The Journal of Psychology and Financial Markets
2001, Vol. 2, No. 4, 190–200

Momentum, Rational Agents and Efficient Markets

John Crombez

Descriptive behavioral models explain the momentum anomaly by assuming that financial agents are irrational. However, investors are not tested to be susceptible to the cognitive failures observed in psychological experiments. We consider an environment where financial agents are rational, markets are efficient as defined by the Grossman–Stiglitz [1980] efficiency, and there are market imperfections in the information market. Based on a simulation experiment, we find that returns on momentum strategies can exist in this environment because of the noise in expert information. We empirically find that even in a sample of large and liquid stocks, this noise is still observable and, hence, momentum can be empirically found for these samples even when agents are rational and markets are efficient.

The literature holds few examples of successfully explaining anomalies by risk factors or changes in fundamentals. Because of this observation, many researchers question the efficiency of financial markets and assume that investors are susceptible to cognitive failures.[1] These assumptions are the basis for descriptive behavioral models in finance (how agents behave).

This paper benchmarks these models by starting from the assumption that markets are efficient and investors are rational. We show that in such a normative framework (how agents should behave), anomalies can still be observed.

The existence of momentum in stock prices is one market anomaly that is difficult to explain. Past literature has explained momentum by descriptive behavioral models, where the assumption is that financial agents are irrational (Daniel, Hirschleifer, and Subrahmanyam, 1998 and Barberis, Shleifer, and Vishny, 1998). Stocks with strong past performance continue to outperform stocks with poor past performance in the next period with an average excess return of about 1% per month (Jegadeesh and Titman, 1993, 1999). Barberis, Shleifer, and Vishny [1998] explain the existence of momentum as an underreaction to news signals. The news is not quickly reflected in the price and so continues to have an impact in subsequent periods. In this case, however, because of their underreaction, investors are not considered rational.

There is little evidence in the past literature that the anomalies can be explained by fundamentals or risk. For example, the Fama–French [1993] three-factor model does not explain these momentum returns. Furthermore, Dreman and Lufkin [2000] find that fundamentals are relatively stable when large price changes are observed for favored and unfavored stocks. They conclude that only psychological factors can explain these large price changes.

This paper presents numerical evidence that anomalies such as momentum effects can exist even when markets are efficient and investors are ex ante rational in their decision-making. We show that distortions in the information market can explain the existence of anomalies. When investors use historical price changes and expert information as their information set, the level of disagreement among experts determines the reaction to news.

We also show that expert information is presented to the investor differently depending on the size of the firm and the exposure of the stock to market risk. This is an important theoretical consideration because it benchmarks the assumptions made in descriptive behavioral models from a normative point of view. In other words, if anomalies can be observed when markets are efficient and investors are rational, the assumptions of cognitive failures should be proven true for investors before they are applied. If this is not the case, the descriptive behavioral model cannot be regarded as a possible alternative model to explain stock returns.

The Financial Environment

As DeBondt and Thaler [1995] say, the principle of assumed optimality states that actual decision-making is not often studied, implying that decision processes do not affect outcomes. This is certainly true for finance. Descriptive behavioral models start from the assumption that generally observed cognitive failures

John Crombez is Professor at Hogeschool Gent teaching Investment Analysis. He is also the Director of Ghent Finance Center (KPMG and Ghent University). He received his Masters Degree in Statistical Sciences (Université de Neuchâtel, CH) and his PhD in Economics.

Requests for reprints should be sent to: John Crombez, Ghent University, Department of Financial Economics, Sint-Pietersplein, 4, 9000, Ghent, Belgium. Email: john.crombez@rug.ac.be

(especially in psychological experiments) are true for financial agents as well.

The problems with this assumption are twofold. First, we do not know how susceptible financial agents are to cognitive failures observed in other disciplines such as psychology. Second, as Fama [1998] notes, descriptive behavioral models work well on the anomalies they are designed to explain, and these models rarely come up with an alternative to the efficient market hypothesis (EMH). We summarize these two statements with the following argument: To consider human judgment suboptimal without discussion of the limitations of optimal models is naïve (Einhorn and Hogarth, 1988).

We empirically study rational judgments about expected returns by financial agents in an environment that maintains its realism. Rational expectations are formed when the financial agent makes Bayesian forecasts. This means that the agent observes a set of evidence in the market and weighs this evidence according to Bayes' rule. The decision problem in this paper is the judgment about the next payoff or the return. If the news is good, prices will be bid up, the payoff will be positive, and the expected return will decline.

The setting in this environment is as follows. A financial agent makes a judgment about the next period's return based on his information set. This information set consists of historical price changes and information about the value of the stock communicated by experts. We believe this information set is relevant because the investor aims to simplify a very complex environment in order to make a judgment. There are experts in the markets, such as analysts, who do not trade, but who communicate their opinions by such methods as earnings forecasts and buy and sell signals. In forming their opinions, the experts use costly information (they visit companies, contact management, use complex models) that is less accessible for the other agents (because the cost is high).

According to the agent defined in the Hong–Stein model [1999], an uninformed investor uses historical price changes and makes linear forecasts. If the historical price changes are the likelihood (or the probability distribution of the data) for the investor and returns are normally distributed, this agent's forecast about the next period's payoff will be the sample mean.

Next, this financial agent is aware that the decision problem is complex and seeks aid from experts. He uses the prediction of the next period's payoff by the experts as prior evidence for his base rates. Using all this evidence, the investor applies Bayes' rule and makes a rational forecast about the next period's payoff. As described, the evidence the agent uses to make the judgment about the future price change is the historical price change (B) and the expert opinion about the future price change (A).

$$P(A/B) = \frac{P(B/A)P(A)}{P(B)}$$

To summarize, the investor in our model makes a judgment about the next period's return based on historical price changes and expert opinions. He evaluates the information by applying Bayes' rule. In doing so, this investor forms rational expectations.

This paper is organized as follows. The second section presents the assumptions we make in order to determine what environment the decision is made in. The third section shows how we numerically solve the decision problem given the environment. The fourth section provides the results of the computation, and the final section offers our conclusions.

The Decision-Making Environment

We first define the environment in which the financial agent makes his rational judgments.

Assumption 1: There is asymmetric information in financial markets.

This assumption implies that some information must be acquired against a cost. Part of the public information is costly, but it does not only concern fundamental firm information. According to decision-making theory, when the task is difficult, experts also acquire additional information because they use models and theory to improve their judgment (Griffin and Tversky, 1992). With respect to financial markets, this seems an important extrapolation of costly public information. Investment banks and brokers, for example, invest a great deal in the search for better models and theories.

Assumption 2: Financial markets are efficient according to the Grossman–Stiglitz [1980] formulation of efficiency.

The second assumption implies that market imperfections exist because of the noise in the precision of the information[2] from informed market participants. Note that Griffin and Tversky [1992] report that the level of dispersion in this information (or the inverse of the strength of the evidence) is important in weighing the evidence by a Bayesian forecaster or a rational investor. The strength of the evidence is more important than the number of market participants (the weight of the evidence) who give information when the decision problem is difficult.

Consider a case where two experts communicate their opinion about the future return of a stock. If they communicate the same prediction, this will be more important to the decision-maker than the fact that there are only two experts.

The market for information is competitive, but following from the second assumption, this market is imperfect. Together with Assumption 1, this implies that prices fully reflect all costless public information. Other public information is costly.

Assumption 3: There exists a theoretical relationship between news on expected returns and prices.

If we combine Assumption 3 with Assumption 2, we see the importance of the precision of this type of news. In this framework, news comes from experts who have costly public information and these experts, such as analysts, do not necessarily trade. This is relevant for financial markets.

Rational Expectations About the Next Period's Return

Aspects of the Framework

If Bayes' rule is the appropriate language[3] for investors to use to form their beliefs, then it is worthwhile to discuss it here further. By 1952, Markowitz was already developing the idea that, to form a belief about the parameters needed to apply his expected return variance rule, it would be necessary to combine signals from statistical parameters with the judgment of experts. In doing so, we could use Bayes' rule to formalize his probabilities.

Our framework here suggests how agents should form their beliefs in order to obtain a rational prediction for the price change of the assets of interest for the next period. In order to test this setting numerically, we simulate the rational weighing of the evidence. This allows us to evaluate the actual weighing function applied by a rational investor and its implications for the behavior of the financial agent.

The Choice of the Prior

In finance, depending on the decision problem being addressed, there are several possible priors that are relevant (such as earnings yield, expert surveys, and so on), but the choice of a prior is also one of the limitations of each tested model (Lenk and Wedel, 2001). A discussion of each prior is vital for the evaluation of each model.

Here, we use one-year analyst consensus forecasts of earnings yield (FY1) as the prior or expert information. Because of the theoretical link between accounting multiples and expected returns, this seems a reasonable choice. But Dechow, Hutton, and Sloan [1999] also find that FY1 is an informative variable for determining the true value of the stock.

We illustrate the use of this prior information with straightforward examples. Suppose the Gordon growth model [$P = D/(r - g)$] holds. A price-to-dividend ratio of 25 means that the expected return r minus the growth rate g equals 4%. A 1-percentage point change in expected return translates into a 25-percentage point change in price (this example is taken from Cochrane, 2001).

The same example can be applied to Ohlson's [1995] residual income model (RIM). We suppose the following price dynamics: Abnormal returns are persistent and other information is not. The RIM takes the following form: $P = e_{t+1}/r$, where e_{t+1} is the one-year analysts' earnings forecast (see Dechow, Hutton, and Sloan, 1999). Suppose that the earnings yield forecast is 5%. A 1-percentage point change in the expected return causes a 17-percentage point change in the price.

The range of analyst forecasts determines the prior scale (as in Kinney, Burgstahler, and Martin, 1999). It is considered an appropriate proxy for the strength of the expert evidence. Comparable to the range measure, we use the highest (H) and lowest (L) earnings yield forecasts on a stock to compute the proxy for the strength of evidence (τ). The variance of the expert opinion is calculated from the Parkinson [1980] measure for extremes. The variance of the expert opinions is given by $0.361(H - L)^2$.

Using a range of expert opinion and not considering the number of analysts making a forecast is defended by findings in the psychology literature and the finance literature. Griffin and Tversky [1992] suggest that the strength of the evidence (the level of divergence in the expert opinions) is more important than the weight of the evidence (the number of analysts making a forecast). Hong, Lim, and Stein [2000] find that the diffusion of information is lower for small-caps. Information diffusion is proxied by the number of analysts covering a firm. This number is best explained by the size of the firm. Other variables such as book-to-market or residual analyst coverage (the impact of one additional analyst) explain little of the total analyst coverage.

Model Specifications

To simulate the judgments given the financial agent's available information in the same way as a Bayesian forecaster, we apply the Bayesian bootstrap regression procedure (BBR). The main reason for using this method is the flexibility of the prior choice. In doing so, we reduce the loss of information by choosing an algorithm that allows for flexible individual priors (as opposed to conjugate priors and common priors).

The basics of Bayesian sampling analysis can be found in Kloek and Van Dijk [1978]. But the foundation of the BBR approach itself is attributable to Geweke's [1986] solution to the constrained normal linear regression model. The procedure itself generates outcomes from the likelihood that can be used for integration procedures of expectations with respect to the posterior distribution (Heckelei and Mittelhammer, 1996).

The major difference between our procedure and Geweke's is that we do not make the same parametric assumptions that he does. The BBR procedure is used as an engine to develop Bayesian forecasts about the decision problem of choosing the ex ante optimal expected return or payoff for an asset or asset class. For a rational agent, this engine represents the human mind weighing the evidence as a pan balance.

Suppose that asset prices follow a random walk. In this case, any linear forecasting rule for the future price change based on historical price changes alone is not efficient (Campbell, Lo, and MacKinlay, 1997). Indeed, empirical analysis shows that returns are hard to predict at short horizons. Since we start from a decision problem about expected returns that is solved at a relatively high frequency (e.g., monthly), it is reasonable to assume that historical prices will convey little or no information about future price changes. Moreover, this high frequency of decisions is relevant to the real world, where, for example, a monthly rebalancing of portfolios is not unusual.

Suppose that P_t denotes the observed time series logarithmic prices of an asset at time t. If prices follow a random walk with drift, the price process can be described by Equation (1).

$$P_t = \mu + P_{t-1} + \varepsilon_t \tag{1}$$

Changing the log price in this equation to the left side shows that the time series of returns of the asset (r_t) is described by a constant μ and a white noise error term. In fact, this expression is the same as a normal linear regression model of the returns on a unit vector (ι_t) [see Equation (2)]. The parameter μ denotes the expected return for the uninformed agent in the Hong–Stein model.

$$r_t = \mu\iota_1 + \varepsilon_t \tag{2}$$

This regression model is the point we start from to apply the BBR methodology as described by Heckelei and Mittelhammer [1996]. Classical inference methods using a likelihood function are often difficult to evaluate. If the problem is solved in a Bayesian framework, analytical solutions are scarce (Geweke, 1989). In most cases, priors and functions of interest are not tractable, and, in finance, this problem is important. Often, uncertainty and estimation risk are reduced by applying Bayesian methods with a common shrinkage parameter (Jorion, 1991).

For example, in portfolio selection, the return on the minimum-variance portfolio is used as the prior parameter for all assets in the Bayes–Stein estimator. As Bawa, Brown, and Klein [1979] state, the use of a common shrinkage parameter implies a loss of information about individual assets or asset classes.

Bayesian Bootstrap Regression

Starting from the normal linear regression model in (2), we denote the return series or the series of historical price changes as the likelihood $f(r|\mu, \sigma) = L(\mu, \sigma|r)$. From here, we drop the subscript t from r_t for reasons of notational simplicity, with r still denoting the vector of historical price changes. In the Bayesian bootstrap approach, the residual series in (2) follows, in non-parametrical form, an unknown empirical distribution $h(\varepsilon|\sigma)$ determining the parameter σ in the likelihood. First, we develop the general solution to the problem. Next, we extend the application to its non-parametrical form. This extension is used for the estimation of the expected return in this paper.

Using Bayes' rule as the rational calculus to make a judgment for the financial agent, we define the posterior density for the parameters as in Equation (3):

$$\begin{aligned} p(\mu,\sigma|r) &= \frac{p(\mu,\sigma)f(r|\mu,\sigma)}{p(r)} \\ &\propto p(\mu,\sigma)f(r|\mu,\sigma) \end{aligned} \tag{3}$$

Here, $p(\mu, \sigma|r)$ is the posterior density, $p(\mu, \sigma)$ denotes the prior information on the parameters, and $p(r)$ denotes the data distribution. Since we want to model decisions on the expected return (μ), we concentrate on the marginal posterior distribution for μ [see Equation (4)]. Throughout the remainder of the text, independence between the prior parameters μ and σ is assumed.

$$\begin{aligned} p(\mu|r) &= \int_0^\infty p(\mu,\sigma|r)d\sigma \\ &\propto p(\mu)\int_0^\infty p(\sigma)f(r|\mu,\sigma)\,d\sigma \end{aligned} \tag{4}$$

In order to solve (4) numerically, the likelihood function is weighted by the prior evidence on σ and then the parameter σ is integrated out. Next, the expression is normalized to unit total mass, dividing by the integral with respect to μ. This reformulation is made in order to make it a proper density (Heckelei and Mittelhammer [1996]). We obtain $f_\sigma(r|\mu)$, which can be regarded as the posterior distribution of μ when $p(\mu)$ equals 1. In this case, it is possible to formalize posterior expectations of any function of μ as an integral involving this function of interest and the marginal posterior (Geweke, 1989).

We denote $E[g(\mu)]$ as the posterior expectation of the function of interest on the parameter μ. Next, we substitute for the distribution of $f_\sigma(r|\mu)$, by an importance function $I(\mu)$. We also control the computations for convergence characteristics. If there is slow convergence—observable when $E[g(\mu)]$ is not stable as the number of replications increases—then the choice of $I(\mu)$ is poor. For these simulations, $E[g(\mu)]$ is stable.

If, in this case, μ_i^* with i = 1, … , n are i.i.d. outcomes from $f_\sigma(r|\mu)$, $E[g(\mu)]$ can be approximated by a prior weighted average. The weights $p(\mu)$ are determined by the relevance of the simulated μ_i^* relative to the prior evidence (the expert consensus forecast and the scale of the prior evidence determined by the level of disagreement among analysts).

$$E[g(\mu)] = \frac{1/n\sum g(\mu_i^*)p(\mu_i^*)}{1/n\sum p(\mu_i^*)} \tag{5}$$

Finally, from Equation (5), we can estimate the necessary expectations of the function of interest. In (6), we show the expectation of 1) the posterior mean, and 2) the numerical standard error (nse) using the Bayesian bootstrap regression (with N denoting the number of replications) (see Heckelei and Mittelhammer, 1996).

$$(a) \qquad \hat{E}(\mu) = \frac{\sum_{i=1}^{N}\mu_i^* p(\mu_i^*)}{\sum_{i=1}^{N} p(\mu_i^*)}$$

$$(b) \qquad nse = \frac{\sum_{i=1}^{N}(\mu_i^* - \hat{\mathrm{E}}(\mu))^2 p(\mu_i^*)^2}{\left[\sum_{i=1}^{N} p(\mu_i^*)\right]^2} \tag{6}$$

Empirical Evaluation of the Environment and Rational Expectations

The Case of Expert Evidence

Before we assess the hypothesis in this section, we examine a practical case of expert evidence based on the environment as presented in the second section. We focus on the strength of the expert evidence for a sample of stocks that can be considered for investment purposes.

Carhart [1997] finds no evidence for beneficial momentum investment strategies after accounting for transaction costs. This is rather confusing, because Rouwenhorst [1998] finds a return on a European momentum strategy of 1% per month. These findings are based on return data for the period 1980–1995, and stocks in the sample cover about 60% to 90% of the included European countries. Apparently, no really small stocks are included and the individual stocks should not be susceptible to infrequent trading. Nevertheless, momentum is observed.

This implies that it is important to consider whether there are differences in the strength of the expert evidence. We investigate this question for a European database of liquid stocks. We collect relevant information for all European stocks in the MSCI Europe index between March 1992 and August 2000 from the IBES database, as well as the mean consensus forecast, the highest and lowest forecast, and the number of analysts covering the stock. As outlined before, we use the earnings yield forecast as the expert opinion for the next period's price change. The database we use is an intersection database between the IBES data and the Datastream dataset of market values and returns for the stocks in the MSCI Europe index.

The average number of stocks in the index during this period is 588. The average number of MSCI Europe index stocks covered by analysts is 471. Between 1992 and 2000, the coverage of index stocks in Europe increased steadily. In 1992, an average of sixteen analysts announced a one-year-ahead forecast on a stock. This increased to a maximum of twenty-one analysts in 1997, and dropped to nineteen in 2000. Given the increasing coverage of stocks, we can assume that at the end of the sample more stocks were covered, on average, by fewer analysts.

A last important feature of the dataset is that Portuguese stocks entered the index only at the end of 1997. There are almost no stocks covered by only one analyst (in that case, τ is equal to two times the standard deviation of the market portfolio).

For all stocks, we calculate the prior scale as a proxy for the strength of the expert evidence. Kinney, Burgstahler, and Martin [1999] suggest using the range of expert opinions or the standard deviation of opinions as a measure of precision.

The Parkinson measure is computed for all stocks covered by analyst forecasts. All stocks are ranked according to their market capitalization (MV) and their beta (computed using the CAPM and a market value-weighted portfolio), and ten deciles are formed for each ranking parameter (an average of fifty stocks). For each decile and for each month, the equally weighted strength of evidence, $\bar{\tau}$, is computed. For the cross-section in each month for each decile, we calculate the mean by applying the robust mean absolute deviation (MAD) algorithm. The data reveal that there are extreme outliers. Applying MAD prevents bias from the outliers.

We find that even in a sample of liquid stocks, there are large differences in the strength of expert evidence.

Table 1. *Stock Characteristics and Signal Strength*

	Ranking Variable: MV Reported Variable: $\bar{\tau}$	Ranking Variable: beta Reported Variable: $\bar{\tau}$
Q1	0.02914	0.00169
Q2	0.02801	0.00158
Q3	0.00331	0.00077
Q4	0.00219	0.00047
Q5	0.00085	0.00077
Q6	0.00040	0.00077
Q7	0.00042	0.00023
Q8	0.00036	0.00093
Q9	0.00028	0.00108
Q10	0.00008	0.00526

Note: Average characteristics for firms in decile portfolios ranked by market capitalization (MV) or beta. For the MV and beta deciles, the Parkinson [1980] measure (τ) is estimated. The Q1 portfolios have low values for the ranking variable.

In the Q1 and Q2 portfolio based on size (the smallest stocks), the precision of the expert information is extremely low (Table 1). The variance of the information is around 2.85% (a standard deviation of 17%). This implies that stocks in these deciles are surrounded by noisy expert signals (even after accounting for outliers). The noise of the expert opinion decreases with the size of the firm.

For the decile of the most liquid stocks, the variance of the expert evidence is 0.0008 (a standard deviation of 0.9%). These figures confirm the finding of Hong, Lim, and Stein [2000] that news diffuses at a higher speed for larger firms, and implies that news from experts is quickly reflected in the stock prices of large firms. Expert evidence is also weak on average for low beta stocks and high beta stocks. The variance of the Q1 beta decile is 0.00169 (a standard deviation of 4%), and 0.00526 for the high beta decile (a standard deviation of 7%).

Scenarios and Weighing Functions

We evaluate three scenarios. The consensus expert opinion in each is that the return for the next period will be 4%. In other words, the FY1 is 0.04, so an imaginary stock price of 100 will rise to 104. This expert opinion is communicated with a different strength. For the three scenarios, we estimate the Bayesian forecasts using the BBR. The Appendix shows the general algorithm presented by Heckelei and Mittelhammer [1996] to explain how the Bayesian bootstrap regression is conducted. We use 5,000 bootstraps to mimic the weighing of evidence by the human mind using the BBR.

The likelihood used to make a decision consists of sixty months of historical price changes on the Belgian stock market. Returns are taken from Datastream using the total market index for Belgium and are computed as the monthly percentage of changes in the return index. These returns are the monthly returns for the period 1996–2000.

The difference in the strength of the evidence is presented by the figures in Table 1 for size deciles. The highest strength of evidence is shown by the Q10 number (0.00008, or a standard deviation of 0.8944%). The medium strength is provided by the Q5 number (0.00085, or a standard deviation of 2.9155%), and the weakest evidence is provided by the Q3 number (0.00331, or a standard deviation of 5.7533%). We do not use the Q1–Q2 numbers because of their extreme nature.

In the analysis in Table 1, negative earnings forecasts are included and explain the extreme observations in the Q1–Q2 deciles. Hence we use the Q3 number. We have also computed different scenarios (pessimistic expert opinion, different parameters), but the conclusions made below are unchanged.[4]

Table 2 presents the three scenarios and the most important characteristics of the judgment. The first column of Table 2 displays the strength of the evidence (we denote this as the scale τ). The higher the τ, the lower the strength of evidence. The assumption is that the financial agent obtains a consensus opinion about the expected price change for the next period with a level of disagreement displayed by the experts. The expert opinion is that given all costless and costly public information, the stock price will rise 4%. We assume that this consensus opinion is true in the sense that, given all available information, the rational expert valuation of the imaginary stock price is 104. Disagreement among experts can exist because the experts use different models to make a forecast and have different relationships with the firm. We suggest that this is a realistic representation of the financial environment.

In the second column, the square root of this scale or the standard deviation of the expert opinions is given. Hence, when experts disagree, the standard deviation of their opinions is as least as large as the standard deviation of the sample returns (in this example, 5.75% compared to 4.17% for the historical sample). In the third section, we denote this strength of evidence as the precision of the information from informed market participants. The noise in this signal causes the small market imperfections in the Grossman–Stiglitz model [1980].

Columns (3) and (4) give the highest and lowest forecasts as communicated by the experts. These are calculated from the definition of the Parkinson measure and assumed to be symmetric at around 4%. The last column shows the Bayesian forecast made by a rational financial agent who observes the historical price changes, the expert opinion, and the level of disagreement among analysts.

If there is no disagreement among experts, the financial agent's Bayesian forecast will be 4% (all the weight in the prior is on the 4%). When the level of disagreement among experts is low ($\tau = 0.00008$), the fi-

Table 2. *Scenarios*

Strength of the Evidence τ	Standard Deviation $\sqrt{\tau}$	Highest Expert Forecast	Lowest Expert Forecast	Bayesian Forecast
0.00000	0.00000%	4.000%	4.000%	4.000%
0.00008	0.89944%	4.744%	3.256%	2.262%
0.00085	2.91550%	6.426%	1.574%	1.728%
0.00331	5.75330%	8.788%	–0.008%	1.680%

Note: The expert opinion about the next period's return is 4% in all cases. The historical sample estimate for the next period's return is 1.668% (the standard deviation of 4.17%). For each opinion, a signal strength (scale) is given in the first column. The second column shows the square root of the strength of the evidence or the standard deviation of expert opinions. The third and fourth columns show the highest and lowest forecast based on the strength of the evidence when the two opinions are symmetrical around 4%. The last column shows the Bayesian forecast for the next period's return for the three scenarios. Numbers in the first column are expressed in percentage points.

nancial agent expects the return on the stock to be 2.262%. Looking at Table 1, this is on average the case for large stocks.

Figure 1 shows the weighing function for the financial agent we model and his environment. In the left panel, the weighing function is shown for the case when the precision of the information is high.

The right panel shows the likelihood, as well as the range of expert opinions and the distribution of possible outcomes. These possible outcomes are thought of as the possible returns the financial agent has in mind based on the historical price changes. In the simulation, these possible outcomes are the bootstraps for the mean. The weights of the weighing function are ranked according to the 5,000 possible outcomes for the next period's payoff and are accordingly ranked on the x-axis. The left part of the weighing function denotes the weights assigned to the lowest possible outcomes by a rational investor.

From Figure 1, some interesting features emerge. When the information signal is strong (or, in other words, when the strength of the evidence is high), the weighing function of the agent is to a large extent influenced by the optimistic expert opinion. Most of the decision weight goes to the larger possible outcomes for the next period's return. In this case, the Bayesian forecaster sets his expectation about the next period's return at 2.262%, which is above the expectation of the financial agent who does not use expert information (1.668%). Also, in the right tail of the possible judgment density, the weights sharply increase. The deviation of the updated judgment from the mean is in this case economically important (0.594% per month).

If the expert signal is weaker, however, the weighing function is far more conservative, even with an optimistic expert opinion of 4% (Figure 2). In this case, the standard deviation of expert opinions is still lower than the standard deviation of historical returns (2.92% compared to 4.17%). We see that weights only slightly increase for the larger possible judgments (Figure 2). On the lower tail, the agent displays a weak belief in a bad outcome of the return.

FIGURE 1
Low Level of Disagreement Among Experts

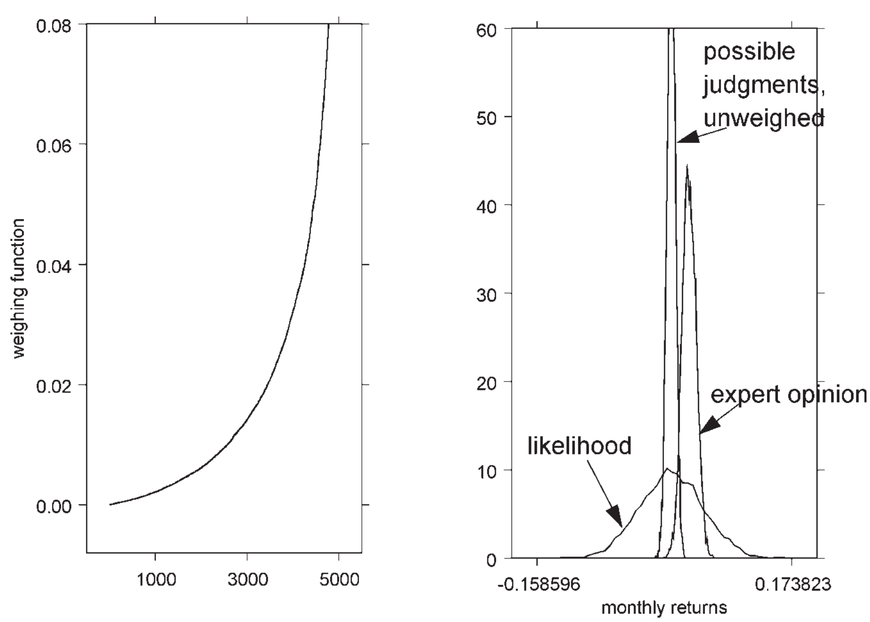

FIGURE 2
Moderate Level of Disagreement Among Experts

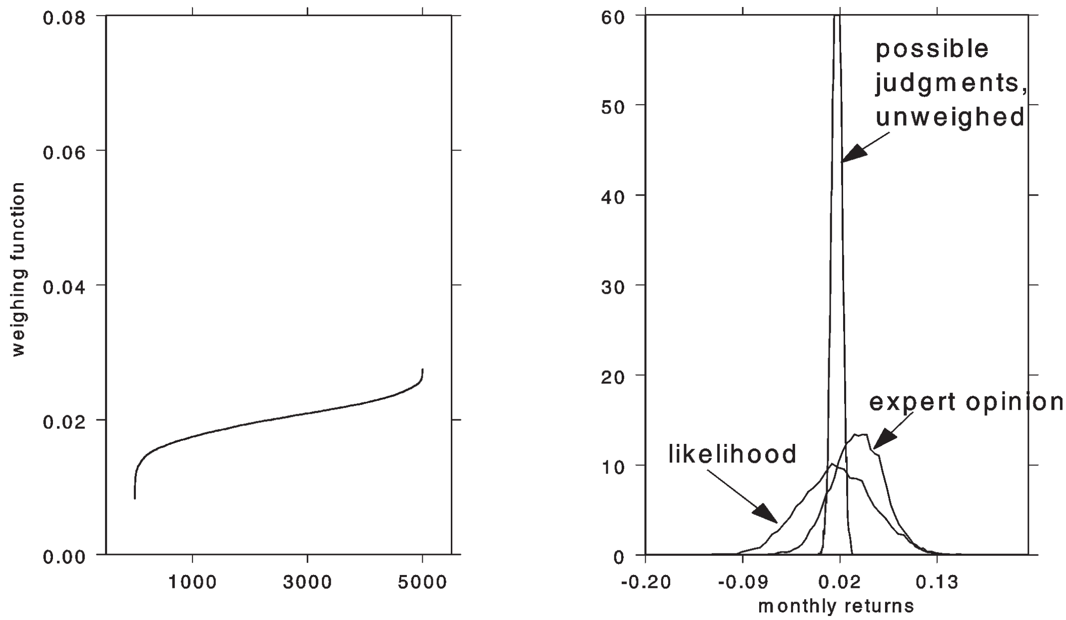

The large difference in the steepness of the weighing function is surprising. This means that the rational agent already puts more weight on the statistical evidence than on the expert opinion. The expected return in this case is above the price trader's expectation (1.728%, or 0.060% per month).

When there is a lot of disagreement among experts, almost all the weight that a rational agent assigns to the evidence goes to the historical observations of returns (Figure 3). The weighing function is rather flat, indicating that even with an optimistic expert opinion of 4%, the rational financial agent will have no strong preference for any of the possible outcomes he has in mind. The only effect the optimistic expert opinion will have on his judgment will be a slightly lower probability for a bad outcome than a good outcome. The Bayesian forecast of the return is 1.680%, or 0.012% per month above the historical mean.

FIGURE 3
High Level of Disagreement Among Experts

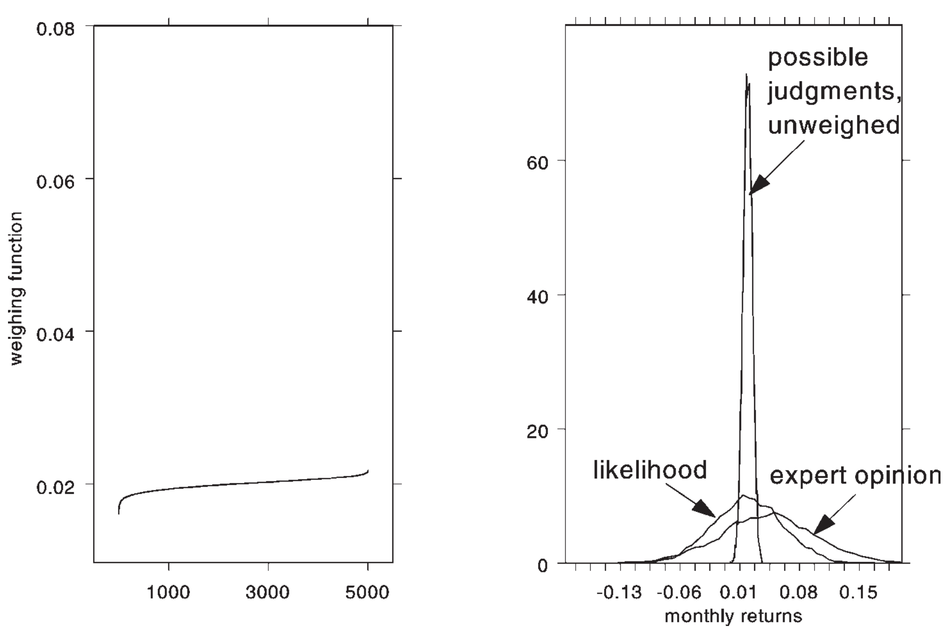

Is There a Case for Momentum Effects?

What does this imply for the existence of momentum? When the strength of the evidence is high, the rational financial agent expects that the price of the asset will go up and hence bids up its price. In the second and third scenario, the rational agent does not give a lot of weight to the expert opinion because the strength of the evidence is low. In this case, the expected return will not change much (the financial agent expects that the price will not increase much, in contrast with the optimistic expert opinion). When the expert evidence is strong, the precision of the information will be high and the prices of these stocks will be bid up. Hence, new information will quickly be reflected in the price and autocorrelation in returns will not be observed.

The next question is to what extent the diffusion of information induces momentum effects. To answer this question, we conducted the following simulation. In all three scenarios, the rational financial agent makes a judgment about the expected return. With respect to his consumption behavior, he will accordingly buy additional units of the asset based on this judgment. We assume there is no additional information about the asset in the next periods and, hence, the expert continues to value the price of the asset at 104. We simulate the return after two periods when prices are bid up relative to the expected return. If the financial agent bids up the price of the stock (when the strength of the expert evidence is high) to 102.262 (expected return of 2.262%), the experts still expect the stock price to rise 1.699%.

In Table 3, we show the simulated forecasts for the three scenarios assuming no additional information in these consecutive periods and that the expert's valuation remains the same. Additionally, we assume that the financial agent observes that the expert consensus remains the same and hence puts more weight on this opinion in the next period. We formalize this increase in weight by doubling the strength of the expert evidence.

Table 3 shows the case of no additional public information, and the case of an increase in the weight a financial agent gives to the expert opinion when this opinion remains the same as in the previous period. In this case, the stock price only rises to its true value of 104 in the first scenario after two periods. In the other two scenarios (where the strength of the evidence is lower), the true stock price (104) will not be attained after two periods. This means that differences in the information diffusion can account for the existence of momentum.

Note that our conclusion concurs with Barberis, Shleifer, and Vishny [1998]. It is possible that news is not so quickly reflected in the prices. However, these authors began with the assumption that investors are irrational.

The cases studied here are intermediate cases. Remember that the stock price will rise to its true value in the first period when there is no disagreement among experts. On the other hand, we observed cases of higher levels of disagreement than the ones tested here (the Q1 and Q2 deciles in Table 1). Hence, in the imaginary case that the expert forecast is the true forecast and in the case where financial agents are rational, we still find that momentum effects can exist because of frictions in the market of information. Here we define these frictions as the difference in the expert opinion because of the use of different complex models and different relationships between the experts and the firms (as documented by Lim, 2001).

Booms, Crashes and Overreaction

Can this story be relevant for what we observe in financial markets? We present two more examples. First, we observe that upward stock markets are slow, and that prices continue to rise for quite a long period. On the other hand, crashes are fast. We observe severe price declines in only a couple of days. Looking at the strength of the evidence and remembering the simulations for the three scenarios, we find that there is a possibility that these observations can be explained when markets are efficient. The strength of the evidence or the level of agreement among analysts is larger for each size decile when markets decline. This implies that the expert opinion will be more quickly reflected in the prices in downward markets and, again, that crashes are faster than booms.

Second, if we continue the simulation in Table 3 for more periods, we see that there is a case for overreaction when markets are efficient. Suppose there is still no additional news and the experts still believe that the true value of the stock is 104. When the experts disagree a lot about this value (scenario 3), the flat weigh-

Table 3. *Bayesian Forecasts in Consecutive Periods*

P **at** $t = 1$	**Expert Forecast**	τ	**Bayesian Forecast at** $t = 2$	**Price at the End of** $t = 2$
102.26	1.699%	0.00004	1.671%	103.97
101.73	2.223%	0.00043	1.693%	103.45
101.68	2.281%	0.00166	1.671%	103.38

Note: The first column displays the stock price at the beginning of $t = 2$. The second column shows the expert opinion at the end of $t = 1$ when there is no additional information. The third column gives the strength of the evidence, which is double that in the previous period. The fourth column shows the Bayesian forecast for the second period. The fifth column shows the price at the end of $t = 2$.

ing function implies that the investor will put almost all of his weight on the statistical evidence. In the third period, the expert consensus is that the stock price will rise 0.60% and the consecutive investor's forecast is 1.62%, bidding up the price of the stock to 105.1. Because the expert information containing costly public information is presented rather fuzzily, the rational investor is not able to filter the news out of the expert information and acts almost as an uninformed investor (who would always forecast the return to be 1.668%). As a result, the investor will overreact and buy additional units of the stock.

Conclusion

Descriptive behavioral models provide explanations for observed anomalies assuming that financial agents are irrational. The violation of the rationality axiom is taken to be true for financial agents relying on general findings from psychology. We explore whether the momentum anomaly (so far unexplained by risk) can also numerically be found when markets are efficient and agents are rational. We assume that there is an information market imperfection in that the strength of the expert evidence is noisy. Hence, the costly public information they reflect through the forecasts is slowly diffused through the markets and prices do not fully reflect all costly public information. According to the Bayesian forecasts and the weighing functions for the evidence we computed, momentum can exist in efficient markets where agents are rational. This is even true for samples of liquid stocks.

Our findings also imply that when the strength of the evidence increases, momentum strategies will be less beneficial. Indeed, for U.S. data, Kinney, Burgstahler, and Martin [1999] report that the accuracy of expert opinions has increased during the 1990s. Also, the precision of the communicated evidence increases. Hence, this study of normative behavior also implies that an improvement of this segment of the market microstructure will increase market efficiency. More importantly, experts should be able to communicate fully independent opinions regardless of their reputation. If that is the case and they have a strong consensus about the stock's value or payoff, prices will better reflect available information.

Notes

1. Examples of cognitive failures or irrationalities are overconfidence (when people overestimate the reliability of their own knowledge) (De Bondt and Thaler, 1995), and conservatism (the slow updating of beliefs relative to new evidence) (Edwards, 1968).
2. The noise in the precision of the information is also denoted as the inverse strength of the evidence and stands here for the dispersion in the expert opinion.
3. We denote a language as the calculus used to translate evidence into a judgment.
4. Results are available upon request.

References

Barberis, N., A. Shleifer, and R. Vishny. "A Model of Investor Sentiment." *Journal of Financial Economics,* 49, 1998.

Bawa, V.S., S.J. Brown, and R.W. Klein. "Estimation Risk and Optimal Portfolio Choice." In Arnold Zellner and Joseph B. Kadane, eds., *Studies in Bayesian Econometrics,* Vol. 3. Amsterdam: North Holland, 1979.

Campbell, J.Y., A.W. Lo, and A.C. MacKinlay. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press, 1997.

Carhart, M.M. "On Persistence in Mutual Fund Performance." *Journal of Finance,* 52, 1997.

Cochrane, J. *Asset Pricing*. Princeton, NJ: Princeton University Press, 2001.

Daniel, K., D. Hirschleifer, and A. Subrahmanyam. "A Theory of Overconfidence, Self-Attribution, and Security Market Under- and Over-reactions." *Journal of Finance,* 53, 1998.

De Bondt, W.F.M., and R. Thaler. "Financial Decision-Making in Markets and Firms: A Behavioral Perspective." In R.A. Jarrow, V. Maksimovic, and W. Ziemba, eds., *The Handbook of Operational Research and Management Science*. Elsevier Science, 1995.

Dechow, P., A.P. Hutton, and R.G. Sloan. "An Empirical Assessment of the Residual Income Valuation Model." *Journal of Accounting and Economics,* 26, 1999.

Dreman, D.N., and E.A. Lufkin. "Investor Overreaction: Evidence that its Basis is Psychological." *Journal of Psychology and Financial Markets,* Vol. 1, No. 1, 2000.

Edwards, W. "Conservatism in Human Information Processing." In B. Kleinmutz, ed., *Formal Representation of Human Judgment*. New York: John Wiley and Sons, 1968.

Einhorn, H.J., and R.M. Hogarth. "Behavioral Decision Theory: Processes of Judgment and Choice." In D.E. Bell, H. Raiffa, and A. Tversky, eds., *Decision Making, Descriptive, Normative, and Prescriptive Interactions*. Cambridge: Cambridge University Press, 1988.

Fama, E.F., and K.R. French. "Common Risk Factors in the Returns on Stock and Bonds." *Journal of Financial Economics,* 33, 1993.

Fama, E. "Market Efficiency, Long-Term Returns, and Behavioral Finance." *Journal of Financial Economics,* 49, 1998.

Geweke, J. "Exact Inference in the Inequality Constraint Normal Linear Regression Model." *Journal of Applied Econometrics,* 1, 1986.

Geweke, J. "Bayesian Inference in Econometric Models Using Monte Carlo Integration." *Econometrica,* 57, 1989.

Griffin, D., and A. Tversky. "The Weighing of Evidence and the Determinants of Confidence." *Cognitive Psychology,* 24, 1992.

Grossman, S., and J.E. Stiglitz. "On the Impossibility of informationally Efficient Markets." *American Economic Review,* 70, 1980.

Heckelei, T., and R.C. Mittelhammer. "Bayesian Bootstrap Inference via Regression Structure Likelihood." In R. Carter Hill, ed., *Bayesian Computational Methods and Applications*. *Advances in Econometrics,* Vol. 11, 1996, part A, and Greenwich, CT and London: JAI Press.

Hong, H., T. Lim, and J.C. Stein. "Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies." *Journal of Finance,* 55, 2000.

Hong, H., and J.C. Stein. "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets." *Journal of Finance,* 54, 1999.

Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance,* 48, 1993.

Jegadeesh, N., and S. Titman. "Profitability of Momentum Strategies: An Evaluation of Alternative Explanations." NBER Working paper #7159, 1999.

Jorion, P. "Bayesian and CAPM Estimators of the Means: Implications for Portfolio Selection." *Journal of Banking and Finance,* 15, 1991.

Kinney, W., D. Burgstahler, and R. Martin. "The Materiality of Earnings Surprises." Working paper, University of Texas at Austin, University of Washington, and Indiana University, 1999.

Kloek, T., and H.K. Van Dijk. "Bayesian Estimates of Equation System Parameters: An Application of Integration by Monte Carlo." *Econometrica,* 46, 1978.

Lenk, P., and M. Wedel. "Bayesian Econometrics: A Reaction to Geweke." *Journal of Econometrics,* 100, 2001.

Lim, T. "Rationality and Analysts' Forecast Bias." *Journal of Finance,* 56, 2001.

Markowitz, H. "Portfolio Selection." *Journal of Finance,* 7, 1952, pp. 77–91.

Ohlson, J.A. "Earnings, Book Values and Dividends in Equity Valuation." *Contemporary Accounting Research,* 11, 2, 1995.

Parkinson, M. "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business,* 53, 1980.

Rouwenhorst, G. "International Momentum Strategies." *Journal of Finance,* 53, 1998.

Appendix

FIGURE A-1
The General Bayesian Bootstrap Regression Algorithm

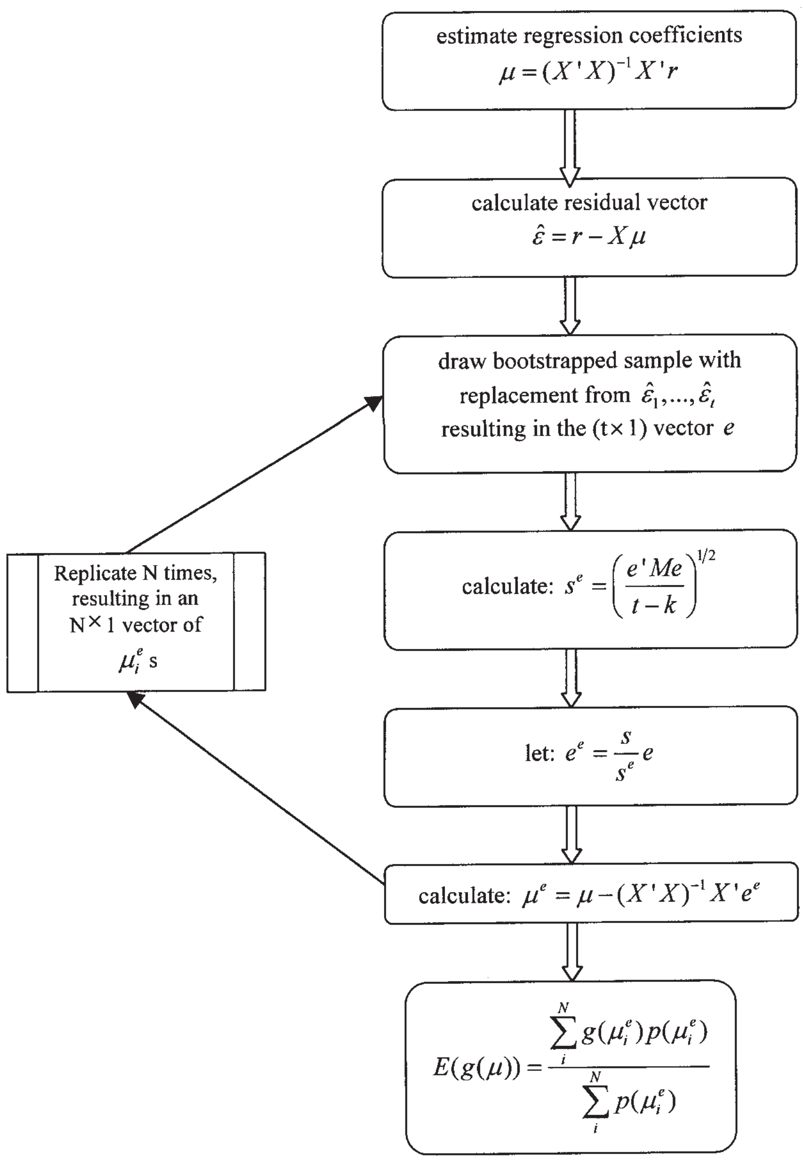

Note: N denotes the number of replications and the superscript e denotes a drawing from the empirical distribution.