

Hindsight bias and investment performance

Bruno Biais (Toulouse University) & Martin Weber (Mannheim University)¹

January 2007

Abstract

Once they have observed information, hindsight biased agents fail to remember how ignorant they were initially. This bias has been amply documented by the psychological literature. To study its economic consequences, we formulate a theoretical model of the hindsight bias, providing a foundation for previous empirical measures. We show theoretically that hindsight biased agents overreact to information signals, conduct incorrect learning about volatility, and engage in suboptimal investment decisions. To test empirically the hypothesis that the hindsight bias reduces performance in important economic situations, we collect psychometric and performance data from a sample of investment bankers in London and Frankfurt. We find strong evidence that these agents are hindsight biased, and that more biased agents have lower performance. These findings are robust to differences in location, information and experience.

¹Many thanks for thoughtful comments to participants in the Toulouse and Amsterdam Universities seminars, the Warwick Conference on Behavioral Finance (2006) and the Munich Summer School on Economics and Psychology (2006), in particular Denis Hilton, Luis Garicano, Robin Hogarth, Florencio Lopez de Silanes, Thomas Mariotti, Paola sapienza, Jean Tirole and Luigi Zingales.

Hindsight bias and investment performance

1 Introduction

Experimental psychologists have extensively documented the prevalence of cognitive biases. Are these relevant for economic decisions? Experimental studies document behaviour which cannot be explained by standard rational choice theory. Are these experimental results relevant for actual financial markets? Experienced professionals, operating in their natural habitat and with a lot at stake are likely to be more rational than lay participants in experiments with minor stakes. Biases could be expected to be weeded out by learning and competition in the market environment, so that psychological biases would be irrelevant for economic thinking. To shed light on these issues, we collected field data from highly-paid bankers in Frankfurt and London relating one important psychological bias to their economic performance.

Decision making in financial markets relies crucially on information processing and learning. Efficient learning requires comparing new information to previous expectations. The hindsight bias, which is the inability to correctly remember one's prior expectations after observing new information, hinders information processing by decision makers in financial markets. Biased agents are not surprised by new information, as "they knew it all along." Starting with Fischhoff (1975) and Fischhoff and Beyth (1975), this bias has been amply documented by the psychology literature.² This literature has shown that the hindsight bias arises in a large variety of contexts and that even the knowledge of this bias does not eliminate it. Markets reduce but don't eliminate this bias, as demonstrated experimentally by Camerer, Loewenstein and Weber (1989).

The goal of this paper is to study the consequences of the hindsight bias for an important type of economic activity: investment and trading. To provide a theoretical framework for our investigation, we model information processing by hindsight biased agents. We show that these agents overreact to new information. The intuition is the following. Rational agents form conditional means by averaging their prior expectations with their signals. Hindsight biased agents, when trying to form such averages, incorrectly remember their prior expectations. They reconstruct biased priors, contaminated by their new knowledge. This leads to a form of double counting of their signal.

Using this model of biased information processing, we analyze theoretically financial decision making by hindsight biased agents. When choosing portfolios, these agents rely on biased expected returns, exaggerating the information content of recent signals. Also, hindsight biased agents fail to assess variances correctly. When volatility is stochastic, traders need to update their assessment of risk, based on return realizations. On observing unexpectedly positive or negative returns, rational agents raise their volatility estimates. Hindsight biased agents, who "knew it all along," fail to understand that such returns

²For an extended review of the literature and a discussion of psychological explanations of the bias see Christensen-Szalanski and Wilham (1991), Hawkins and Hastie (1990) and Roesse (2004).

were unexpected, and thus underestimate variances. The conclusion of our theoretical analysis is therefore that hindsight biased agents, because they form incorrect conditional expectations and variances, will choose inefficient portfolios and achieve inferior financial performance.

To test this hypothesis, we collect data from traders working for a large international investment bank. To assess their hindsight bias, we use the same psychometric approach as in Camerer, Loewenstein and Weber (1989). We form two groups, A and B. We ask group B participants to form assessments about a random variable. Then we ask group A participants to predict the average group B answer, *after telling them the realization of the random variable*. Hindsight biased agents in group A are excessively influenced by this new information. They offer inaccurate predictions of group B estimates, too close to the realization. In this context, Camerer, Loewenstein and Weber (1989) propose an index to measure the hindsight bias via the extent to which the participants tilt their predictions towards the true answer. We show that, within our theoretical framework, this index measures precisely the hindsight bias parameter of the agents.

To estimate their hindsight bias, we collected psychometric data from 41 bankers in Frankfurt and 49 in London, all working on issues relating to trading and investment, for the same large international investment bank. We found these bankers to be significantly hindsight biased. And yet the questions we asked them related to their own field of expertise, such as stock market returns, macro variables or characteristics of the investment banking industry. We also found that more experienced bankers or bankers with more precise information were not less biased.

To examine if this bias affected economic performance, we asked the bank for data about the performance of the respondents. The investment bank sorted them into three categories, according to their overall compensation (including wage and bonus.) We used this categorization as a proxy for performance. We related this proxy to our measure of the hindsight bias. Consistent with the hypothesis that the hindsight bias reduces performance, we found that the bankers in the highest earnings category had the lowest bias on average. We first obtained this finding in the Frankfurt subsample. Then, to test its robustness, we collected additional data for London. In this new sample we found the same result. To evaluate the statistical significance of the result, we performed a non-parametric Wilcoxon rank test. We found that high earning bankers were significantly less biased than low-earnings bankers and than mid-earnings bankers. The difference between low- and mid-earnings bankers was not significant, however.

To further document the link between the hindsight bias and performance, we computed a contingency table, with 9 cells corresponding to the three earnings categories and a categorization of the bankers in 3 hindsight bias groups of equal size. The cells of this table are empirical estimates of the probability that the banker is in one of the three performance group (high earnings, medium earnings or low earnings), conditional on his bias (high bias, medium bias or low bias). If earnings and bias were independent, conditional probabilities should be equal to unconditional ones, and each cell in the matrix should take the value $1/3$. In contrast with this null hypothesis, we found that the diagonal terms in the matrix were the highest. Bankers with low bias, are most likely to be in the high earnings categories. The probability that they are is 46%, clearly above the 33% probability expected under the null hypothesis. Similarly, bankers with medium (resp. high) bias are most likely to be in the medium (resp. low) earnings category. The corresponding

conditional probabilities were estimated to 39% for the medium category and 48% for the high-bias category. A Chi-square test confirmed the statistical significance of these results.

One might wonder if the link between performance and bias could be spurious, and reflect other characteristics, such as experience, type of job or information. To test the robustness of our findings with respect to experience, we split the bankers in two sub-groups: the first one included bankers with less than 10 years of experience, the other one bankers with 10 years or more. We then computed the average hindsight bias of bankers with different earnings and different experiences. We found that the link between bias and performance was still prevalent after controlling for experience. Both for experienced bankers and for more junior ones, the high earnings bankers were found to be less biased than the medium earnings and low earnings bankers. We also found the link between performance and bias to be robust to differences in jobs and information.

Our paper complements three recent streams of literature in behavioral finance and economics.³ First, as discussed in Harrison and List (2004), experimental economists increasingly rely on field data collected from practitioners, rather than from students. Thus, Haigh and List (2005) experimentally find that myopic loss aversion is even stronger for CBOT traders than students. Sarin and Weber (1993) find that ambiguity aversion reflected in experimental market prices is also prevalent among investment bankers. Harrison, Lau and Williams (2002) estimate time preferences using field data from the Danish population. Second, our economic model of biased agents is in the line of the behavioral finance models of Daniel, Hirshleifer and Subrahmanyam (1998), Hirshleifer (2001), Cao, Coval and Hirshleifer (2002), and Rabin (1999 and 2002). Third, we empirically relate psychological constructs to economic variables, as Camerer (1987), Fenton O’Creevy (2003), Biais, Hilton, Mazurier and Pouget (2005), and Glaser, Langer and Weber (2005).

In the next section we first survey the psychology literature, on the hindsight bias. In Section 3, we present our theoretical model. In Section 4, we present our data and empirical results. Section 5 concludes. Proofs omitted in the text are in the appendix.

2 Literature

The research on hindsight bias started with Fischhoff (1975) and Fischhoff and Beyth (1975).⁴ Let \tilde{v} be a random variable and I_0 the decision maker’s information at some point in time t_0 . Let t_1 be a point in time, when the random event has been resolved. At time t_1 the decision maker possesses information I_1 , including the original information I_0 as well as the outcome of the random variable. The ex-ante expectation of the random variable is : $E(\tilde{v}|I_0)$. The ex-post recollection of this ex-ante expectation is: $E[E(\tilde{v}|I_0)|I_1]$. For a rational decision maker, both expectations have to be equal. In contrast, hindsight biased agents fail to remember the initial expectation. They forget they knew only I_0 , they feel they already knew I_1 (or something close to I_1) at time 0. Wasserman et al (1991, p 30) describe this phenomenon as “a projection of new knowledge into the past accompanied by a denial that the outcome information has influenced judgements.” Correspondingly,

³Camerer (2005) offers an interesting survey of behavioral economics.

⁴For a review of the literature and a discussion of psychological explanations of the bias see Fischhoff (1982), Christensen-Szalanski and Wilham (1991), Hawkins and Hastie (1990) and Roesse (2004).

for hindsight biased agents the ex–post recollection of the initial belief will be closer to the realization than the true ex–ante expectation.

The notion of hindsight bias was initially developed in the context of binary variables: $v \in \{0, 1\}$. In that case, the expectation is the probability that the variable takes the value one. Hindsight bias arises if the ex–post recollection of the ex–ante probability is greater when the event actually occurred. The bias arising in the general case, where \tilde{v} is not necessarily binary, is often referred to as the “curse of knowledge.” In the present paper we will not differentiate curse of knowledge and hindsight bias.

Numerous papers in experimental psychology have shown the prevalence of the hindsight bias and discussed its measurement (see e.g., Hawkins and Hastie, 1990.) Musch (2003) finds empirically that individual differences in hindsight bias exist. This suggests that this bias acts as a trait, consistently influencing individual behavior in various environments.

Some papers have also provided insights into the consequences of the hindsight bias. Fischhoff (1982) argues that the hindsight bias will prevent one from rejecting one’s hypotheses about the world.⁵ Buksar and Conolly (1988) study advanced strategic planning students analyzing business cases. They find that, in this context, the hindsight bias hinders learning from past experience. Mangelsdorff and Weber (1998) show that, in a principal agent relation, the hindsight bias will prevent the principal from correctly evaluating the performance of the agent.⁶ Indeed, biased principals fail to remember what was known when the agent’s decision was taken.⁷ Holzl et al (2002) offer evidence of hindsight bias about the economic advantages of the European Monetary Union.

Three main different empirical designs have been used to demonstrate the hindsight bias.

1. In a within person design, subjects are first asked to report their ex–ante expectations. Then, they learn the realization of the variable. Then they are asked to report their ex–post recollection of their ex–ante expectations. Fischhoff and Beyth (1975) provide evidence of hindsight–bias in this context.
2. In a between subjects design, subjects each have to report their ex–ante expectation of an event. Two groups are formed. In group one, participants receive no information. In group two, participants are told the true outcome of the event, and yet are asked to report their ex–ante expectation. Fischhoff (1975) offers evidence of hindsight bias in this context.
3. In the third design also two groups are formed. In group two, subjects are told the true outcome of the event and asked to estimate the average expectation of group one (knowing that that group had no information). Camerer, Loewenstein

⁵“When we attempt to understand past events, we implicitly test the hypotheses or rules we use both to interpret and anticipate the world around us. If, in hindsight, we systematically underestimate the surprises that the past held and holds for us, we are subjecting those hypotheses to inordinately weak tests and, presumably, finding little reason to change them.” (Fischhoff, 1982, page 343.)

⁶See also Camerer (2005).

⁷This is in line with the finding of Baron and Hershey (1988). They asked subjects to evaluate decisions. They found that subjects rated the decision maker better when the outcome was favorable than when it was not.

and Weber (1989) use this approach, which we will revisit below in the context of our model.

Bias in design 1) could reflect memory effects. In design 1) and in 2), bias could arise from a person's desire to maintain high levels of public esteem (see e.g., Campbell and Tesser, 1983). In our study, we will use design 3) where these effects should not arise. In addition, as Rabin (1998) points out, as economists, we care mostly about a person's belief about other people not about herself. Design 3) captures this aspect.

3 Theory

3.1 Modelling the hindsight bias

We now present a statistical model of the hindsight bias, in line with Camerer, Loewenstein and Weber (1989). This provides a modelling framework to analyze the consequences of the hindsight bias for economic decision making and social interactions.

3.1.1 Mean

Consider a random variable, denoted by \tilde{v} . Denote its a priori mean by μ . Suppose the agent has been told the realization v of \tilde{v} and is asked to report what was his a priori belief about the mean of \tilde{v} . If rational, he reports $E(\tilde{v}) = \mu$. If hindsight biased, he fails to remember how ignorant he was initially about \tilde{v} . By a process of memory reconstruction he incorporates the new knowledge into his remembrance of the prior expectation (see e.g., Carli, 1999.) Thus he forms a biased recollection of the prior mean, tilted towards the realization of the random variable. We model this by specifying the reconstructed prior mean of the biased agent as a weighted average of the true prior mean and of the realization of the random variable:

$$(1) \quad \hat{E}_v(\tilde{v}) = \omega v + (1 - \omega)\mu = \mu + \Delta\mu(v)$$

where $\Delta\mu(v) = \omega(v - \mu)$, \hat{E}_v denotes the biased expectation, the subscript denotes that the biased expectation is reported after observing the realization v of \tilde{v} , and the constant ω measures the magnitude of the bias. This formulation is similar to equation (1) in Camerer, Loewenstein and Weber (1989). When asked to report the mean of the distribution, the biased agents partially forget the ex-ante expectation. Rational agents are not influenced by the realization of the random variable. They fully remember the prior mean, and thus place no weight on the realization ($\omega = 0$). In contrast, completely biased agents totally forget the rational prior mean and place all the weight on the realization of the random variable ($\omega = 1$).

3.1.2 Density

To model the hindsight bias as simply as possible, we assume it only affects the ex-post reconstruction of the a priori mean of the distribution. We assume it does not affect the perception by the agent of the other moments of the prior distribution. Thus, once they are told the realization v of a random variable \tilde{v} , hindsight biased agents simply shift its

distribution by $\Delta\mu(v)$. Denote by $f(\cdot)$ the true a priori density of \tilde{v} . The hindsight biased density is: $\hat{f}_v(x) = f(x - \Delta\mu(v))$.

From the perspective of hindsight biased agents having observed the realization of \tilde{v} , the reconstructed a priori expectation of this variable is:

$$\hat{E}_v(\tilde{v}) = \int_{-\infty}^{\infty} \hat{f}_v(x) x dx = \int_{-\infty}^{\infty} f(x - \Delta\mu(v)) x dx.$$

Denote: $t = x - \Delta\mu(v)$. With this change of variable, the hindsight biased expectation becomes:

$$\hat{E}_v(\tilde{v}) = \int_{-\infty}^{\infty} f(t)(t + \Delta\mu(v)) dt = \mu + \Delta\mu(v),$$

as in equation (1).

Consider for example the case of a normal distribution, with true density:

$$f(x) = \frac{1}{(2\pi)^{1/2}\sigma} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}.$$

The hindsight biased density is:

$$\hat{f}_v(x) = f(x - \Delta\mu(v)) = \frac{1}{(2\pi)^{1/2}\sigma} e^{-\frac{1}{2}\left(\frac{x-(\mu+\Delta\mu(v))}{\sigma}\right)^2}.$$

3.1.3 Updating

Suppose the agent observes a signal \tilde{s} on the variable \tilde{v} :⁸

$$\tilde{s} = \tilde{v} + \tilde{\epsilon},$$

where all the variables are jointly normal. $E(\tilde{\epsilon}) = 0$, which implies that $E(\tilde{v}) = E(\tilde{s}) = \mu$. Rational agents update their expectation of \tilde{v} using the projection theorem:

$$(2) \quad E(\tilde{v}|\tilde{s} = s) = E(\tilde{v}) + \frac{\text{cov}(\tilde{v}, \tilde{s})}{\text{var}(\tilde{s})}(s - E(\tilde{s})) = \delta s + (1 - \delta)\mu,$$

where

$$\delta = \frac{\text{cov}(\tilde{v}, \tilde{s})}{\text{var}(\tilde{s})}.$$

Hindsight biased agents will use a different updating formula, because they fail to remember correctly their prior expectation. They will use a biased reconstruction of the prior mean, contaminated by this new knowledge. To model this, we assume that the agent's reconstruction of the a priori mean is a weighted average of the true prior mean (μ) and the expectation conditional on the new information. When the agent exactly observes the realization of the random variable, the conditional mean is v , and we obtain equation (1). When the agent only observes a signal s , the conditional expectation is as given in equation (2), and the reconstructed prior mean is:

$$(3) \quad \hat{E}_s(\tilde{v}) = \omega[\delta s + (1 - \delta)\mu] + (1 - \omega)\mu = \omega\delta s + [\omega(1 - \delta) + (1 - \omega)]\mu,$$

⁸At this stage, since we only consider a single person context, it does not matter if the signal is private or public. Later in the paper we analyze the two cases in turn.

where, similarly to equation (1), the subscript denotes that the biased expectation is reported after observing the realization s of \tilde{s} . Comparing equations (3) and (1), we see that perfect signals influence agents more than noisy signals. In the limit, when the signal is pure noise and δ goes to 0, the agent pays no attention to it, and it does not bias the agent's remembrance of the prior mean.

On the other hand, since the bias only amounts to a shift in the distribution, it does not affect the variance or the covariance. Thus, the hindsight biased version of the projection theorem is:

$$\hat{E}_s(\tilde{v}|\tilde{s} = s) = \hat{E}_s(\tilde{v}) + \frac{\text{cov}(\tilde{v}, \tilde{s})}{\text{var}(\tilde{s})}(s - \hat{E}_s(\tilde{s})),$$

where the covariance and the variance are the same as for the rational agent. After simple manipulations, we obtain the following lemma:

Lemma 1: *After observing signal s , the hindsight biased agent, with parameter ω , forms the following biased conditional expectation:*

$$\hat{E}_s(\tilde{v}|\tilde{s} = s) = \delta(1 + (1 - \delta)\omega)s + (1 - \delta)[1 - \omega\delta]\mu.$$

The difference between the weight placed on s by the hindsight biased agent and that placed by the rational agent is:

$$\delta(1 + (1 - \delta)\omega) - \delta = \omega\delta(1 - \delta) \geq 0.$$

Hence, we can state our next lemma.

Lemma 2: *Hindsight biased agents overreact to private signals, i.e., they put more weight on them than Bayesian rationality would entail.*

The interpretation of the lemma is the following: Hindsight biased agents engage in a form of double counting. They use the signal to update their expectation, like rational agents, but they start from a biased prior, also affected by the signal.

Put in other words, Lemma 2 states that hindsight biased traders overestimate the informativeness of the signal \tilde{s} . One might wonder if this could result from other cognitive biases. Take for example miscalibration. Miscalibrated agents underestimate variances. In the present context, there are two relevant random variables for the agent: \tilde{v} and $\tilde{\epsilon}$. If the agent underestimates the variance of $\tilde{\epsilon}$, but not that of \tilde{v} , he will overestimate the informativeness of the signal, similarly to a hindsight biased agent. But if the agent equally underestimates the two variances, then miscalibration does not lead to overreaction.

3.2 Measuring the hindsight bias

As noted above, Camerer, Loewenstein and Weber (1989) measure the hindsight bias using a “between subjects design.” They form two groups (A and B) and ask a question to the participants in group B . Denote the true answer to this question by v . Denote by W_B the average answer of group B participants. The experimenter gives the answer v to participants in group A . Then the experimenter asks these participants to assess the average answer of group B .

Denote by $\hat{E}_v^a(W_B)$ the assessment of W_B by participant a in group A . The index of hindsight bias used by Camerer, Loewenstein and Weber (1989) is:

$$(4) \quad h^a = \frac{\hat{E}_v^a(W_B) - W_B}{v - W_B},$$

i.e., the difference between a 's assessment of the average group B 's answer minus the actual average, divided by the difference between the true answer and group B 's average answer. As noted by Camerer, Loewenstein and Weber (1989) rational agents should on average correctly estimate W_B . Hence, for such agents, on average $h^a = 0$. In contrast, hindsight biased agents will be influenced by the revelation of v . If they are completely biased, they will totally forget the answer was difficult to find. They will believe the others gave the exact right answer: v . Thus, for completely biased agents, $h^a = 1$.

As shown below, this reasoning is confirmed in our statistical framework. Before they are told the true answer by the experimenter, the participants don't know it. For them it is a random variable, which we denote \tilde{v} . Assume there is a common prior about this random variable, and denote its mean by: $E(\tilde{v}) = \mu$. We first analyze the case where the agents have no private signals about \tilde{v} , and thus simply work from the common prior. Second we consider the case where there are private signals.

3.2.1 When there are no private signals

With a common prior and without private signals, all participants in group B should give the same answer: μ . Hence, $W_B = \mu$. As stated in equation (1), once told the true answer, participant a will reconstruct a biased expectation of \tilde{v} . Rationally anticipating that group B participants should answer the common prior, but incorrectly believing that this prior is the expectation given in equation (1), a will state:

$$\hat{E}_v^a(W_B) = \omega^a v + (1 - \omega^a)\mu,$$

where ω^a is the hindsight bias parameter of agent a in A . Substituting W_B and $\hat{E}_v^a(W_B)$ in equation (4), we obtain our next lemma.

Lemma 3: *In our framework, if participants don't observe any private signals, the Camerer, Loewenstein and Weber (1989) index for each participant (h^a) is equal to his (her) hindsight bias parameter (ω^a).*

3.2.2 When there are private signals

It is likely, however, that participants have also observed private signals, which will influence their answers. Hence, we extend our analysis to study the following questions: Could it be that tilting one's answer towards v is a rational reaction to the presence of information asymmetries? Does this imply that, in this context, the Camerer, Loewenstein and Weber (1989) index no longer captures the hindsight bias of the participants?

Assume each participant b in group B has observed a private signal: $\tilde{s}_b = \tilde{v} + \tilde{\epsilon}_b$, where $E(\tilde{\epsilon}_b) = 0$. For simplicity assume all signals have the same precision, all random variables are jointly normal, and conditionally on the realization of \tilde{v} , the participants' private signals are independent (i.e., $\tilde{\epsilon}_b$ independent from $\tilde{\epsilon}_{b'}$ for $b \neq b'$).

Each participant i is hindsight biased with parameter ω^i . While we assume each participant is unaware of his or her own bias (otherwise his reasoning would be inconsistent), we assume that, in all other respects, agents are rational. Hence, they rationally anticipate that the others are biased. Participants do not necessarily know the exact biases of the others, however. Thus, for each participant the bias parameters of the others are random variables. We assume that participants rationally anticipate that these bias parameters are i.i.d, with expectation ω , and independent from the other random variables.

In this context, the average answer of group B is characterized in the next lemma:

Lemma 4: *In our framework, when participants observe private signals, as the number of participants goes to infinity, the average answer of participants in group B goes to:*

$$W_B = (\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)E(\tilde{v}).$$

We hereafter focus on the case where the number of participants is large, so that the law of large numbers is a good approximation. Using Lemma 4 we can now characterize the answer of participant a in group A about the average answer of group B. If a was rational, his estimate would be (up to our approximation):

$$\begin{aligned} E(W_B|v) &= W_B = (\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)E(\tilde{v}) \\ &= (\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)\mu. \end{aligned}$$

This would indeed involve some tilting away from the unconditional mean (μ) towards the realization of the random variable (v), to take into account the fact that group B participants have observed private signals.

But what if participant a is hindsight biased? In that case, a fails to remember properly $E(\tilde{v}) = \mu$, and instead reconstructs it as: $\hat{E}_v(\tilde{v}) = \omega_a v + (1 - \omega_a)\mu$. Using this biased reconstruction and Lemma 4, the answer of a is readily obtained, and correspondingly the value of the Camerer, Loewenstein and Weber (1989) index can be computed. It is given in the next proposition:

Proposition 1: *In our framework, even when participants observe private signals, as the number of participants goes to infinity, the Camerer, Loewenstein and Weber (1989) index for each participant (h^a) goes to his or her hindsight bias parameter (ω^a).*

When predicting the average answer of group B participants, to account for their private signals, one needs to compute a weighted average of the unconditional mean and the realization of the random variable. Hindsight biased agents, however, incorrectly remember the unconditional mean. Their reconstruction is contaminated by their observation of v . Hence their prediction of the answer of group B participants is too close to v . This is captured by the Camerer, Weber and Loewenstein (1989) index (hereafter referred to as the CLW index).

One might wonder if our result that the Camerer, Loewenstein and Weber (1989) index captures the hindsight bias is due to the simple informational structure we consider. Would it be possible to design a more complicated information structure in which the answers of fully rational agents would give rise to positive values for the Camerer,

Loewenstein and Weber index? The general answer is no. With rational agents, the Camerer, Loewenstein and Weber index must on average equal 0, since its numerator is equal to the difference between a variable (W_B) and its expectation ($E(W_B|v)$). To illustrate this point with an example, relax the assumption that agents observe independent signals and consider correlated signals: $s_i = v + \theta + \epsilon_i$, where $E(\theta) = E(\epsilon_i) = 0$ and v , θ and the ϵ_i are independent. In that parametrization, agents make a common mistake: θ and also individual mistakes ϵ_i . In the appendix we prove that, even with such correlated errors, the Camerer, Loewenstein and Weber (1989) index will be 0 on average if all agents are rational. So experimental observation of average positive index values is not consistent with rationality.

3.3 Adverse consequences of the hindsight bias in financial markets

We now analyze, in the context of our model of the hindsight bias, two important financial tasks: learning about variances and investment choices.

3.3.1 The hindsight bias reduces the performance of investment strategies

Consider a standard multi-asset mean variance economy, à la CAPM. There are N assets, with final payoff \tilde{v} . The random vector \tilde{v} is jointly normal with mean μ . There is a continuum of competitive agents, A . For simplicity assume they all have CARA utility with the same risk aversion coefficient γ .

Three dates must be considered in the analysis. At time 0, all agents observe a public signal \tilde{s} on the variable \tilde{v} . Assume $\tilde{s} = \tilde{v} + \tilde{\epsilon}$, all the variables are jointly normal and $E(\tilde{\epsilon}) = 0$. Denote by τ the inverse of variance covariance matrix of the asset payoffs conditional on s , i.e., the conditional precision. Because the variables are normal, conditional precision is constant.

At time 0, each agent a also receives an endowment z_a in the risky assets and C_a in cash. In the simple framework we consider, we don't need to make any particular assumption about the distribution of the shocks z_a . Denote by z_m the aggregate endowment shock:

$$z_m = \int_{a \in A} z_a da,$$

In addition to the individual endowments, there is also an exogenous endowment x of the asset, with inelastic supply. At time 1 trading takes place. Denote by q_a the vector of asset demands by agent a . Finally at time 2 the value of the asset is realized and consumption takes place. The risk free rate between time 1 and time 2 is denoted by r . Agents differ with respect to their endowments, but, as will be obvious below, this does not affect the results qualitatively. More importantly, we assume the agents differ in terms of rationality. To be precise, we assume that each agent a is hindsight biased with bias parameter $\omega^a \in [0, 1]$. While this bias will affect the perception of expectations by agents, it does not affect the other moments of the distribution.

The final wealth of agent a at time 2 can be written as:

$$W_a = (q_a + z_a)'v + (C_a - q_a'p_a)(1 + r),$$

which, conditionally on the signal, is normal. Thus, with CARA utilities and jointly normal distributions, the program of agent a at time t is:

$$\text{Max}_{q_a} \hat{E}_s[W_a|s] - \frac{\gamma}{2} \text{Var}[W_a|s].$$

First order conditions and market clearing yield the following lemma:

Lemma 5: *In our framework the demand of agent a is:*

$$(5) \quad q_a = \frac{\tau}{\gamma} \left\{ \hat{\delta}^a s + (1 - \hat{\delta}^a) \mu - p(1 + r) \right\} - z_a,$$

where:

$$\hat{\delta}^a = \delta(1 + (1 - \delta)\omega^a).$$

and the equilibrium price is:

$$(6) \quad p = \frac{1}{1 + r} \left\{ \hat{\delta}^m (s - \mu) + \mu - \gamma \tau^{-1} (x + z_m) \right\}.$$

The second and third terms in equation (6) are the standard expectation and risk premium terms arising in this CAPM like economy. The first term reflects how biases of the agents affect equilibrium prices through the demand of the agents. As in the standard model, the demand of agent a reflects his endowment shock z_a , as agents seek to trade away from their undiversified endowments, to hold more balanced portfolios. The distinctive feature of this analysis is that the demand of agent a also reflects his hindsight bias, which affects the way in which the agent interprets the public signal.

Substituting the equilibrium price into the agent's demand function we obtain his equilibrium holdings:

$$(7) \quad q_a + z_a = (x + z_m) + \frac{\tau}{\gamma} \delta(1 - \delta)(\omega^a - \omega^m)(s - \mu).$$

Contrast these with the optimal holdings for agent a , that is the holdings he would choose if he rationally formed his beliefs:

$$(8) \quad (x + z_m) - \frac{\tau}{\gamma} \delta(1 - \delta)\omega^m(s - \mu).$$

These optimal holdings are equal to the market portfolio, which the agent would hold if everybody was rational, corrected by a term aimed at exploiting the mispricing due to the irrationality of the others. Comparing equations (7) and (8), we obtain the following proposition:

Proposition 2: *In our framework, the actual holdings of agent a differ from those which, given prices and information, would be optimal for him. The wedge between actual and optimal holdings is:*

$$\frac{\tau}{\gamma} \delta(1 - \delta)\omega^a(s - \mu),$$

i.e., the biased agent puts too much weight on the assets for which the public signal was good. The greater the agent's bias, the greater the deviation from optimal holdings, and the lower the investment performance of the agent.

The greater the agent's bias, the more he will deviate from optimal holdings. This deviation will lead him to overweight the assets for which the signal was good, and underweight the assets for which the signal was bad. This will reduce his investment performance.

3.3.2 The hindsight bias hinders learning about variances

Risk is one of the most important elements of the environment of financial decisions. Correctly factoring risk in these decisions is made difficult by the fact that volatility fluctuates randomly. Because of these fluctuations, it is important for agents to conduct efficient learning about volatility. In this subsection we show that the hindsight bias reduces the ability of agents to accurately learn about risk, and leads to variance underestimation.

Assume that, with probability λ the random variable \tilde{v} is normally distributed with variance $\underline{\sigma}^2$. With probability $1 - \lambda$ it is normal with variance $\bar{\sigma}^2$, where $\bar{\sigma} > \underline{\sigma}$. After observing one realization of \tilde{v} , the agent has to update her beliefs about the true variance. Applying Bayes law, rational agents would update the probability that the variance is low to:

$$\begin{aligned} P(\underline{\sigma}|v) &= \frac{f(v|\underline{\sigma})\lambda}{f(v|\underline{\sigma})\lambda + f(v|\bar{\sigma})(1-\lambda)} \\ &= \frac{\frac{\lambda}{\underline{\sigma}} e^{-\frac{1}{2}\left(\frac{v-E(\tilde{v})}{\underline{\sigma}}\right)^2}}{\frac{\lambda}{\underline{\sigma}} e^{-\frac{1}{2}\left(\frac{v-E(\tilde{v})}{\underline{\sigma}}\right)^2} + \frac{1-\lambda}{\bar{\sigma}} e^{-\frac{1}{2}\left(\frac{v-E(\tilde{v})}{\bar{\sigma}}\right)^2}} \\ &= \frac{1}{1 + \frac{1-\lambda}{\lambda} \frac{\underline{\sigma}}{\bar{\sigma}} e^{-\frac{1}{2}\left(\frac{1}{\bar{\sigma}^2} - \frac{1}{\underline{\sigma}^2}\right)(v-E(\tilde{v}))^2}}. \end{aligned}$$

Hence:

$$(9) \quad P(\underline{\sigma}|v) = \frac{1}{1 + \frac{1-\lambda}{\lambda} \frac{\underline{\sigma}}{\bar{\sigma}} e^{-\frac{1}{2}\left(\frac{1}{\bar{\sigma}^2} - \frac{1}{\underline{\sigma}^2}\right)(v-\mu)^2}}.$$

Hindsight biased agents will proceed to a similar updating, except that, instead of using $E(\tilde{v}) = \mu$, they will rely on the biased expectation: $\hat{E}_v(\tilde{v}) = \omega v + (1 - \omega)\mu$. This leads to the following biased probability:

$$\hat{P}_v(\underline{\sigma}|v) = \frac{1}{1 + \frac{1-\lambda}{\lambda} \frac{\underline{\sigma}}{\bar{\sigma}} e^{-\frac{1}{2}\left(\frac{1}{\bar{\sigma}^2} - \frac{1}{\underline{\sigma}^2}\right)(v-\hat{E}_v(\tilde{v}))^2}},$$

where, as in the case of expectations or densities, the hat denotes that the belief is hindsight biased, and the subscript v denotes that the belief is formed after observing v .

Substituting the biased expectation, we obtain:

$$(10) \quad \hat{P}_v(\underline{\sigma}|v) = \frac{1}{1 + \frac{1-\lambda}{\lambda} \frac{\underline{\sigma}}{\bar{\sigma}} e^{-\frac{1}{2}\left(\frac{1}{\bar{\sigma}^2} - \frac{1}{\underline{\sigma}^2}\right)(1-\omega)^2(v-\mu)^2}}.$$

Comparing (9) and (10), we obtain our next proposition:

Proposition 3: *In our framework, when agents must conduct learning about volatility, those with hindsight bias underestimate volatility.*

The interpretation of this result is the following. Agents increase the probability that variance is large when they observe realizations that differ a lot from their prior expectations, i.e., when they are surprised. Hindsight biased agents tend to not be surprised: “They knew it all along.” Hence they underestimate volatility.

Such underestimation will undermine the quality of their decisions: They will underestimate the value of options and hedging strategies. They will inaccurately assess risk return tradeoffs. They will incorrectly estimate risk premia. Hence, as in the previous subsection, we conclude that the hindsight bias reduces performance.

4 Empirical analysis

4.1 Data Collection

The psychometric data used in the present study was collected as part of a larger seven page questionnaire where other psychological constructs were also measured. These variables, together with similar ones collected from another bank are analysed in Glaser, Langer and Weber (2005). Our subject pool consists of bankers of a large European Bank. The participants in the study are from the trading department of the bank. While the participants have various tasks (sales, research and trading), they all participate in the elaboration of portfolio allocation, trading and investment decisions.

The data was collected by inviting groups of participants into one of two conference rooms in the bank. There, they filled out the questionnaire under the supervision of an experimenter. After finishing the questionnaire, which took about 30 minutes, participants were asked not to talk to their colleagues until the end of the data collection. The overall data collection effort took about two hours.

The data was collected anonymously with each questionnaire being labelled with a number. This number was used to sort the questionnaires into earnings categories. For each number the personnel department of the bank informed us whether the respondent was in the top, middle or low earning category. Earnings meant overall compensation, including bonus. Sorting subjects into three categories was a compromise between our wishes (as much performance data as possible) and the privacy requirements of the participants and institution. It should be noted that the distribution of earnings is skewed (as suggested by common wisdom and an interview with the human resource department). Therefore earnings in the low earnings and medium classes are less variable than in the high earnings class.

We first collected the data in the Frankfurt branch of the bank, on December 5, 2001. We had 41 respondents (each member of the department who was present that day). Then, to assess the robustness of our results, we collected similar data, in the London branch of the bank, on October 9, 2002. There we had 49 respondents, corresponding to the vast majority of the department.⁹ Bankers in both locations had similar jobs. In Frankfurt, the questionnaire had to be approved by the management and by the worker’s

⁹In London we received 49 questionnaires. However, only 44 participants were grouped into earnings categories. Hence, as regards the London subsample, the test of hypotheses is based on 44 subjects whereas the pure hindsight bias results are based on 49 subjects.

council. We made it very explicit that the study was done by two universities and that the management of the bank would not see any specific questionnaire. The only data made available to participants and management were aggregate data. No payments were given.

4.2 Descriptive statistics

Table 1 offers some descriptive statistics on the population of bankers in our sample. It shows that low and mid earnings bankers have approximately the same age in Frankfurt and London. High earnings bankers are considerably older in London. In Frankfurt the bankers are relatively equally distributed across the different departments (sales, trading, research and other). Unfortunately, the London traders did not fill in our questionnaire. Hence, the bankers in our London subsample only belong to sales and research.

As explained in Section 2, we chose the same experimental design as Camerer, Loewenstein and Weber (1989) to estimate the hindsight biases of the investment bankers. This involved two types of tasks for each participant. In the first task, traders were asked to estimate the true value of some unknown variables (own estimate.) In the second task, traders had to guess the average answers of the other group to some questions (others estimate.) In that second task, the participants were given the true answer to the question asked to the other group. Each subject had to answer ten questions, five in the own estimate condition and five in the others estimate condition. We formed two groups in each location, hereafter referred to as Group A and B.

The items we used to measure hindsight bias were taken from the natural habitat of the bankers. Nevertheless, we used questions for which we thought they did not know exactly the correct answer. The items were similar for Frankfurt and London. Whenever it was reasonable, the questions were exactly the same for the two locations. The second sample however, was collected 10 months later than the first and in a different country. Hence, we had to change some questions to ensure comparability. For example, while in the first sample we had a question about the German consumer price index, in the second sample we changed it to a question about the change in UK retail price index. The questions we used to measure the hindsight bias as well as the average answers are listed in Table 2 and 3. The numbers in Tables 2 and 3 are rounded to two digits. For each participant we computed a measure of the hindsight bias for each of the five question. Then we took the average of these five numbers.

The results in Tables 2 and 3 show that the participants in our study are hindsight biased on average. This finding is similar to those of previous psychometric studies (see e.g., Hawkins and Hastie 1990). Note that the bias is observed in spite of the fact that the questions are from the natural habitat of the bankers and that the participants are experienced professionals.¹⁰ The degree of hindsight bias is different for different questions but between 0 and 1 in the majority of the cases. Table 4 documents the mean and the median of hindsight bias of the participants, aggregated across the 5 questions. The average hindsight bias is very similar in London (.547) and in Frankfurt (.579). It is

¹⁰We also collected answers to the same questionnaire from economics and management students in Mannheim University. The mean hindsight bias of the students (.92) is greater than its counterpart for bankers. But this mean is driven by a few outliers with large bias. The median hindsight bias for students (.25) is lower than its counterpart for bankers.

significantly lower than 1 and above 0 in both locations. The t-statistics are: 3.60 and 4.35 respectively for London, and 3.43 and 4.72 respectively for Frankfurt.

4.3 Measurement and psychometric issues

An important psychometric issue is whether our measure of the hindsight bias really captures this bias, or spuriously reflect another trait. One possibility would be that our measure of the hindsight bias would simply mirror the precision of the knowledge of the participants regarding the questions we asked them. To check this we estimated a measure of this precision. In the own estimate treatment, for each participant and each question, we computed the absolute value of the difference between the participant’s answer and the true answer. Then we summed these mistakes across questions. Finally, we normalized this total mistake of each participant by the average mistake of the participants in his group. The greater this ratio, which we herafter refer to as the mistake ratio, the lower the precision of the information of the participant regarding the questions we asked him. To examine if the hindsight bias is related to the mistake ratio, we computed the correlation between the two. The Spearman correlation coefficient is -.0837. The probability to observe such a value under the null that the two variables are uncorrelated is 44.6%. Hence, the mistake ratio and the hindsight bias are weakly and insignificantly negatively correlated. Thus, if we find that our measure of the hindsight bias predicts performance, it cannot be because it proxies for information precision.¹¹

Another possibility could be that our measure of the hindsight bias would actually reflect underconfidence. To see this, consider the case where the agents are fully rational as regards their first order beliefs (about v) but can have biased second order beliefs and overestimate the precision of signals of the others. This would be a form of underconfidence, as the agent would believe that his own rationally anticipated precision, measured by δ , could be lower than that of the others, irrationally expected to be: $\hat{\delta} > \delta$. An underconfident agent would believe that the average answer of the agents in the other group would be:

$$\hat{\delta}v + (1 - \hat{\delta})\mu,$$

and his Camerer, Loewenstein and Weber (1989) index would be:

$$\frac{(\hat{\delta}v + (1 - \hat{\delta})\mu) - (\delta v + (1 - \delta)\mu)}{v - (\delta v + (1 - \delta)\mu)} = \frac{\hat{\delta} - \delta}{1 - \delta} \geq 0.$$

Thus, *in theory*, underconfidence could generate positive values for the Camerer, Loewenstein and Weber (1989) index. *In practice*, however, this is unlikely to be driving the results in our experiment. For the participants in that experiment, Glaser, Langer and Weber (2005) also collected the answers to a “better than average” questionnaire. In that questionnaire, participants were asked to answer what percentage of their colleagues they

¹¹The measures of hindsight bias are not correlated to other biases assessed by Glaser, Langer and Weber 2005. The correlation between miscalibration (a measure of overconfidence) and hindsight bias is -.1 (London sample) and .16 (Frankfurt sample) with both values not significantly different from zero. Similarly, the correlation between the hindsight bias and the “better than average effect” (another measure of overconfidence) is not significantly different from 0.

expected to perform better than themselves along several dimensions (trading skills, communication skills, vision, technical skills.) If the agents were rational, then the average answer should be 50%. But, in the population of bankers we studied, the average answer was significantly below 50%, indicating that bankers were overconfident on average. In addition, we estimated the correlation across bankers of the overconfidence score with the hindsight bias index. If the Camerer, Loewenstein and Weber (1989) actually captured underconfidence, it should be significantly negatively to the overconfidence score. In our data the correlation was positive and not significantly different from 0. In Frankfurt the correlation was 0.055 (with a p-value of 0.733) and in London it is 0.0975 (p-value = 0.51.) Thus, it is unlikely that our measure of the hindsight bias spuriously captured underconfidence.

4.4 Empirical results

Up to that point we have demonstrated that investment bankers exhibit hindsight bias. We now want to address our main empirical question and investigate to what degree the hindsight bias is related to performance.

4.4.1 Main result

As can be seen from Table 5, the hindsight bias (median as well as mean) is lower for the high earnings category than for the two other categories. The average hindsight bias is somewhat lower for the low earnings category than for the middle earnings category. Note however that when one focuses on medians the difference is not large. Table 5 also shows that, for the middle earnings category, the median bias is much lower than the mean, while for the high earnings category, the mean bias is much lower than the median, and for the low earnings category, there is no clear pattern. This suggests that the mean in the high earnings category is driven down by a few participants with very low bias, while the mean in the middle earnings category is driven up by a few participants with very high bias.

In addition, Table 5 shows that the relation between the hindsight bias of bankers and their performance is prevalent in both locations. Both in London and in Frankfurt the bankers with the highest performance are also those with the lowest bias. Such robustness speaks to the out-of-sample validity of the result. It is particularly striking since, as was mentioned above, we first designed and implemented the experiment with Frankfurt participants, and then, to evaluate the robustness of our results, we replicated the same experiment with London participants.

Table 5 also suggests that the result is not driven by other variables, such as age or experience. Indeed, the age and experience structure differs in the two locations. While in London the high earnings bankers have longer experience than the middle earnings bankers (18 versus 9.9), in Frankfurt this is not the case (11.8 versus 12.5. See Table 1).

Tables 6 and 7 provide information about the significance of the differences in hindsight bias between the three earnings categories. In Table 6, to evaluate significance we rely on a non-parametric Wilcoxon rank test (z-stat).¹² The difference between the bias of the

¹²This type of statistics is likely to be appropriate to the extent that it is robust to noise in the data and outliers.

high earnings bankers and that of the middle earnings bankers is significant. The difference between high earners and low earners is also significant. The difference between low and mid earners is not significant. This lack of significance is consistent with the skewness of the earning distribution mentioned above. Some stars receive very high compensation, rewarding exceptional performance, while the bulk of the bankers have relatively similar compensation. Consequently, there is only a small difference between the monetary compensation of the low earnings categories and that of the middle earnings categories. This small difference is likely to be related to insignificantly different performances, and, correspondingly, biases.

Table 7 reports the estimated probabilities of the performance of the bankers conditionally on their biases. To construct this table we split the bankers in three categories in terms of bias (the number of participants was the same in the three categories). Then, we estimated conditional probabilities based on empirical frequencies. In the table, the probabilities in each column add up to one. Consistently with the hypothesis that the hindsight bias reduces performance, the largest number in each column is on the first diagonal: bankers with the lowest bias are most likely to be high-earners, bankers with medium bias are most likely to be mid-earners, and bankers with the highest bias are more likely to be low-earners. To test the null hypothesis that performance is independent from bias, we performed a Chi square test. Since there are three rows and three columns in Table 7, the relevant statistic is a Chi Square with 4 degrees of freedom, for which the critical value at the 5% level is 9.49. The distance between the empirical contingency table and the theoretical one (computed under the null) was found to be 10.05. The null hypothesis that performance and bias are independent can thus be rejected at the 5% level.

4.4.2 Robustness

Experience Experience could influence the bias of the bankers and its impact on their performance. For example, List (2003) provides evidence on the consequences of market experience on market anomalies. To control for experience we split the sample of bankers in two groups. The first group includes the bankers with less than 10 years of experience and the second group includes the bankers with 10 years of experience or more. Table 8 reports the average bias in each of these two groups. Both for the low-earnings and the high-earnings bankers, the average bias is higher for the bankers with 10 years of experience or more. This suggests that experience does not eliminate the hindsight bias.

Table 8 also offers information on the robustness of the link between bias and performance after controlling for experience. Both in the junior bankers sub-sample and in the senior bankers sub-sample, the high earnings bankers are the least hindsight biased. In fact, the most successful bankers in our sample, i.e., those who are in the highest earnings category although they are still quite young, are the less biased. The figures in the table suggest there is no interaction between the effect of hindsight bias on performance and experience.

Jobs Some participants in the experiment were employed in research, others in sales, and yet others in trading (or sales and trading). As an additional robustness check, we analyzed if our results obtained for each of these three different occupational categories.

Table 9 reports the estimated probabilities of the performance of the bankers conditionally on their biases, for the different occupational categories. As for Table 7, to compute the frequencies in Table 9, we split the bankers in three categories in terms of bias (the number of participants was the same in the three categories for each job). We classified participants as having low (resp. high bias) bias if they were among the 33% least (resp. most) biased of their occupational category. Conditional probabilities were estimated as empirical frequencies.

The results in Table 9 suggest the link between bias and performance is reasonably robust across occupational categories in the bank. Consistently with the hypothesis that the hindsight bias reduces performance, bankers with the lowest bias are most likely to be high-earners for research and for sales. Bankers with medium bias are most likely to be mid-earners, for sales and for trading. And bankers with the highest bias are more likely to be low-earners for research and for trading. In addition, for research and in trading, no banker with high bias has high earnings.

Mistakes Yet another robustness check is to examine whether our results still hold after one controls for the precision of the information of the participants, relative to the questions asked. As discussed above, for each participant, we computed a mistake ratio, decreasing with the precision of the answers of the agent in the own answer treatment. Interestingly, the mistake ratio varies somewhat with performance. Participants with high performance have an average mistake ratio of .890, while participants with medium performance have a mistake ratio of 1.004, and participants with low performance have a mistake ratio of 1.051. Yet, the correlation between the mistake ratio and the hindsight bias is not significant. The Wilcoxon rank sum test comparing the mistake ratio of high and medium earnings is $z = .549$, with a p -value of .583. The Wilcoxon rank sum test comparing the mistake ratio of high and low earnings is $z = 1.206$, with a p -value of .2278. And the Wilcoxon rank sum test comparing the mistake ratio of medium and low earnings is $z = .595$, with a p -value of .5516. Furthermore, as mentioned above, the mistake ratio is not significantly correlated with the hindsight bias.

To control for mistake ratios, we divided the population of bankers in three categories. The first category includes the 28 bankers with the lowest mistake ratios (their average mistake ratio is .425). The second category includes the 28 bankers with intermediate mistake ratios (their average mistake ratio is .765). The third category includes the 29 bankers with the highest mistake ratios (their average mistake ratio is 1.745). Table 10, presents the probability of high performance, conditional on bias, estimated for each of the three categories.

Overall Table 10 suggests some robustness of the link between bias and performance, across information precisions. Consistently with the hypothesis that the hindsight bias reduces performance, for the three categories of mistake ratios, bankers with the lowest bias are the most likely to be high-earners.

5 Conclusion

This paper suggests that cognitive biases can affect important decisions in financial markets and banks. We focus on the hindsight bias. Agents who exhibit this bias fail to

remember how ignorant they were before learning outcomes and answers. Our theoretical model predicts that the hindsight bias should hurt financial performance. To test this hypothesis we collect psychometric and performance data about highly paid investment bankers. We find that bankers exhibit hindsight bias when asked questions about economics, banking and finance, and that experience does not reduce this bias. We find that bankers with low bias obtain significantly better performance. We first obtained this result for a sample of bankers based in Frankfurt. Then, to check the out-of-sample robustness of this result, we replicated the experiment with a new set of bankers, based in London. The link between the hindsight bias and performance was found to be similar in the two locations. We also found it was robust to other variables, such as information and experience.

It would be interesting to further check the robustness of these conclusions with other experiments in the field and in the lab. Laboratory experiments could be particularly useful to provide a finer analysis of the economic consequences of the hindsight bias on cognitive processes and decisions. It would also be interesting to provide a more in depth theoretical analysis of this bias. Could it emerge from cognitive constraints, such as limited memory? We leave these issues for further research.

References

- Baron, J., and J. Hershey, 1988, "Outcome Bias in decision evaluation," *Journal of Personality and Social Psychology*, 54, 569-579.
- Biais, B., D. Hilton, K. Mazurier, and S. Pouget, 2005, "Judgemental overconfidence, self-monitoring and trading performance in an experimental financial markets," *Review of Economics Studies*.
- Bukszar, E. and T. Connolly, 1988, "Hindsight bias and strategic choice: Some problems in learning from experience," *Academy of Management Journal*, 31, 628-641.
- Camerer, C. F. (1987). "Do Biases in Probability Judgment Matter in Markets - Experimental-Evidence." *American Economic Review* 77(5): 981-997.
- Camerer, C. F., G. Loewenstein, and M. Weber. October 1989. "The Curse of Knowledge in Economic Settings: An Experimental Analysis." *Journal of Political Economy*, 97, 1232-54.
- Camerer, C. F. (2005). "Behavioral Economics." Lecture presented at the World Congress of the Econometric Society, London, 2005.
- Campbell, J. and A. Tesser, 1983, "Motivational Interpretations of hindsight bias: An individual difference analysis," *Journal of Personality*, 51, 605-620.
- Cao, H., J. Coval and D. Hirshleifer (2002), "Sidelined Investors, Trading-Generated News, and Security Returns," *Review of Financial Studies*, 15(2), (2002): 615-648.
- Carli, L., 1999, "Cognitive Reconstruction, Hindsight, and Reactions to Victims and Perpetrators," *Personality and Social Psychology Bulletin*, 25, 966-979.
- Christensen-Szalanski, J. and C. Wilham, 1991, "The hindsight bias: A meta-analysis," *Organizational Behavior and Human Decision Processes*, 48, 147-168.
- Daniel, K., D. Hirshleifer and A. Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Over- Reactions," *Journal of Finance*, pp. 1839-85.
- Fenton O'Creavy, M., N. Nicholson, E. Soane, and P. Willman, 2003, "Individual and contextual influences on the market behaviour of finance professionals," *Journal of Occupational and Organizational Psychology*.
- Fischhoff, B., 1975, "Hindsight \neq Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty." *Journal of Experimental Psychology: Human Perception and Performance*, 1. 288-299.
- Fischhoff, B., 1982, "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight." In D. Kahneman, P. Slovic, & A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, NY: Cambridge University Press, 332-351.
- Fischhoff, B., and R. Beyth, 1975, "I knew it would happen - Remembered probabilities of once future things." *Organizational Behavior and Human Performance*, 13, 1-16.
- Glaser, M., Langer, and M. Weber, 2005, "Overconfidence and trading volume," Working Paper, Mannheim University.
- Haigh, M., and J. List, 2005, "Do professional traders exhibit myopic loss aversion? An experimental analysis," *The Journal of Finance*, 60, 523-534.
- Harrison, G., and J. List, 2004, "Field experiments," *Journal of Economic Literature*, 4, 1009-1055.
- Harrison, G., M. Lau, and M. Williams, 2002, "Estimating Individual Discount Rates in Denmark: A Field Experiment," *American Economic Review*, 92(5), 1606-1617.

- Hawkins, S., and R. Hastie, 1990, "Hindsight: Biased Judgments of Past Events After the Outcome is Known." *Psychological Bulletin*, 107, 311-327.
- Hirshleifer, D., "Investor Psychology and Asset Pricing," *Journal of Finance*, 56(4), August (2001):1533-1598.
- Holzl, E., E. Kirchler, and C. Rodler, "Hindsight Bias in Economic Expectations: I Knew All Along What I Want to Hear." *Journal of Applied Psychology* 87 (2002): 437-443.
- List, J., "Does market experience eliminate market anomalies?" *Quarterly Journal of Economics*, 118, 41-71.
- Mangelsdorff, L. and M. Weber, 1998, "Hindsight Bias im Prinzipal Agent Kontext," *Buchbeitrag in Organisation im Wandel der Märkte*.
- Musch, J., 2003, "Personality differences in hindsight bias," *Memory*: 473-489.
- Rabin, M., 2002, "Inference by Believers in the Law of Small Numbers," *Quarterly Journal of Economics* 117(3), August 2002, 775-816.
- Rabin, M., and J. Schrag, 1999, "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics* 114(1), 37-82.
- Rabin, M. (1998), Psychology and Economics, *Journal of Economic Literature*, Vol. XXXVI, 11-46.
- Roese, N. J., 2004, "Twisted pair: Counterfactual thinking and the hindsight bias." in D. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 258-273). Oxford: Blackwell.
- Sarin, R. K. and M. Weber, 1993, "The Effect of Ambiguity in Market Setting," *Management Science*, 39, 135-149.
- Wasserman, D., Lempert, R.O., & Hastie, R., 1991, "Hindsight and causality," *Personality and Social Psychology Bulletin*, 17, 30-35.

Proofs

Proof of Lemma 1:

$$\begin{aligned}
 \hat{E}_s(\tilde{v}|\tilde{s} = s) &= \hat{E}_s(\tilde{v}) + \frac{\text{cov}(\tilde{v}, \tilde{s})}{\text{var}(\tilde{s})}(s - \hat{E}_s(\tilde{v})) \\
 &= \omega\delta s + [\omega(1 - \delta) + (1 - \omega)]\mu + \delta(s - \omega\delta s - [\omega(1 - \delta) + (1 - \omega)]\mu) \\
 &= \delta(1 + \omega - \omega\delta)s + (1 - \delta)[\omega(1 - \delta) + (1 - \omega)]\mu.
 \end{aligned}$$

QED

Proof of Lemma 4:

Applying Lemma 1, the answer of participant b in group B is:

$$\hat{E}_{s_b}(\tilde{v}|\tilde{s}_b = s_b) = \delta(1 + (1 - \delta)\omega_b)s_b + (1 - \delta)[1 - \omega_b\delta]\mu.$$

Thus, the average answer of the participants in group B is:

$$W_B = \sum_{b \in B} [\delta(1 + (1 - \delta)\omega_b)s_b + (1 - \delta)(1 - \delta\omega_b)\mu] / N_B,$$

where N_B is the number of participants in the group. The average answer of group B is:

$$\begin{aligned}
 W_B &= (\delta + \delta(1 - \delta)\frac{\sum_b \omega_b}{N_B})v \\
 &\quad + \frac{\sum_b (\delta + \delta(1 - \delta)\omega_b)\epsilon_b}{N_B} + (1 - \delta)(1 - \delta\frac{\sum_b \omega_b}{N_B})\mu.
 \end{aligned}$$

Under our assumptions, as N_B goes to infinity, by the law of large numbers, this goes to:

$$(\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)E(\tilde{v}).$$

QED

Proof of Proposition 1:

If a is hindsight biased, he fails to remember properly $E(\tilde{v})$ and instead reconstructs it as: $\hat{E}_v(\tilde{v}) = \omega_a v + (1 - \omega_a)\mu$.

Thus, his report on the average answer of group B is:

$$\begin{aligned}
 \hat{E}_v(W_B|v) &= (\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)\hat{E}_v(\tilde{v}) \\
 &= (\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)(\omega_a v + (1 - \omega_a)\mu).
 \end{aligned}$$

That is:

$$\hat{E}_v(W_B|v) = [\delta + (1 - \delta)(\delta\omega + (1 - \omega\delta)\omega_a)]v + (1 - \delta)(1 - \omega\delta)(1 - \omega_a)\mu.$$

$$\hat{E}_v(W_B|v) = [\delta(1 + (1 - \delta)\omega) + (1 - \delta)(1 - \omega\delta)\omega_a]v + (1 - \delta)(1 - \omega\delta)(1 - \omega_a)\mu.$$

Thus:

$$\begin{aligned}\hat{E}_v(W_B|v) - W_B &= [\delta(1 + (1 - \delta)\omega) + (1 - \delta)(1 - \omega\delta)\omega_a]v \\ &\quad + (1 - \delta)(1 - \omega\delta)(1 - \omega_a)\mu \\ &\quad - [(\delta + \delta(1 - \delta)\omega)v + (1 - \delta)(1 - \omega\delta)\mu].\end{aligned}$$

That is:

$$\begin{aligned}\hat{E}_v(W_B|v) - W_B &= (1 - \delta)(1 - \omega\delta)\omega_a v - (1 - \delta)(1 - \omega\delta)\omega_a \mu \\ &= (1 - \delta)(1 - \omega\delta)\omega_a [v - \mu].\end{aligned}$$

Furthermore,

$$\begin{aligned}v - W_B &= v - (\delta + \delta(1 - \delta)\omega)v - (1 - \delta)(1 - \omega\delta)\mu \\ &= (1 - \delta)[1 - \delta\omega](v - \mu).\end{aligned}$$

Hence:

$$h^a = \frac{\hat{E}_v^a(W_B) - W_B}{v - W_B} = \frac{(1 - \delta)(1 - \omega\delta)\omega_a [v - \mu]}{(1 - \delta)(1 - \omega\delta)[v - \mu]} = \omega_a.$$

QED

Proof that the Camerer, Loewenstein and Weber index will be 0 on average when agents are rational and errors are correlated:

In this parametrization, since all agents are rational, we have:

$$W_B = \sum_{b \in B} \frac{E(v|s_b)}{N_B} = \delta \sum_{b \in B} \frac{s_b}{N_B} + (1 - \delta)\mu,$$

where:

$$\delta = \frac{\text{cov}(s_b, v)}{V(s_b)} = \frac{V(v)}{V(\theta) + V(\epsilon_b)}.$$

Now,

$$\sum_{b \in B} \frac{s_b}{N_B} = v + \theta + \sum_{b \in B} \frac{\epsilon_b}{N_B},$$

which goes to: $v + \theta$ as N_B goes to infinity. Hence, when N_B goes to infinity, W_b goes to:

$$\delta(v + \theta) + (1 - \delta)\mu,$$

Correspondingly, for the agent a in group A , $E(W_b|v, s_a)$ goes to:

$$\delta v + (1 - \delta)\mu + \delta E(\theta|s_a, v).$$

Thus, his Camerer, Loewenstein and Weber (1989) index becomes:

$$h^a = \frac{[\delta v + (1 - \delta)\mu + \delta E(\theta|s_a, v)] - [\delta(v + \theta) + (1 - \delta)\mu]}{v - [\delta(v + \theta) + (1 - \delta)\mu]} = \frac{\delta(E(\theta|s_a, v) - \theta)}{(1 - \delta)(v - \mu) - \delta\theta}.$$

Hence, the average of all the Camerer, Loewenstein and Weber (1989) indexes in group A is:

$$\sum_{a \in A} \frac{h^a}{N_A} = \frac{\delta(\sum_{a \in A} \frac{E(\theta|s_a, v)}{N_A} - \theta)}{(1 - \delta)(v - \mu) - \delta\theta}.$$

As N_a goes to infinity, $\sum_{a \in A} \frac{E(\theta|s_a, v)}{N_A}$ goes to θ . Consequently, the average Camerer, Loewenstein and Weber (1989) index goes to 0.

QED

Proof of Lemma 5:

The objective of the agent can be rewritten as:

$$\text{Max}_{q_a} \hat{E}_s[(q_a + z_a)' \tilde{v} - q_a' p^t(1 + r)|s] - \frac{\gamma}{2} \text{Var}[(q_a + z_a)' \tilde{v}|s].$$

The optimal demand from the point of view of agent a is:

$$q_a = \frac{\tau}{\gamma} \left\{ \hat{E}_s[\tilde{v}|s] - p(1 + r) \right\} - z_a.$$

Relying on Lemma 1 this yields:

$$q_a = \frac{\tau}{\gamma} \left\{ \hat{\delta}^a s + (1 - \hat{\delta}^a)\mu - p(1 + r) \right\} - z_a,$$

where:

$$\hat{\delta}^a = \delta(1 + (1 - \delta)\omega^a).$$

The market clearing condition is:

$$\int_a q_a da = x.$$

Substituting the demand of the agents, this condition amounts to:

$$\int_a \left\{ \frac{\tau}{\gamma} \left\{ \hat{\delta}^a s + (1 - \hat{\delta}^a)\mu - p(1 + r) \right\} - z_a \right\} da = x$$

That is:

$$\hat{\delta}^m s + (1 - \hat{\delta}^m)\mu - p(1 + r) = \gamma\tau^{-1}(x + z_m),$$

where:

$$\hat{\delta}^m = \int_a \hat{\delta}^a da.$$

This yields the market-clearing price, stated in the lemma.

QED

Table 1: Bankers population: Number, age and experience in Frankfurt and in London

Location	Earnings	Number	Average Age	Average Experience
FRANKFURT	High	12	34	11.83
	Mid	12	34.50	12.50
	Low	17	29.76	7.06
LONDON	High	14	41.86	18
	Mid	14	34.07	9.89
	Low	16	32.38	6.36

Table 2: Hindsight bias questions and results for Frankfurt

Question	True answer	Other group answer	This group prediction (given true answer)	Bias
FRANKFURT A				
Consumer price change Germany 10/2000 to 10/2001	2.5	1.68	2.9	1.49
Drop of Swiss stock market from all time high to 10/2001	30	38.55	35.24	0.39
Change in price of gold 10/2000 to 10/2001	6.2	13.69	6.53	0.96
Number of bankers at Lazard 8/2001	200	3477.27	691.74	0.85
% of Merrill Lynch's earnings from asset management in 2001	10	31.32	20.29	0.52
FRANKFURT B				
Euro/Yen exchange rate 2/13/2001	109	135.85	103.41	1.21
Ratio of foreign debt to GDP in Brazil at the end of 2000	38.4	176.69	59.5	0.85
Net revenues drop Texas Instruments end of third quarter 2001	1285	1034.06	925.05	-0.43
Growth rate of real GDP Russia 2000	8.3	5.57	5.02	-0.20
Growth rate of real GDP OECD countries, 2000	4.1	3.5	3.66	0.27

Table 3: Hindsight bias questions and results for London

Question	True answer	Other group answer	This group prediction (given true answer)	Bias
LONDON A				
Retail Price Index change 10/2000 to 10/2001	1.57	5.37	5.61	-0.06
Drop of Swiss stock market from all time high to 8/2002	39.17	32.06	37.4	0.75
Change in price of gold from 10/2000 to 10/2001	6.2	21.88	10.48	0.73
Bankers at Lazard in August 2001	200	1876.36	641.88	0.74
% Merrill Lynch's earnings from asset management, 2001	10	22.08	18.88	0.27
LONDON B				
Euro/Yen exchange rate 2/13/2002	116.1	135.5	122.22	0.68
Ratio of foreign debt to GDP Brazil, end of 2000	38.4	126.24	57.83	0.78
Net revenue increase TI 2 nd quarter 2002	1.06	-4.35	-3.47	0.16
Growth rate of real GDP Russia, 2001	5	4.25	4.92	0.89
Growth rate of real in GDP OECD countries 2001	1.2	3.75	2.12	0.64

Table 4: Average Hindsight Bias in London and Frankfurt

	Bias in London	Bias in Frankfurt
Median	0.635	0.5
Mean	0.547	0.579
Std. Deviation	0.88	0.786
Number of Observations	49	41

Table 5: Hindsight bias for the three earnings categories

Earnings	Pool			Frankfurt			London		
	Low	Mid	High	Low	Mid	High	Low	Mid	High
Median	0.636	0.684	0.375	0.512	0.63	0.303	0.636	0.809	0.429
Mean	0.546	0.924	0.192	0.678	0.883	0.133	0.546	0.959	0.242
Std. Dev.	1.009	0.833	0.638	0.737	0.948	0.473	1.009	0.757	0.766
Nb obs	33	26	26	17	12	12	16	16	14

Table 6: Significance of the difference between biases across earnings categories

	z-stat	p-value in percent
High versus middle	3.33	.09
High versus low	2.27	2.29
Middle versus low	1.02	30.6

Table 7: Estimated conditional probabilities

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.46	.32	.14
Prob(Mid Earnings bias)	.14	.39	.38
Prob(Low Earnings bias)	.39	.29	.48

Table 8: Hindsight bias controlling for experience.

Experience		Earnings		
		Low	Mid	High
Less than 10 years	Number of observations	27	13	4
	Average experience (in years)	5.1	6.17	5.78
	Median bias	.61	.85	.055
	Mean bias	.55	1.12	-.055
10 years or more	Number of observations	6	13	22
	Average experience (in years)	13.83	15.61	16.87
	Median bias	.77	.64	.43
	Mean bias	.69	.72	.23

Table 9: Probability of performance, conditional on hindsight bias, for the three occupational categories

Panel A: Research

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.57	.14	.00
Prob(Mid Earnings bias)	.0	.14	.50
Prob(Low Earnings bias)	.43	.71	.50

Panel B: Sales

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.57	.29	.21
Prob(Mid Earnings bias)	.14	.43	.43
Prob(Low Earnings bias)	.29	.29	.36

Panel C: Trading and Sales-trading

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.20	.40	.00
Prob(Mid Earnings bias)	.20	.40	.25
Prob(Low Earnings bias)	.60	.20	.75

Table 10: Probability of performance, conditional on hindsight bias, for the three mistake ratio categories

Panel A: Lowest mistake ratio

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.55	.20	.33
Prob(Mid Earnings bias)	.11	.40	.50
Prob(Low Earnings bias)	.33	.40	.17

Panel B: Medium mistake ratio

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.45	.20	.33
Prob(Mid Earnings bias)	.22	.60	.20
Prob(Low Earnings bias)	.33	.20	.47

Panel C: Largest mistake ratio

	Low bias	Mid bias	High bias
Prob(High Earnings bias)	.34	.20	.20
Prob(Mid Earnings bias)	.22	.20	.45
Prob(Low Earnings bias)	.44	.60	.35