next section to test key neural predictions 2 and 3. Recall that in each trial, a subject observes the market return before observing his peer's investment allocation (Figure 1). If a subject observes a positive (negative) market return, then conditional on investing in the risky asset, he will experience an increase (decrease) in wealth. This change in wealth carries news about future consumption, and hence a positive (negative) market return should generate a positive (negative) prediction error in the vSt. Therefore, when the market return is revealed to a subject, we expect vSt activity to positively correlate with the subject's change in wealth, given by $R_t x_{i,t}$.

To check whether vSt activity is consistent with this prediction, we estimate a general linear model (GLM) of BOLD activity in every subject and voxel.

$$b^v(t) = \alpha + \beta_1^v I_{\mathrm{PD}}(t) \times R_t\,(x_{j,t} - x_{i,t}) + \beta_2^v I_{\mathrm{PD}}(t) \times [(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})] + \beta_3^v I_{MKT}(t) \times (R_t\, x_{i,t}) + \beta_4^v controls + \varepsilon(t) \quad (4)$$

This is the same model that we will use to perform tests of predictions 2 and 3, so it is useful to understand its components in detail. $b^v(t)$ denotes the BOLD signal at time t in voxel v. I_{PD} is an indicator function that equals one if, at time t, the peer decision screen is revealed; I_{MKT} is an indicator function that equals one if, at time t, the market screen is revealed; the first nonconstant regressor, $I_{PD}(t) \times R_t(x_{j,t} - x_{i,t})$, represents the relative change in wealth that is revealed at the peer decision screen; the second nonconstant regressor, $I_{PD}(t) \times [(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})]$, represents the social learning signal described in section I.C. The third nonconstant regressor, $I_{MKT}(t) \times (R_t x_{i,t})$, represents the absolute change in wealth that the subject experiences during the "market screen" at time *t.* The "controls" vector includes the following variables: 1) an indicator function denoting the onset of an allocation screen, 2) an indicator function denoting the onset of a market screen, 3) an indicator function denoting the onset of a market screen interacted with the risky asset return, 4) an indicator function denoting the onset of a peer decision screen, 5) an indicator function denoting the onset of a new market and 6) 12 movement regressors that control for subject head movement inside the scanner.

Controls (1) - (5) are convolved with the HRF, whereas control (6) is not. We need to include these controls because the BOLD signal is affected for up to 24 seconds after the initial neural impulse generated by the onset of any single event. Therefore, even though we are concerned only with the neural activity generated at the moment when a peer decision or market return is revealed, the observed BOLD signal at this time is still affected by the onset of several preceding screens.

The next step in testing whether vSt activity correlates with changes in wealth is to perform inference about the extent to which the signal of interest is encoded in a given voxel. We

do this by carrying out a one-sided *t*-test against zero of the average of the individually estimated coefficients. In other words, we compute the average β_3^v across subjects and perform hypothesis tests under the null that the coefficient is zero; finally, we correct for multiple comparisons within the pre-specified vSt region of interest.

Figure 4 shows the results of this hypothesis. Within the pre-specified 56-voxel vSt target region (those voxels colored in yellow and orange) associated with the computation of prediction errors in previous studies, we find a cluster of 15 voxels (those voxels colored in orange and red) where β_3^v, averaged across subjects, is significantly positive ($p<0.05$ SVC).

C. Test of Neural Predictions 2 and 3

We now turn to Predictions 2 and 3, which investigate the extent to which neural activity in the vSt reflects prediction errors that are generated under the two different peer effects mechanisms outlined in section I.C. These are the key predictions of the paper, because they allow us to test between competing mechanisms of peer effects that are difficult to test using only trading data. We proceed in testing Predictions 2 and 3 by performing hypothesis tests about the estimated coefficients from equation (4). Figure 5A shows that within the pre-specified 56-voxel vSt target region (those voxels colored in yellow and orange) associated with the computation of prediction errors in previous studies, we find a cluster of 9 voxels where β_2^v, averaged across subjects, is significantly positive ($p<0.05$ SVC). This indicates that at the time when a peer decision screen is revealed, activity in the vSt positively correlates with the theoretical prediction error that is generated under the social learning mechanism; hence the vSt activity that we observe is consistent with Prediction 2.

Our test of Prediction 3, which investigates the relative wealth mechanism, proceeds in a very similar manner. Figure 5B shows that within the pre-specified 56-voxel vSt target region (again, those voxels colored in yellow and orange), we find a cluster of 19 voxels where β_1^v, averaged across subjects, is significantly negative ($p<0.05$ SVC). This indicates that at the time

when a peer decision screen is revealed, activity in subject *i's* vSt significantly correlates with the change in his relative wealth.

One concern with the previous result is that changes in relative wealth can be highly correlated with changes in absolute wealth, especially if a peer invests only a small amount of wealth in the risky asset. Therefore, a prediction error that is consistent with a relative wealth mechanism may instead be driven by a change in absolute wealth. However, there is a key aspect of the experimental design that allows us to rule out this interpretation. Recall that a prediction error is generated only in response to unexpected information, and therefore, it will be equal to zero in response to information that has already been revealed in the past. In our setting, a subject observes the market return before observing his peer's decision (Figure 1); it follows that any change in expected discounted utility that is driven by a change in absolute wealth must be reflected in the prediction error at the time the market return is revealed. This is precisely the signal that is displayed in Figure 4. Therefore, the prediction error generated at the peer decision screen cannot be driven by changes in absolute wealth, because this information was previously revealed.

Taken together, the results from this section provide support in favor of both Prediction 2 and Prediction 3. It is important to point out that each mechanism explains a unique portion of the variation in vSt activity at the peer decision screen: the two key variables that parameterize the prediction errors of each mechanism are entered into the same general linear model, and are allowed to compete for explained variance. Therefore, our data suggest that peer effects are simultaneously driven through a preference channel and an information channel.

D. Neural Test of a Preference for Conformity

While the neural data is consistent with a specific type of relative wealth concern mechanism, there are other preference-based theories that can potentially generate peer effects in our setting. Here we examine one alternative explanation based on a direct preference for

conformity (Asch (1951)). If subject *i* in our experiment is influenced by subject *j* because of a direct preference for conformity, then on each trial, *i* should derive utility from observing that *j* made a similar investment allocation. In other words, at the time a peer decision is revealed, a direct taste for conformity should generate utility that is decreasing in the distance between subject *i's* allocation and subject *j's* allocation. As such, when the peer decision is revealed, the vSt of subject *i* should compute a prediction error that negatively correlates with $|x_{i,t} - x_{j,t}|$. We can test this prediction by estimating a GLM very similar to the model in equation (4). In fact, the only change we make is to substitute the first nonconstant regressor with the hypothesized prediction error signal that is generated under the conformity theory, $|x_{i,t} - x_{j,t}|$. Specifically, the model we estimate is:

$$\begin{aligned} b^v(t) = \alpha + \beta_1^v I_{\mathrm{PD}}(t) \times |\, x_{j,t} - x_{i,t} \,| + \beta_2^v I_{\mathrm{PD}}(t) \times [(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})] \\ + \beta_3^v I_{MKT}(t) \times (R_t\, x_{i,t}) + \beta_4^v controls + \varepsilon(t) \end{aligned} \tag{5}$$

where the vector of controls is exactly the same as specified in equation (4) and described in section III.B. Contrary to the conformity hypothesis, we do not find any activity in our pre-specified vSt region that is significantly associated with the distance between peer investment allocations at our omnibus threshold of $p<0.05$ SVC. In summary, this null result casts some doubt on the ability of the conformity hypothesis to explain the observed peer effects.

IV. Discussion

In this paper we show how neural data can be used to test the mechanisms that underly peer effects. In nearly all settings in which peer effects have previously been studied in finance, it is extremely difficult to distinguish between different mechanisms because of the researcher's lack of control over information (Manski (1993)) and the endogeneity associated with selective communication (Han and Hirshleifer (2013)). In contrast, a large degree of experimental control

combined with neural data enable us to perform tests of the competing mechanisms. Overall, we find strong evidence of peer effects in the trading data, and the neural data suggest that both a social learning mechanism and a relative wealth mechanism contribute to the observed behavior. This evidence is important because it can guide the development of newer models of financial markets that have begun to incorporate social interactions among investors (Hirshleifer (2014); Shiller (2014)).

While there has been very little work on examining the mechanisms underlying peer effects in finance, we are not the first to tackle this question directly. A recent paper by Bursztyn et al. (2014) uses a field experiment in Brazil to distinguish between a preference-based and a belief-based explanation for peer effects. The authors find that both channels contribute to peer effects, and they show that less financially sophisticated investors tend to use the social learning mechanism more. Our results complement those of Bursztyn et al. (2014) as we use a completely different methodology, and also find that both channels are at work in generating peer effects. Moreover, while these authors abstract away from the outcome phase of the investment problem, we provide subjects with the necessary information to assess the change in their peer's portfolio value, which is an important ingredient in the relative wealth mechanism.

It is also important to note that we do not provide a general test of the preference-based explanation for peer effects, as is done in Bursztyn et al. (2014). Instead we focus only on one specific theory, namely a relative wealth mechanism where agents have utility over the difference between their final wealth and their peer's final wealth (Abel (1990); Gali (1994); Hong et al. (2014)). As described in the previous section, there are other preference-based explanations that can generate peer effects, such as a direct preference for conformity. Although we provide some evidence against the conformity theory, these results are based on a failure to reject the null hypothesis, and thus should be interpreted with caution.

Finally, because our experimental design does not induce subjects to have utility over their peer's wealth, we are testing an exogenous relative wealth mechanism, instead of an

endogenous one that may arise due to scarce resources in the economy (DeMarzo et al. (2008)). This provides a relatively conservative test of the theory as we rely on "hard-wired" preferences over relative wealth, which can be challenging to directly test without access to neural data (Fliessbach et al. (2007)).

REFERENCES

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the joneses, *The American Economic Review* 38-42.

Asch, Solomon E., 1951, Effects of group pressure upon the modification and distortion of judgments, *Groups, leadership, and men* 222-236.

Ashby, F. Gregory, 2011. *Statistical analysis of fmri data* (The MIT Press).

Beshears, John, James Choi, David Laibson, Brigitte Madrian, and Katherine Milkman, 2014, The effect of providing peer information on retirement savings decisions, *Journal of Finance* Forthcoming.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, 1992, A theory of fads, fashion, custom, and cultural change as informational cascades, *Journal of Political Economy* 992-1026.

Bossaerts, Peter, and Elise Payzan-Le Nestour, 2014, Learning about unstable, publicly unobservable payoffs, *Review of Financial Studies* Forthcoming.

Boynton, Geoffrey M., Stephen A. Engel, Gary H. Glover, and David J. Heeger, 1996, Linear systems analysis of functional magnetic resonance imaging in human v1, *The Journal of Neuroscience* 16, 4207-4221.

Brown, Jeffrey R., Zoran Ivković, Paul A. Smith, and Scott Weisbenner, 2008, Neighbors matter: Causal community effects and stock market participation, *Journal of Finance* 63, 1509-1531.

Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman, 2014, Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions, *Econometrica* 82, 1273-1301.

Caplin, Andrew, and Mark Dean, 2008, Dopamine, reward prediction error, and economics, *The Quarterly Journal of Economics* 123, 663-701.

Caplin, Andrew, Mark Dean, Paul W. Glimcher, and Robb B. Rutledge, 2010, Measuring beliefs and rewards: A neuroeconomic approach, *Quarterly Journal of Economics* 125, 923-960.

De Martino, Benedetto, John P. O'Doherty, Debajyoti Ray, Peter Bossaerts, and Colin Camerer, 2013, In the mind of the market: Theory of mind biases value computation during financial bubbles, *Neuron* 79, 1222-1231.

DeMarzo, Peter M., Ron Kaniel, and Ilan Kremer, 2008, Relative wealth concerns and financial bubbles, *Review of Financial Studies* 21, 19-50.

Duflo, Esther, and Emmanuel Saez, 2003, The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment, *Quarterly Journal of Economics* 815-842.

Eyster, Erik, and Matthew Rabin, 2010, Naive herding in rich-information settings, *American economic journal: microeconomics* 2, 221-243.

Fliessbach, Klaus, Bernd Weber, Peter Trautner, Thomas Dohmen, Uwe Sunde, Christian E. Elger, and Armin Falk, 2007, Social comparison affects reward-related brain activity in the human ventral striatum, *Science* 318, 1305-1308.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel, 2014a, Using neural data to test a theory of investor behavior: An application to realization utility, *Journal of Finance* 69, 907-946.

Frydman, Cary, Colin Camerer, and Antonio Rangel, 2014b, Neural measures of regret and reinforcement learning from investment experiences, *Working Paper*.

Gali, Jordi, 1994, Keeping up with the joneses: Consumption externalities, portfolio choice, and asset prices, *Journal of Money, Credit and Banking* 1-8.

Han, Bing, and David Hirshleifer, 2013, Self-enhancing transmission bias and active investing, *Working Paper*.

Hare, Todd, John O'Doherty, Colin F. Camerer, Wolfram Schultz, and Antonio Rangel, 2008, Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors, *Journal of Neuroscience* 28, 5623-5630.

Hirshleifer, David, 2014, Behavioral finance, *Working Paper*.

Hirshleifer, David, and Siew Hong Teoh, 2009, Thought and behavior contagion in capital markets, *Handbook Of Financial Markets: Dynamics And Evolution, Handbooks in Finance* 1-46.

Hong, Harrison, Wenxi Jiang, Na Wang, and Bin Zhao, 2014, Trading for status, *Review of Financial Studies*.

Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, 2004, Social interaction and stock market participation, *Journal of Finance* 59, 137-163.

Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, 2005, Thy neighbor's portfolio: Word of mouth effects in the holdings and trades of money managers, *Journal of Finance* 60, 2801-2824.

Huettel, Scott, Allen Song, and Gregory McCarthy, 2004. *Functional magnetic resonance imaging* (Sinauer Associates).

Ivković, Zoran, and Scott Weisbenner, 2007, Information diffusion effects in individual investors' common stock purchases: Covet thy neighbors' investment choices, *Review of Financial Studies* 20, 1327-1357.

Kaustia, Markku, and Samuli Knüpfer, 2012, Peer performance and stock market entry, *Journal of Financial Economics* 104, 321-338.

Knutson, Brian, G. Elliott Wimmer, Camelia M. Kuhnen, and Piotr Winkielman, 2008, Nucleus accumbens activation mediates the influence of reward cues on financial risk taking, *NeuroReport* 19, 509-513.

Kuhnen, Camelia, 2014, Asymmetric learning from financial information, *Journal of Finance* Forthcoming.

Kuhnen, Camelia M., and Brian Knutson, 2005, The neural basis of financial risk taking, *Neuron* 47, 763-770.

Levy, Ifat, Stephanie C. Lazzaro, Robb B. Rutledge, and Paul W. Glimcher, 2011, Choice from non-choice: Predicting consumer preferences from blood oxygenation level-dependent signals obtained during passive viewing, *The Journal of Neuroscience* 31, 118-125.

Lin, Alice, Ralph Adolphs, and Antonio Rangel, 2012, Social and monetary reward learning engage overlapping neural substrates, *Social Cognitive and Affective Neuroscience* 7, 274-281.

Lo, Andrew W., and Dmitry V. Repin, 2002, The psychophysiology of real-time financial risk processing, *Journal of Cognitive Neuroscience* 14, 323-339.

Lohrenz, Terry, Meghana Bhatt, Nathan Apple, and P. Read Montague, 2013, Keeping up with the joneses: Interpersonal prediction errors and the correlation of behavior in a tandem sequential choice task, *PLoS computational biology* 9.

Penny, William, Karl Friston, John Ashburner, Stefan Kiebel, and Thomas Nichols, 2006. *Statistical parametric mapping: The analysis of functional brain images* (Academic Press).

Pessiglione, Mathias, Ben Seymour, Guillaume Flandin, Raymond J. Dolan, and Chris D. Frith, 2006, Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans, *Nature* 442, 1042-1045.

Poldrack, Russell A., Jeanette Mumford, and Thomas Nichols, 2011. *Handbook of functional mri data analysis* (Cambridge University Press).

Pool, Veronika, Noah Stoffman, and Scott Yonker, 2014, The people in your neighborhood: Social interactions and mutual fund portfolios, *Journal of Finance* Forthcoming.

Preuschoff, Kerstin, Peter Bossaerts, and Steven R. Quartz, 2006, Neural differentiation of expected reward and risk in human subcortical structures, *Neuron* 51, 381-390.

Rutledge, Robb B., Mark Dean, Andrew Caplin, and Paul W. Glimcher, 2010, Testing the reward prediction error hypothesis with an axiomatic model, *The Journal of Neuroscience* 30, 13525-13536.

Schultz, Wolfram, Peter Dayan, and P. Read Montague, 1997, A neural substrate of prediction and reward, *Science* 275, 1593-1599.

Shiller, Robert, 2014, Speculative asset prices, *American Economic Review* 104, 1486-1517.

Smith, Alec, B. Douglas Bernheim, Colin Camerer, and Antonio Rangel, 2014, Neural activity reveals preferences without choices, *American Economic Journal: Microeconomics* 6, 1-36.

Smith, Alec, Terry Lohrenz, Justin King, P. Read Montague, and Colin F. Camerer, 2014, Irrational exuberance and neural crash warning signals during endogenous experimental market bubbles, *Proceedings of the National Academy of Sciences* 111, 10503-10508.

Sokol-Hessner, Peter, Ming Hsu, Nina G. Curley, Mauricio R. Delgado, Colin F. Camerer, and Elizabeth A. Phelps, 2009, Thinking like a trader selectively reduces individuals' loss aversion, *Proceedings of the National Academy of Sciences* 106, 5035-5040.

Tversky, Amos, and Daniel Kahneman, 1974, Judgment under uncertainty: Heuristics and biases, *Science* 185, 1124-1131.

Table 1. Effects of Peer Investment and Sharpe Ratio on Risky Asset Allocation. Dependent variable is subject *i's* investment on trial *t*. *Peer Previous Investment* is subject *j's* investment on trial *t-1*. *Conditonal Sharpe* is the conditional Sharpe ratio using data from periods *{1,2,…,(t-1)}, Previous Investment* is subject *i's* investment on trial *t-1*. *Social* is a dummy that takes on the value 1 if the subject is in the *social treatment.* Standard errors are clustered by subject and *t*-statistics are in parentheses below estimated coefficients. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level respectively.

	(1) Social Treatment	(2) Control Condition	(3) Pooled
Peer Previous Investment	0.053 (2.50)**	0.015 (1.24)	0.149 (1.24)
Conditional Sharpe	0.003 (2.26)**	0.012 (4.67)***	0.011 (4.76)***
Previous Investment	0.597 (9.08)***	0.401 (9.63)***	0.495 (10.80)***
Social*Peer Previous Investment			0.066 (2.38)**
Social*Conditional Sharpe			-0.008 (-2.90)***
Constant	0.133 (5.85)***	0.201 (12.19)***	0.165 (9.99)***
Observations	4320	4320	8640
R^2	0.461	0.276	0.363
Subj. Fixed Effects	Y	Y	Y

Figure 1. Schematic of the three screens displayed in a typical trial in the fMRI experiment. In each of the two hundred trials in the experiment, subjects saw a series of three screens, nearly identical to those shown below. In the first screen, subjects were asked to enter a fraction of their wealth to allocate to the risky asset. Subjects entered this fraction using a handheld device while inside the fMRI scanner, and were allowed to invest in increments of 0.1. On the second screen, subjects saw the realized return of the risky asset, and their updated portfolio value as a function of the current period's risky asset allocation. Finally on the third screen, if the subject was in the *social treatment*, he saw his peer's risky asset allocation for that trial. If the subject was in the *control condition,* he saw a randomly generated number that was uniformly distributed over [0, 1].

Figure 2. BOLD measurements of neural activity. (A) Because fMRI measures the blood oxygen level dependent (BOLD) response, and not neural activity itself, we need a mapping from neural activity to the BOLD response in order to make inferences about changes in neural activity. This mapping is known as the canonical hemodynamic response function, and is shown here as a function of one arbitrary unit of instantaneous neural activity at time 0. (B) This figure shows the BOLD response (the dashed line) that results from three sequential sources of neural activity (the solid line). The BOLD response combines linearly across multiple sources of neural activity.

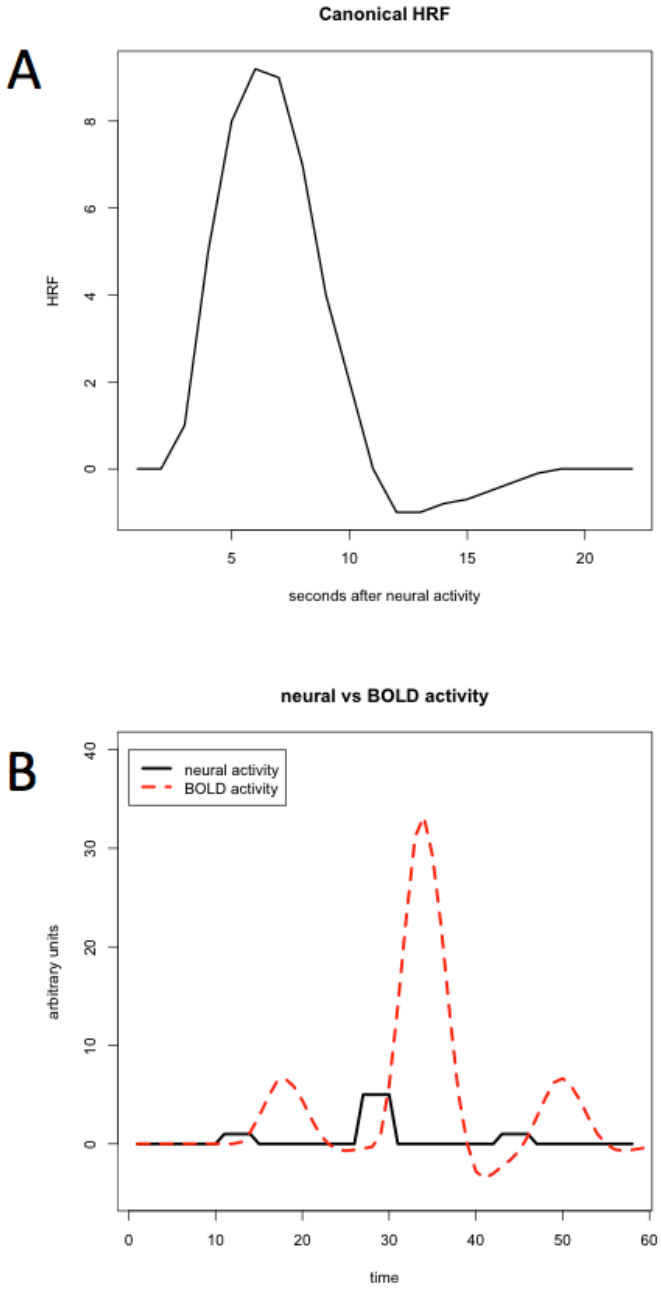

Figure 3. Subject-specific correlations between investment allocation and previous peer investment allocation. For each of the *i=1,…,24* subjects in the *social treatment*, we compute two correlation measures. First we compute $WP(i) = corr(x_{i,t}, x_{j,t-1})$ which represents the *within pair* correlation between subject *i's* investment allocation in period *t* and his peer's investment allocation in period *t-1*. This *within pair* correlation is colored in blue. Next, we compute a similar correlation, but instead of using *i*'s peer, we use subject *k* from the *control condition.* We compute this correlation for each subject *k* in the *control condition,* and then take the average. This leads to the *across pair* correlation*:* $AP(i) = \frac{1}{24}\sum_{k=1}^{24} corr(x_{i,t}, x_{k,t-1})$, which is colored in red. The figure shows that for all but two subjects, the within pair correlation is higher than the across pair correlation.

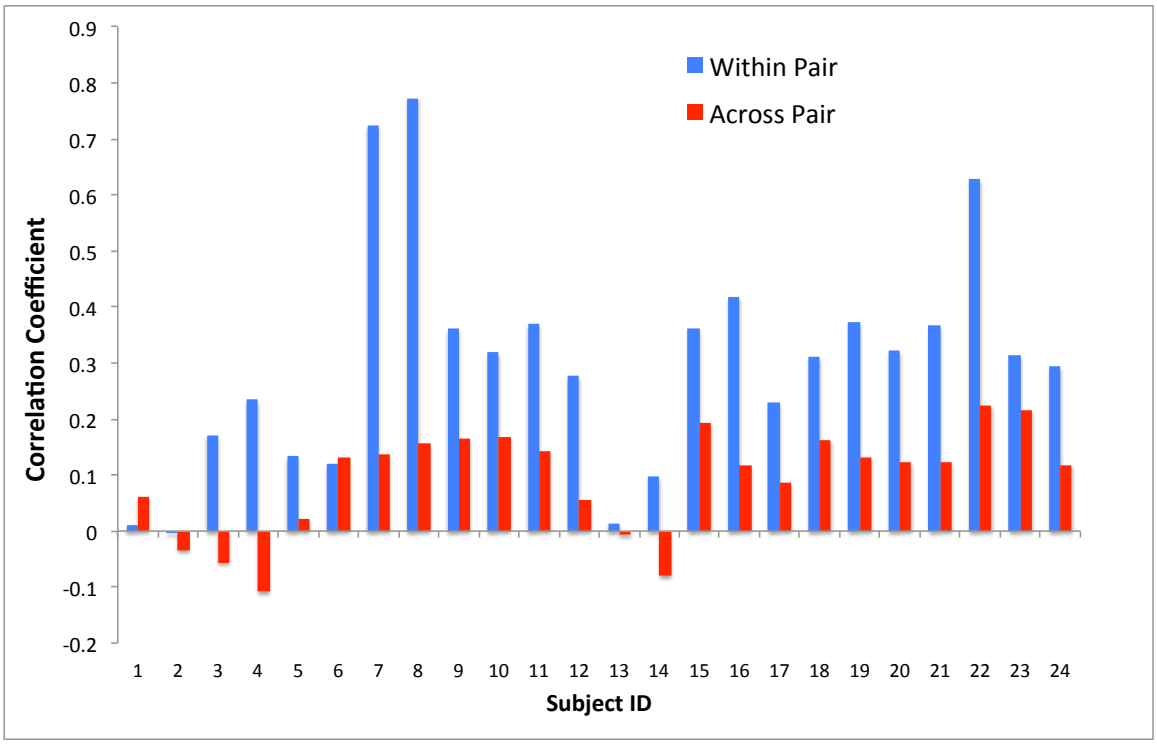

Figure 4. vSt reflects prediction errors generated from changes in wealth. The figure presents estimation results from equation (4):

$$b^v(t) = \alpha + \beta_1^v I_{\text{PD}}(t) \times R_t\,(x_{j,t} - x_{i,t}) + \beta_2^v I_{\text{PD}}(t) \times [(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})]$$
$$+ \beta_3^v I_{MKT}(t) \times (R_t\, x_{i,t}) + \beta_4^v controls + \varepsilon(t)$$

Yellow and orange voxels are those that are in our pre-specified region of interest in the vSt. Red and orange voxels are those that exhibit activity at the time of a market screen onset that significantly correlate with the change in a subject's wealth that is generated from the realized market return. In other words, the voxels in red and orange are those for which we can reject the null hypothesis that $\beta_3 = 0$. All statistics are small volume corrected at p<0.05 using FWE. The $y = 16$ coordinate indicates the two-dimensional plane shown in the brain maps.

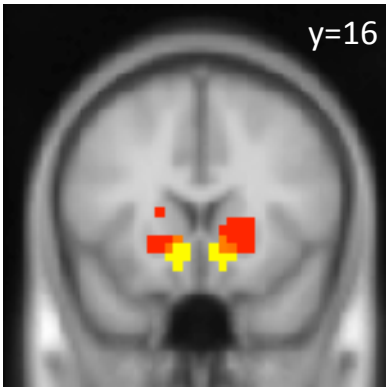

Figure 5. vSt reflects prediction errors generated under multiple peer effect mechanisms. The figure presents estimation results from equation (4):

$$b^v(t) = \alpha + \beta_1^v I_{\text{PD}}(t) \times R_t\,(x_{j,t} - x_{i,t}) + \beta_2^v I_{\text{PD}}(t) \times [(x_{j,t} - x_{j,t-1}) - (x_{i,t} - x_{i,t-1})]$$
$$+ \beta_3^v I_{MKT}(t) \times (R_t\, x_{i,t}) + \beta_4^v controls + \varepsilon(t)$$

Yellow and orange voxels are those that are in our pre-specified region of interest in the vSt. (A) Red and orange voxels are those that exhibit activity at the time of peer decision screen onset that significantly correlate with the social learning prediction error. These are voxels for which we can reject the null hypothesis that $\beta_2 = 0$. (B) Red and orange voxels are those that exhibit activity at the time of peer decision screen onset that significantly correlate with the relative wealth prediction error. These are voxels for which we can reject the null hypothesis that $\beta_1 = 0$. All statistics are small volume corrected at p<0.05 using FWE. The $y = 16$ coordinate indicates the two-dimensional plane shown in the brain maps.

A

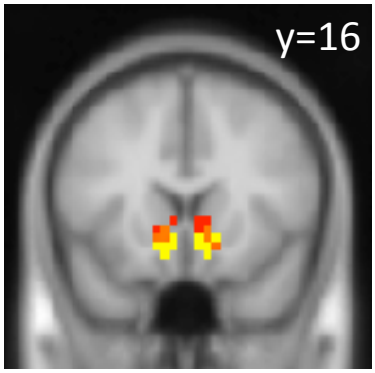

B