Target prices forecast quality and analysts' forecast performance*

Patrice Fontaine† Tristan Roger‡

February 11, 2014

Abstract

We show, in this paper, that measuring the forecast error (*i.e.* the accuracy) of a target price is not sufficient to assess its quality, because the stock price predictability (which depends on the stock return volatility and on the forecast horizon) is likely to vary across stocks and over time. We argue that the evidence of time persistent differences in accuracy, obtained in previous studies, cannot be interpreted as a proof of analysts possessing differential abilities to forecast stock prices. We show that, when replacing the empirical target prices by naive forecasts, persistent differences in forecast errors remain. Our analysis indicates that the persistence of differential forecast errors is driven by persistence in stock return volatility. We introduce a measure of target price forecast quality that considers both the forecast error and the difficulty of issuing a correct forecast. We provide a methodology to estimate the difficulty of forecasting stock prices. Our empirical analysis reveals that, when forecasting difficulty is taken into account, financial analysts do not exhibit differential abilities to forecast future stock prices.

Keywords: Financial analysts, Target prices, Forecast quality, Forecast accuracy, Stock price predictability, Persistence of volatility

*The authors wish to thank Francesco D'Acunto, Brian Ayash, Mark Bradshaw, Patricia Dechow, François Degeorge, François Derrien, Michel Dubois, Ivan Guidotti, Isaac Hacamo, Sonia Jimenez-Garces, Alastair Lawrence, Martin Lettau, Franck Moraux, Terrance Odean and Richard Sloan, as well as participants of the UC Berkeley Finance pre-seminar, Dauphine Finance Seminar and 2013 AFFI Conference for their many helpful comments and suggestions.

†CNRS (Eurofidai). Email: patrice.fontaine@eurofidai.org

‡Corresponding author. Paris Dauphine University. Email: tristan.roger1@gmail.com. This article was written in part when Tristan Roger was a Visiting Scholar at the Haas School of Business, University of California, Berkeley. Tristan Roger gratefully acknowledges the financial support of Région Rhône-Alpes

1 Introduction

Professional investors, [...], fail a basic test of skill: persistent achievement.
Daniel Kahneman (2011)

We show, in this article, that measuring the forecast error (*i.e.* the accuracy) of a target price is not sufficient to assess whether it is a good or a bad forecast, because the difficulty to issue an accurate forecast is likely to vary across stocks and over time. The importance of predictability (*i.e.* the difficulty to issue an accurate forecast) is a well-known concern in the literature on earnings forecasts (Huberts and Fuller, 1995; DeBondt and Forbes, 1999; Beckers et al., 2004). For instance, Jacob et al. (1999) note that "forecasting difficulty is [...] like to differ cross-sectionally". Similarly, Hong et al. (2000) state that "some firms are more difficult than others to predict accurately". The existing literature on target prices ignores this issue and considers that a target price is better than another if the forecast error is smaller. We provide evidence that omitting the issue of stock price predictability (which is a function of the stock return volatility and of the forecast horizon) prevents from correctly evaluating the ability of financial analysts to forecast future stock prices. Furthermore, we show that the controls often used in multivariate analysis cannot be used to neutralize differences in predictability as the relationship between volatility and target price accuracy is nonlinear.

We propose, in this paper, a new measure to evaluate the forecast quality of target prices and a comprehensive framework to evaluate the ability of financial analysts to forecast future stock prices. The main contributions of our approach are the following: (1) our new measure takes into account both the forecast error (the accuracy) and the difficulty of issuing the target price (the predictability); and, (2) our measure can be used in a dynamic setting (as opposed to the *ex-post* design of traditional measures). When taking into account differences in forecast difficulty across target prices, we show that analysts do not exhibit differential abilities to forecast future stock prices. Our results contrast with previous studies (Bradshaw et al., forthcoming; Bilinski et al., 2013).

Analysts play a key role in financial markets. Ivkovic and Jegadeesh (2004) identify two sources of value that analysts bring to the market. First, they extract useful information for investors from public information. Second, through a careful examination of accounting documents and contacts with firm managers, they acquire information previously unknown to other market participants. As such, their reports and forecasts are of great importance since they render private information public. Analysts' reports are typically composed of three main figures: earnings forecasts, purchase recommendations and target prices. Though the latter seems to be of greatest interest to investors (as it gives a precise indication as to whether a stock is under- or overvalued), it has also received the least attention from academics. However, a few articles point out the important role of target prices both for individual investors and practitioners. For instance, Brav and Lehavy (2003) and Asquith et al. (2005) report significant market reaction to target price revisions, even after controlling for recommendations and earnings forecast revisions. Lawrence et al. (2012) analyze the web traffic of a leading website of analyst report information and find that target prices are the type of analyst information most requested by investors.

In the first part of the paper, we provide evidence that the forecast error by itself is unsuitable to evaluate the ability of financial analysts to forecast future stock prices. We show that forecast errors strongly depend on stock return volatility. We demonstrate that this relationship between forecast errors and volatility is mechanical, and not the result of analysts being particularly good at issuing accurate target prices for low volatility stocks. We establish that the relationship between forecast errors and volatility is nonlinear. This nonlinearity prevents from using simple controls such as including the volatility in multivariate regressions. Finally, we provide evidence that ignoring the dependence between accuracy and predictability has important consequences when determining whether financial analysts exhibit differential skills to forecast future stock prices. Previous studies (Bradshaw et al., forthcoming; Bilinski et al., 2013) report that analysts' forecast errors exhibit persistent differences. The authors interpret these findings as a proof of analysts possessing differential skills to issue target prices.[1] We provide a simple proof that these persistent differences in accuracy cannot be interpreted as analysts exhibiting differential skills. Our approach is the following one. We replicate the analysis by replacing the actual target prices by naive forecasts. Because naive target prices are determined by a mechanical rule, our (naive) analysts cannot exhibit differential skills. Hence, if we observe persistent differences in forecast errors, it means that these differences cannot be interpreted as evidence of differential skills. Our results indicate that our (naive) analysts do exhibit persistent differences in accuracy. Thus, differences in forecast errors arise mainly from differences in the firms covered by analysts. Financial analysts who cover a pool of

[1] However, Bradshaw et al. (forthcoming) stress out that, although statistically significant, their results are economically weak and unlikely to motivate investors to utilize target prices of individual analysts based on past performance.

stocks with low volatility tend to exhibit higher accuracy (lower forecast errors).

In the second part of the paper, we introduce a new measure to evaluate the forecast quality of target prices. Our measure of forecast quality takes into account both the target price accuracy and the forecast difficulty. We define the accuracy as the absolute forecast error $|S_T - TP_{t,T}|$, where $TP_{t,T}$ is the target price issued at time t with horizon $T-t$ and S_T is the stock price at the end of the horizon.[2] A simple way to measure the difficulty of issuing an accurate target price is to estimate the expected value of the absolute forecast error $E_t[|S_T - TP_{t,T}|]$. The expected value of the absolute forecast error is an increasing function of the stock return volatility and the length of the forecast horizon. We then define the target price forecast quality as the abnormal absolute forecast error, denoted $E_t[|S_T - TP_{t,T}|] - |S_T - TP_{t,T}|$ where $E_t[|S_T - TP_{t,T}|]$ measures the difficulty of issuing an accurate forecast (the expected forecast error) and $|S_T - TP_{t,T}|$ measures the forecast accuracy (the realized *ex-post* forecast error).

An important issue here is to find a way to estimate the expected value of the absolute forecast error $E_t[|S_T - TP_{t,T}|]$. First, we note that the absolute forecast error $|S_T - TP_{t,T}|$ corresponds exactly to the final payoff of a a portfolio containing a call option and a put option on the same underlying stock; the two options are characterized by the same strike price $TP_{t,T}$ and the same maturity $T-t$. This type of options portfolio is called a straddle in the option literature. Second, we note that the price of the straddle at time t is equal to the discounted expected value of the final payoff. It follows that the difficulty of issuing an accurate target price can be estimated by the price of a straddle. When issuing a target price, an analyst acts as if she shorts a straddle. Selling a straddle implies that she receives up front the discounted expected value of the absolute forecast error $E_t[|S_T - TP_{t,T}|]$. She will then pay, at the end of the horizon, the realized absolute forecast error $|S_T - TP_{t,T}|$ (which corresponds to the target price accuracy). The gain or the loss that results from this operation is exactly equal to our measure of target price forecast quality.

The contributions of our approach are the following. First, we provide a coherent way to incorporate both the target price accuracy and the difficulty of issuing accurate target prices. Second, we provide a way to estimate the target price forecast quality at any point in time: the target price forecast quality at time $t+\tau$, $\tau \in [0; T-t]$ is simply equal to $E_t[|S_T - \Phi_{t,T}|] - E_{t+\tau}[|S_T - \Phi_{t,T}|]$, that is, the price of the straddle at time t

[2]The absolute forecast error is usually defined as $\frac{|S_T - TP_{t,T}|}{S_t}$. However, to simplify notations, we normalize so that $S_t = 1$.

(capitalized until time $t+\tau$) minus the price of the straddle at time $t+\tau$. This last feature is particularly important as it provides a simple and consistent solution to the issue of measuring, when a revision occurs, the forecast quality of the initial target price.

In the last part of the paper, we address the question of whether financial analysts exhibit persistent differential abilities. The design of our new measure ensures that differences in stock price predictability do not interfere with our analysis. Kahneman (2011) exposes two basic conditions for the possible existence of expertise skills: (1) an environment that is sufficiently regular to be predictable; and (2) an opportunity to learn the regularities through prolonged practice. Although the second condition is met, the first condition is not likely to be fulfilled. Indeed, one can hardly define the stock market as a regular environment. Even without assuming any kind of market efficiency, predicting stock prices at a 12-month or longer horizon is an extremely difficult task. We therefore do not expect financial analysts to demonstrate differential skills in forecasting future stock prices. Our analysis indicates that the differences in target price forecast quality are not persistent. Thus, financial analysts do not have differential skills to forecast future stock prices. Our findings are robust to a number of changes such as restricting the sample to experienced financial analysts, or restricting the validity of target prices to a shorter period of time.

2 Data and descriptive statistics

Our primary dataset consists of a total of 686,863 target prices issued by 10,137 analysts (620 brokers) on 7,646 U.S. stocks for the 2000-2010 period. The provider of the target prices is I/B/E/S. For each forecast, we have the code of the analyst (and the broker code) who issues the forecast, the issue date, the horizon in months (usually 6 or 12 months), and the target price. We remove from the database the forecasts for which the stock price is not available on the issue date (20,766 forecasts), or is less than one dollar (2,044 forecasts). We also delete from the database the forecasts for which the ratio of the target price over the stock price is in the bottom one percent of the distribution (7,468 forecasts) and the forecasts for which this ratio is higher than four (2,313 forecasts). Finally, we discard the observations for which the price history is too short to compute an acceptable estimation of the historical volatility[3] (4,801 forecasts). After deleting these 37,392 observations, 770

[3]We delete the observations for which there are less than 2 months of price history prior to the forecast.

analysts are removed from the database as they are left with no forecasts. Our final sample consists of 649,471 target prices issued by 9,367 analysts (583 brokers) on 7,268 stocks.

Our secondary dataset consists of the prices, market capitalizations and volumes of trading for the 7,268 stocks considered. This second dataset comes from CRSP. Target prices and stock prices are adjusted for splits and corporate actions.

Table 1 reports for each year the number of forecasts, number of active analysts, the average, median and maximum number of active analysts per stock, and the average, median and maximum number of stocks covered per analyst. We observe that the number of forecasts per year more than doubles over the sample period while the number of active analysts remains roughly constant. It appears that the inclusion of a target price in analysts' reports is an increasingly popular practice. An analyst typically covers 4 different stocks at the beginning of the sample period; this number increases to 7 in the last years. Conversely, the number of analysts covering a given stock increases over the sample period from 4 to 6.

On average, the analysts in our sample revise their forecasts approximately every 6 months (137 trading days). The target prices in our sample are on average 23% higher than the current stock price.[4] This statistic is similar to what can be observed for other periods and/or countries. For instance, Brav and Lehavy (2003) find that target prices on U.S. stocks for the 1997-1999 period are on average 28% higher than the current price while Kerl (2011) reports an implicit return of 18.07% for German stocks for the 2002-2004 period. Finally, it appears that the analysts in our sample are mainly optimistic about future stock prices with only 13% of the target prices being below the concurrent price.

3 Target price accuracy and stock price predictability

We show, in this section, that measures of accuracy such as the absolute forecast error (AFE) are not well-suited to evaluate whether a forecast is better than another or to evaluate analysts' performance. The main reasons why measuring the accuracy is unsuitable to assess whether a target price is a good or a bad forecast are: (1) forecast errors strongly depend on stock price predictability (*i.e.*, the difficulty of issuing an accurate target price); and, (2) the *ex-post* design of accuracy measures implies a low feedback speed and pre-

[4]The annualized return for the S&P 500 Composite index for the same period is -1.21%.

vents a correct evaluation of the accuracy when target price revisions occur.[5] We provide both empirical and theoretical evidence of the existence of these different issues. We also show that these issues have an influence when studying the ability of financial analysts to forecast stock prices and can bias the economic findings.

The most popular measure of forecast accuracy is the absolute forecast error (AFE). We note, however, that our results also apply to other measures of accuracy. The absolute forecast error is defined as

$$AFE_{j,t} = \frac{|S_T - TP_{t,T}|}{S_t}, \tag{1}$$

where $TP_{t,T}$ is the value of a target price issued at time t with horizon T, S_T is the stock price at the end of the forecast horizon and S_t is the stock price at the time the target price was issued.

3.1 Relationship between target price accuracy and stock return volatility

The stock price (un)predictability corresponds to the difficulty of issuing an accurate target price. This difficulty is a function of both the stock return volatility and the target price horizon. As a consequence, there exists a mechanical relationship between target price accuracy and stock return volatility.[6] We provide here both empirical and theoretical evidence that target price accuracy is mechanically influenced by stock return volatility. Furthermore, we show that this relationship is nonlinear. In the rest of the paper, we use the Absolute Forecast Error (AFE) measure to evaluate target price accuracy. However, our results hold if we use different specifications for the accuracy.

Each year, we assign target prices to five quintiles with respect to the volatility of the underlying stock. For each quintile, we report the average AFE of the target prices in the quintile. Table 2 provides the average AFE per quintile for the 2000-2010 period. Panel A reports the results using actual target prices. Panel B provides the average AFE using naive forecasts. We build the naive forecasts so that the implied stock return $(TP_{t,T} - S_t)/S_t$ of

[5] The accuracy can be evaluated only at the end of the target price horizon, usually 12 months

[6] We do not consider here the issue of target price horizon as most studies consider only target prices with a 12-month horizon. However, it is easy to show that a negative relationship exists between target price accuracy and forecast horizon.

a 12-month horizon target price is equal to the 12-month risk free rate. We report results using naive forecasts in order to eliminate the possibility that the relationship between AFE and stock return volatility ensues entirely from financial analysts being particularly good at forecasting stock prices for high volatility firms. Our findings indicate a strong monotonic relationship between stock return volatility and AFE. This result holds both for actual data and naive forecasts.

In multivariate analyses, a usual way to account for the dependence between two variables is to incorporate a control variable. However, we show that the solution to use volatility as a control variable does not apply in our context as the relationship between AFE and stock return volatility is nonlinear. In order to demonstrate this nonlinearity, we perform the following regression

$$AFE_{jt} = \alpha + \sum_{k=1}^{10} \beta_k \mathbf{1}_{jt}^{k} \sigma_{jt} + \epsilon_{jt}, \qquad (2)$$

where AFE_{jt} is the absolute forecast error of a target price on firm j issued at time t, σ_{jt} is the stock return volatility of stock j measured at time t and $\mathbf{1}_{jt}^{k}$ is an indicator that takes the value 1 if the stock return volatility σ_{jt} belongs to the k-th volatility decile and 0 otherwise.

The idea underlying this regression is the following one. In case of perfect linearity, all the coefficients β_k will take the same value. On the contrary, if the relationship between AFE and volatility is nonlinear, we will find differences in the values taken by the coefficients β_k. For instance, if the coefficient β_k decreases (increases) with k, the relationship between volatility and AFE is concave (convex). Table 3 reports the results of the regression. The coefficient β_k increases with k, indicating that the relationship between AFE and volatility is nonlinear and convex. We provide, in Appendix A, theoretical evidence of the nonlinearity of the relationship between AFE and volatility.

This nonlinearity prevents from measuring analysts' performance with relative measure of accuracy and prevents from using simple controls such as including the volatility in multivariate regressions.[7]

3.2 Target price accuracy and revisions

Target prices are generally issued with a 12-month horizon. However, in practice, analysts often revise their forecasts before the end of this horizon. The accuracy measures that can be found in the literature on target prices are *ex-post* measures. It follows that, when a revision occurs, the evaluation of the accuracy of the initial target price is problematic. Bonini et al. (2010) provide an interesting insight on the issue of evaluating accuracy when forecast revisions occur. When a target price is revised, the two possible approaches to address the problem are:

(1) to consider the revised target price and the initial target price as being two distinct forecasts (the accuracy is measured on two partially overlapping periods); and,

(2) to assess the accuracy of the initial target price by adjusting the time horizon and to consider the revision as a new target price.

To illustrate the issues associated with these two approaches, we provide the following examples. Assume that an analyst issues the 2^{nd} of July 2012 (date t), a 12-month target price of \$50 on Facebook, with the current price at date t being \$30. One month later, the analyst becomes convinced the stock price is overvalued. She revises her forecast and sets a new 12-month target price at \$10 (while the current stock price is \$20).

If we consider the first approach, we now have two forecasts on Facebook; the first one is at \$50 starting at the beginning of July 2012, and the second one is at \$20 starting at the beginning of August 2012. As noted by Bonini et al. (2010), there are a number of problems associated with this approach. Mainly, there is no obvious economic meaning

[7]The literature on earnings forecasts uses relative measures of accuracy to account for differences in predictability. That is, the accuracy of a forecast is determined with respect to the accuracy of other forecasts issued in similar conditions (*i.e.*, forecasts issued on the same firm and during the same period of time). Clement (1999) proposes to measure the analysts' performance by comparing the analyst's absolute forecast error to the average absolute forecast error of other analysts following the same stock during the same time period. Hong et al. (2000) propose an alternative way to control for differences in earnings predictability. For a given firm and a given year, they rank analysts with respect to the absolute forecast error of their most recent forecast. The rankings are then transformed into scores. In the case of target prices, relative measures cannot be used because the end of the forecast horizon depends on the issue date (in addition to the issue of the nonlinearity).

in considering two opposite forecasts. (One forecast is higher than the stock price and the revised one is lower.) Indeed, if the stock price of Facebook decreases, the analyst is considered simultaneously inaccurate (on the first target price) and accurate (on the second target price).

What happens if we follow the second approach? Two successive forecasts are considered: first, a 1-month forecast with a target price of \$50, and second, a 12-month forecast with a \$10 target price. This second approach also has obvious drawbacks. According to this second approach of target price revisions, the accuracy of the first target price is measured by considering the first forecast as a 1-month target price of \$50. In other words, this assumption transforms an annualized expected return of 67% when issued with a 12-month horizon into an monthly expected return of 67%!

This simple example shows that dealing smoothly with revisions is not so easy when using *ex-post* measures. We will show, in section 4, that it is possible to remedy the problem by using the expected value of the absolute forecast error to estimate the accuracy at the time the revision occurs.

3.3 Persistent differences in accuracy and stock return volatility

We show in this section that the persistence in absolute forecast errors (AFE) found in previous studies is mechanically caused by persistence in volatility. The relationship between AFE and stock price predictability implies that analysts covering less volatile stocks are more accurate. It follows that persistent differences in AFE cannot be interpreted as a proof of persistent differential abilities.

In this subsection, we define the analyst's performance as in Bradshaw et al. (forthcoming) and Bilinski et al. (2013). For a given period $]t-1;t]$, we evaluate the analyst's performance as the average of the AFE of all the target prices she issued during that period. We are also interested in the volatility of the stocks covered by each analyst. For each target price issued during the $]t-1;t]$ period, we calculate the historical stock return volatility for the 6 months preceding the target price issue date. We then take the average of the volatilities computed for all the stocks for which the analyst issued target prices during that period.

Our analysis consists in assigning analysts to five quintiles with respect to their performance over a measurement period $]t-1;t]$. The persistence is then measured by estimating the performance over the test period conditional on their ranking over the measurement period. The test period is defined as $]t+\theta;t+\theta+1]$ where $\theta = 12$ months. A lag $\theta = 12$ months is added between the measurement period and the test period to insure that the two periods do not overlap.[8] We therefore avoid mechanically inducing a positive relation between current and subsequent analysts' AFE. Financial analysts exhibit persistent forecast if the most (least) accurate analysts over the measurement period $]t-1;t]$ are ranked in the highest (lowest) quintile for the test period $]t+\theta;t+\theta+1]$ and if the difference between the AFE of quintile 1 and the AFE of quintile 5 is statistically different from zero.

Table 4 presents the results using the target prices in our sample. As our database is similar to the one of Bradshaw et al. (forthcoming), it is not surprising that we obtain similar results. Table 4 indicates that analysts with a lower average AFE in the measurement period also have a lower average AFE in the test period. When the analysts' performance is computed using a quarterly frequency, analysts in the first quintile (last) exhibit an average AFE of 0.1551 (0.9693) in the measurement period and an average AFE of 0.3322 (0.5779) in the test period. We observe, however, that the stock return volatility is lower for target prices issued by the analysts in the first quintile. This volatility increases with the quintiles, both in the measurement period and in the test period. This result indicates that if we observe persistent differences in AFE, we also observe persistent differences in volatility.

We want to check if the persistence in volatility causes the persistence in AFE. In order to do so, we replace the target prices in our data with naive forecasts. We build our naive forecasts so that the implied stock return $(TP_{t,T} - S_t)/S_t$ of a 12-month horizon target price is equal to the 12-month risk free rate. As our new target prices result from a mechanical rule, observing persistent differences in AFE would imply that accuracy strongly depends on the stock return volatility. It would follow that the accuracy is therefore unsuitable to evaluate the ability of financial analysts to forecast future stock prices.

As can be seen in Table 5, the differences in AFE are persistent even though target

[8] If the test period was $]t;t+1]$, the accuracy of target prices issued at the end of the measurement and the accuracy of target prices issued at the beginning of the test period would be artificially correlated. Indeed, if we consider two target prices issued at time $t-\epsilon$ (in the measurement period) and $t+\epsilon$ (in the test period), there would be an overlap on the period $]t+\epsilon;t-\epsilon+\theta]$. This autocorrelation would cause an artificial persistence.

prices are issued as a result of a mechanical rule. When the analyst's performance is computed using a quarterly frequency, the average AFE ranges from 0.1317 to 0.7614 in the measurement period. In the test period, we observe a significant difference of 0.1424 between the first and the last quintiles. This result indicates that the analysts who were the most (least) accurate in the measurement period are still the most (least) accurate in the test period. We also observe a strong persistence in volatility. Given that we considered naive forecasts, we can conclude that the persistence in AFE is mainly driven by the persistence in volatility. It follows from our analysis that the accuracy cannot be used to evaluate whether analysts exhibit differential skills to forecast future stock prices.

4 Target price forecast quality

4.1 *Ex-post* measure of target price forecast quality

The *ex-post* absolute forecast error is equal to $|S_T - TP_{t,T}|$ and corresponds to the target price accuracy. As shown in the previous section, in order to evaluate whether a target price is a good or a bad forecast, one needs to take into account the stock price predictability, which is a function of the stock return volatility and the length of the forecast horizon. The difficulty of issuing an accurate target price can be evaluated by the expected value of the absolute forecast error $E_t[|S_T - TP_{t,T}|]$. For a given stock price $S_t^A = S_t^B$ and a target price $TP_{t,T}$, the expected value of the absolute forecast error $E_t\left[\left|S_T^A - TP_{t,T}\right|\right]$ on stock A is higher than the expected value of the absolute forecast error $E_t\left[\left|S_T^B - TP_{t,T}\right|\right]$ on stock B if the stock return volatility of stock A is higher than the stock return volatility of stock B. Similarly, the expected value of the absolute forecast error of a target price issued with a longer horizon is higher than the expected value of the absolute forecast error of a target price issued with a shorter horizon.

We define the target price forecast quality as the difference between the forecast difficulty, $E_t[|S_T - TP_{t,T}|]$, and the forecast accuracy, $|S_T - TP_{t,T}|$. The absolute forecast error $|S_T - TP_{t,T}|$ corresponds to the final payoff of a straddle with a strike price equal to $TP_{t,T}$, that is, a portfolio containing a call option and a put option on the same underlying stock; the two options are characterized by the same strike price and the same maturity. A possible way to estimate the expected value of the absolute forecast error $E_t[|S_T - TP_{t,T}|]$ is therefore to consider the price of the straddle at time t (capitalized until time T). We define the *ex-post* measure of target price forecast quality as follows.

Definition 1 *The ex-post forecast quality $TPFQ_{t,T}$ of a target price $TP_{t,T}$ issued at time t on a stock S, with an horizon equal to $T-t$, is defined as*

$$\begin{aligned} TPFQ_{t,T} &= E_t\left[|S_T - TP_{t,T}|\right] - |S_T - TP_{t,T}| \\ &= \left(C_t + P_t\right) e^{r(T-t)} - \left(C_T + P_T\right) \end{aligned} \tag{3}$$

where C_t (P_t) is the price at time t of a call (put) option on the stock S with maturity date T and strike price $TP_{t,T}$.

As we need to compare the forecast quality of target prices issued on stocks with different price levels, we require the measure of forecast quality to be homogeneous of degree one (we do not want the stock price level to influence the measure of forecast quality). This means that we assume the stock price to be equal to 1 at the time the target price is issued. We write $S_t = 1$ and we adjust the target price $TP_{t,T}$ accordingly.

Assuming that the stock price follows a geometric Brownian motion, the *ex-post* forecast quality $TPFQ_{t,T}$ of a target price issued at time t can be calculated according to the Black and Scholes (1973) model (see Appendix B).[9]

4.2 Properties

Our measure is composed of two components: (1) the expected value of the absolute forecast error (which estimates the difficulty of issuing an accurate target price); and, (2) the *ex-post* absolute forecast error (which is the traditional estimation of the accuracy of a target price). The difficulty of issuing an accurate target price is positively related to the stock return volatility σ_t (*e.g.* it is more difficult to forecast the future price of a stock with high volatility than to forecast the future price of a stock with low volatility) and to the length of the forecast horizon $T-t$ (*e.g.* the task of issuing an accurate target price with a 24-month horizon is more difficult than the task of issuing an accurate target price with a 12-month horizon). Therefore, our measure of forecast difficulty $E_t\left[|S_T - TP_{t,T}|\right]$

[9]For simplicity's sake, we choose to compute the value of the straddle using Black and Scholes (1973) model. However, our measure of target price forecast quality could be extended to more complex models of option pricing.

must satisfy the two following requirements: (1) it must increase with the stock return volatility σ_t; and, (2) it must increase with the length of the forecast horizon $T - t$.

Proposition 2 *For a given final stock price S_T and a given target price $TP_{t,T}$, the forecast quality $TPFQ_{t,T}$ is an increasing function of the stock return volatility σ_t and of the length of the horizon $T - t$.*

We provide a proof for this proposition in Appendix D.

4.3 Target price forecast quality in a dynamic setting

We propose a dynamic setting in which the target price forecast quality can be estimated at any point in time. The target price forecast quality at time $t+\tau$, $\tau \in [0; T-t]$ is simply equal to $E_t [|S_T - \Phi_{t,T}|] - E_{t+\tau} [|S_T - \Phi_{t,T}|]$, that is, the price of the straddle at time t (capitalized until time $t + \tau$) minus the price of the straddle at time $t + \tau$.

Definition 3 *The forecast quality $TPFQ_{t,t+\tau}$ at time $t+\tau$, $\tau \in [0; T-t]$ of a target price $TP_{t,T}$ issued at time t on a stock S, with an horizon equal to $T - t$, is defined as*

$$\begin{aligned} TPFQ_{t,t+\tau} &= E_t [|S_T - TP_{t,T}|] - E_{t+\tau} [|S_T - TP_{t,T}|] \\ &= (C_t + P_t) e^{r\tau} - (C_{t+\tau} + P_{t+\tau}), \end{aligned} \tag{4}$$

where C_t (P_t) is the price at time t of a call (put) option on the stock S with maturity date T and strike price $TP_{t,T}$.

At the time the target price is issued ($\tau = 0$), the target price forecast quality $TPFQ_{t,t}$ is equal to 0. This property translates the idea that at the time the target price is issued, one does not know yet whether it is a good or a bad forecast. Note that when $\tau = T - t$, we retrieve the *ex-post* measure of forecast quality defined in equation 3.

Remark 4 *The stock return volatility is set constant for a given forecast. That is, once a target price is set, we use the stock return volatility at the time the forecast was issued to evaluate the value of the straddle (that is, to estimate $E_{t+\tau} [|S_T - TP_{t,T}|]$). This same* *volatility is used until a new target price (revision) is issued. If the volatility was not set constant, an increase in volatility would mechanically lower the target price forecast quality. With this methodology, we distinguish between target prices issued on stocks with different volatilities (cross-section) but the variations of volatility over time do not influence the target price forecast quality (time-series).*

4.4 Target price revisions

In practice, analysts often revise the target prices before the end of the horizon. We consider the initial forecast and the revision as two separate forecasts. Once a revision occurs at time $t + \tau$, the first forecast is no longer valid. However, we need to evaluate the forecast quality of the initial target price over the period $]t; t + \tau]$. It follows from the previous definition that the forecast quality of the initial target price at time $t + \tau$ is simply equal to $E_t [|S_T - TP_{t,T}|] - E_{t+\tau} [|S_T - TP_{t,T}|]$.

Let us consider the following example. An analyst issues a target price $TP^1_{t,T}$ at time t. She then revises her forecast, at time $t + \tau$, and issues a target price $TP^2_{t+\tau,T+\tau}$. The forecast quality over the period $]t; T + \tau]$ is simply equal to

$$TPFQ_{t,T+\tau} = TPFQ_{t,t+\tau} \left(TP^1_{t,T}, \sigma_t \right) + TPFQ_{t+\tau,T+\tau}(TP^2_{t+\tau,T+\tau}, \sigma_{t+\tau}). \tag{5}$$

where $TPFQ_{t,t+\tau}$ is the forecast quality, estimated at time $t+\tau$, of the initial target price $TP_{t,T}$, issued at time t with an horizon equal to $T - t$, and $TPFQ_{t+\tau,T+\tau}$ is the forecast quality, estimated at time $T+\tau$, of the revised target price $TP_{t+\tau,T+\tau}$, issued at time $t+\tau$ with an horizon equal to $(T + \tau) - (t + \tau)$.

4.5 Illustration of the target price forecast quality measure

To gain a clear understanding of how the forecast quality is computed, in Figure 1, we present an example of three target prices made on a stock by a given analyst. In this example, the risk-free rate is equal to 0. A first target price, equal to 45, is issued at time t_1 (while the stock price is equal to 35.76). The second target price (first revision), equal to 30, is issued at time t_2 (while the stock price is equal to 31.88). Finally, at time t_3, the analyst revises her forecast and announces a target price of 33 (while the stock price is equal to 49.34).[10]

At time t_1, we consider the price of a straddle with a strike price equal to 45. The price of the straddle[11] is computed using a 6-month historical volatility of $\sigma_{t_1} = 0.3905$ and is equal to 0.4091. The price of this straddle is used to evaluate the difficulty of the forecast. The target price forecast quality over the period $]t_1; t_2]$ is then equal to the value of the straddle at time t_1 (*i.e.*, the difficulty) minus the value of the straddle at time t_2 (*i.e.*, the accuracy). The price of the straddle at time t_2 is computed using the stock return volatility at the issue date t_1 ($\sigma_{t_1} = 0.3905$). The straddle at time t_2 is worth 0.4249. As a result, the target price forecast quality for the period $]t_1; t_2]$ is equal to $TPFQ_{t_1,t_2} = -0.0158$. As can be seen in Figure 1, the target price forecast quality for the period $]t_2; t_3]$ is equal to $TPFQ_{t_2,t_3} = -0.2901$. It follows that the target price forecast quality for the period $]t_1; t_3]$ is equal to $TPFQ_{t_1,t_3} = TPFQ_{t_1,t_2} + TPFQ_{t_2,t_3} = -0.3059$.

For the period $]t_1; t_2]$, the sensitivity of the forecast quality measure to the stock price (the delta of the forecast quality) is positive (because the stock price is below the target price). The target price forecast quality increases when the stock price increases (*i.e.*, gets closer to the target price) and decreases otherwise. The delta of the target price forecast quality becomes negative for the period $]t_2; t_3]$ as the target price is below the stock price. Because the target prices are close to the stock price, the sensitivity to the horizon is positive.

To illustrate the influence of stock return volatility on target price forecast quality, we assume in the previous example that the volatility is multiplied by 1.2 (everything else being equal). Figure 2 shows the forecast quality computed using the real volatility (solid line) and the forecast quality computed with a volatility set 20% higher (dashed line). Because the sensitivity of our measure to the volatility is positive, it appears on the graph that the target price forecast quality is higher when the volatility is higher.

5 Analysts' abilities to forecast future stock prices

We want to determine whether financial analysts exhibit genuine skills in forecasting future stock prices. Observing a positive value for the target price forecast quality is not sufficient

[10] Every forecast has a 12-month horizon.
[11] Remember that the stock price is normalized to 1.

to conclude to the existence of forecasting abilities. This positive target price forecast quality can be the result of skill, luck or simply the fact that by being overly optimistic, the forecast quality increases when the market rises. A first analysis of the existence of forecasting skills is to check whether analysts exhibit persistent differential forecast qualities. As stated by Kahneman (2011): "the diagnostic for the existence of any skill is the consistency of individual differences in achievement". We will thus consider an analyst to be skilled if she manages to consistently beat the other analysts.

5.1 Test of persistent differences in *ex-post* analysts' forecast performance

We have shown that observing persistent differences in forecast errors does not allow us to make any inference on the existence of differential forecasting abilities. We conduct the same analysis as before but we use our measure of target price forecast quality instead of the target price accuracy (AFE) to evaluate analysts' performance. As our measure of target price forecast quality takes into account the differences in volatility, a significant persistence of differences in target price forecast qualities would imply that financial analysts possess differential abilities to forecast target prices. In order to allow for a direct comparison with the previous results, we use the *ex-post* version of our measure (defined in equation 3). For a given period, we define the analyst's *ex-post* forecast performance ($exAFP$) as the average of the *ex-post* forecast quality of all the target prices she issued during that period.

The results in Table 6 show that when using a measure of forecast performance that accounts for differences in volatility, the persistence vanishes. In the measurement period, the analysts in the first quintile (*i.e.*, best performance) exhibit an average forecast quality of 0.3536 while the analysts in the last quintile (*i.e.*, worst performance) exhibit an average forecast quality of -0.2487. This difference in forecast quality between the first and the last quintile is not significant anymore in the test period. We obtain this result both for quarterly and semiannual periods.

5.2 Test of persistent differences in analysts' forecast performance in a dynamic setting

The main limitation of the *ex-post* measure of target price forecast quality used in the previous subsection is that it is necessary to introduce a 12-month lag between the measurement period and the test period. A second limitation is that the target price forecast quality is evaluated using the stock price at the end of the 12-month horizon. However, in practice, when a revision occurs the first forecast becomes inactive and only the revision is taken into account.

Our new measure, when used in a dynamic setting, allows us to consider revisions and to estimate the variations of analysts' performance on a daily basis. We introduce a measure of forecast performance where the analyst's forecast performance over the period $[t:T]$ is defined as the sum of her daily forecast performance over this period. We write

$$AFP_{t,T} = \sum_{\tau=t+1}^{T} \left(\frac{1}{J} \sum_{j=1}^{J} \left(TPFQ_{\tau_j,\tau} - TPFQ_{\tau_j,\tau-1} \right) \right), \tag{6}$$

where $AFP_{t,T}$ is the analyst's forecast performance over the period $[t:T]$, J is the number of outstanding target prices (defined by $\tau_j < \tau$) issued by the analyst and $TPFQ_{t_j,\tau}$ is the forecast quality of the target price issued on stock j at time t_j.

Contrary to *ex-post* measures, the analyst's forecast performance $AFP_{t,T}$ over a period $[t:T]$ can be measured using only information over this same period.[12] This feature allows us to evaluate the persistence over the short-run; we set the measurement period on $]t-1;t]$ and the test period on $]t;t+1]$ (we do not need anymore to add a lag between the measurement period and the test period). When using our new measure of analyst's forecast performance AFP, we need to restrict the sample period to the 2001-2010 period. Indeed, for 2000, we do not observe the target prices issued in 1999 which could still be outstanding. Therefore, we do not have the full portfolio of target prices for the first year of the sample. These unobserved target prices could influence the analyst's forecast performance (AFP).

The results in Table 7 show that, even over the short run, there are no persistent differ-

[12] When using *ex-post* measures, one needs to have information up to $T + 12$ in order to assess the analyst's forecast performance.

ences in analysts' forecasting abilities. Using both quarterly and semiannual frequencies, we do not observe any significant differences, in the test period, between the analysts' forecast performance AFP in the first quintile and the analysts' forecast performance in the last quintile.

5.3 Robustness check: the impact of learning

One reason why we might not observe any persistent differences in analysts' forecast performance (AFP) is that financial analysts learn over time and subsequently improve their forecast quality. If experience influences target price forecast quality, young inexperienced analysts will be ranked in the lower quintiles when they enter the sample period. They will then gradually move toward the highest quintiles as they learn and acquire experience. These agents would therefore add noise to our analysis of persistent differences in analysts' forecast performance (AFP).

We run the same analysis as above on a restricted sample containing only analysts with at least two years of experience. We compute, at time t, the experience of an analyst by observing the time $t - \tau$ at which she issued her first target price; the experience is then simply equal to $t - (t - \tau) = \tau$. We restrict the sample period to 2003-2010 in order to have enough observations for the first period.

The results presented in Table 8 show that a potential learning process cannot explain the absence of persistent differences in forecast performance. When restricting the sample to analysts with at least two years of experience, we still do not observe, in the test period, any significant differences between the first and last quintiles.

5.4 Robustness check: slow adjustment of target prices

Dechow and You (2013) assume that financial analysts might be slow at adjusting their target prices. One of the reasons for this slow adjustment is that target prices are usually embedded in analysts' reports. Because writing a report is a long and difficult task, financial analysts may not adjust their target prices as often as they should. For example, if the analyst changes her opinion about the future price of a stock one month after her initial forecast, she might have to wait for the next report publication to revise her target price. This feature might cause the analysts to appear less skilled than they actually are.

In order to test this hypothesis, we restrict the validity of the target prices to a shorter period of time (*e.g.* one month). That is, for a given stock and a given target price, we compute the target price forecast quality only for the first month following the issue date. In other words, we consider the forecasts to be inactive after one month. We then compute the analysts' forecast performance AFP using these short-validity target prices.

We conduct the same analysis as before to test for the existence of differential abilities. The (unreported) results - using 1 month, 3 months and 6 months for the validity of the target prices - still indicate that analysts do not exhibit differential forecasting abilities.

5.5 Scoring

An alternative way to check whether some analysts consistently outperform (or underperform) their colleagues is to compute the average relative performance over the sample period. The methodology is the following. Each month, we compute the analysts' forecast performance AFP. We then rank the analysts in five quintiles. We define the score of the analyst as (quintile-1)/4. Thus, the best analysts receive a score of 1 while the worst analysts receive a score of 0. The analysts ranked in the intermediate quintiles obtain a score of 0.25, 0.5 or 0.75. As the 2001-2010 period contains 120 months, we have, for each analyst, up to 120 monthly scores. In order to be able to compare the different scores, we remove the analysts with less than 24 monthly scores (less than 2 years of activity). The total number of analysts in the restricted sample is 5,481. Finally, we compute, for each analyst, the average of her monthly scores. An analyst who always ranks with the 20% best analysts will then have an average score equal to 1. An analyst who oscillates between the first quintile and the second quintile will have an average score between 0.75 and 1, and so on.

Figure 3 shows the average score of all the financial analysts for whom we were able to compute at least 24 monthly scores. We observe that only one analyst has a score higher than 0.75 and none have a score below 0.25. The average score is concentrated around the value 0.5. The solid curve corresponds to a Gaussian distribution. We observe that the distribution of the average score is a close fit to the Gaussian distribution. This additional result confirms our previous findings that financial analysts do not exhibit differential abilities to forecast future stock prices.

6 Conclusion

This article sets a new framework for evaluating the forecast quality of target prices issued by financial analysts. We show that measuring forecast errors (target price accuracy) is not sufficient to evaluate the quality of target prices. Differences in volatility lead to different degrees of difficulty (predictability) when forecasting future stock prices. These differences of difficulty must be incorporated when evaluating the forecast performance of financial analysts.

Our measure of target price forecast quality is defined as the difference between the forecast difficulty (the expected value of the absolute forecast error estimated at the issue date) and the target price accuracy (the realized absolute forecast error). We propose a way to estimate the expected value of the absolute forecast error. We are then able to evaluate the forecast difficulty of each target price.

The contributions of our measure are the following. First, our measure accounts for differences in volatility. Second, it allows us to consider forecasts with different horizons. Third, our measure is a dynamic measure, meaning that we are able to measure the forecast quality of a target price at any point in time. As a consequence of this last feature, our measure also provides a simple and consistent solution to the issue of measuring the forecast quality of a target price when a revision occurs. In the empirical part of our study, we show that, when taking into account the issue of the predictability, financial analysts do not exhibit persistent differential abilities. Our measure of forecast quality is designed to evaluate target prices. However, this measure could also be used to evaluate the forecast quality of other types of forecasts (exchange rates, for instance).

References

Asquith, Paul, Michael B. Mikhail, Andrea S. Au. 2005. Information content of equity analyst reports. *Journal of Financial Economics* **75**(2) 245–282.

Beckers, Stan, Miachel Steliaros, Alexander Thomson. 2004. Bias in european analysts earnings forecasts. *Financial Analysts Journal* **60**(2) 74–85.

Bilinski, Pawel, Danielle Lyssimachou, Martin Walker. 2013. Target price accuracy: International evidence. *The Accounting Review* **88**(3) 825–851.

Black, Fischer, Myron S. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* **81**(3) 637–654.

Bonini, Stefano, Laura Zanetti, Roberto Bianchini, Antonio Salvi. 2010. Target price accuracy in equity research. *Journal of Business Finance and Accounting* **37**(9-10) 1177–1217.

Bradshaw, Mark T., Lawrence D. Brown, Kelly Huang. forthcoming. Do sell-side analysts exhibit differential target price forecasting ability? *Review of Accounting Studies* .

Brav, Alon, Reuven Lehavy. 2003. An empirical analysis of analysts' target prices: Short-term informativeness and long-term dynamics. *Journal of Finance* **58**(5) 1933–1967.

Clement, Michael B. 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* **27**(3) 285–303.

DeBondt, Werner F.M., William P. Forbes. 1999. Herding in analyst earnings forecasts: Evidence from the united kingdom. *European Financial Management* **5**(2) 143–163.

Dechow, Patricia M., Haifeng You. 2013. Determinants of target price forecasts and valuation errors made by analyst. Working paper, Haas School of Business, University of California, Berkeley.

Hong, Harrison, Jeffrey D. Kubik, Amit Solomon. 2000. Security analysts' career concerns and herding of earnings forecasts. *RAND Journal of Economics* **31**(1) 121–144.

Huberts, Lex C., Russell J. Fuller. 1995. Predictability bias in the u.s. equity market. *Financial Analysts Journal* **51**(2) 12–28.

Ivkovic, Zoran, Narasimhan Jegadeesh. 2004. The timing and value of forecast and recommendation revisions. *Journal of Financial Economics* **73**(3) 433–463.

Jacob, John, Thomas Z. Lys, Margaret A. Neale. 1999. Expertise in forecasting performance of security analysts. *Journal of Accounting and Economics* **28**(1) 51–82.

Kahneman, Daniel. 2011. *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Kerl, Alexander G. 2011. Target price accuracy. *BuR - Business Research* **4**(1) 74–96.

Lawrence, Alastair, James Ryan, Yuan Sun. 2012. Investors' demand for sell-side research: Sec filings, media coverage, and market factors. Working paper, Haas School of Business, University of California, Berkeley.

Appendix A

Assuming, as it is common in the financial literature, that stock prices can be modeled as a Geometric Brownian Motion, we have

$$\log(S_T) \sim \mathcal{N}\left(\log(S_t) + \left(\mu - \frac{1}{2}\sigma^2\right)(T-t)\,, \sigma^2\,(T-t)\right), \tag{7}$$

where $\mathcal{N}()$ is the normal distribution, μ is the drift and σ is the volatility.

The probability that the stock price ends up inside an interval $[b_l; b_u]$ at the end of a determined horizon is equal to

$$\begin{aligned}
\Pr\left[\log(b_l) < \log(S_T) < \log(b_u)\right] &= \Pr\left[\tfrac{\log(b_l/S_t)-(\mu-\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t)}} < z < \tfrac{\log(b_u/S_t)-(\mu-\frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t)}}\right] \\
&= \Pr\left[b_l^* < z < b_u^*\right] \\
&= \Phi\left(b_u^*\right) - \Phi\left(b_l^*\right),
\end{aligned}$$

where b_l^* and b_u^* are defined by

$$b_l^* = \frac{\log(b_l/S_t) - (\mu - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t)}} \text{ and } b_u^* = \frac{\log(b_u/S_t) - (\mu - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t)}}, \tag{8}$$

where $\Phi()$ is the cumulative distribution function of a standard Gaussian random variable and z is a standard Gaussian variable.

The probability of the stock price ending up inside a given interval at the end of a determined horizon is a nonlinear function of both the volatility and the horizon. It follows, by extension, that the expected value of the absolute forecast error is a nonlinear function of both the stock return volatility and the target price horizon.

Appendix B

Assuming that the stock price follows a geometric Brownian motion, the *ex-post* forecast quality $TPFQ_{t,T}$ of a target price issued at time t can be calculated according to the Black and Scholes (1973) model as

$$
\begin{aligned}
TPFQ_{t,T} &= e^{r(T-t)}\left[\Phi(d_{1,t}) - \Phi(-d_{1,t}) - TP_{t,T}e^{-r(T-t)}\left(\Phi(d_{2,t}) - \Phi(-d_{2,t})\right)\right] - |S_T - TP_{t,T}| \\
&= e^{r(T-t)}\left[2\Phi(d_{1,t}) - 1\right] - TP_{t,T}\left[2\Phi(d_{2,t}) - 1\right] - |S_T - TP_{t,T}|,
\end{aligned}
\tag{9}
$$

with

$$
\begin{aligned}
d_{1,t} &= \frac{\ln\left(\frac{1}{TP_{t,T}}\right) + \left(r + \frac{1}{2}\sigma_t^2\right)(T-t)}{\sigma_t\sqrt{T-t}} \\
d_{2,t} &= d_{1,t} - \sigma_t\sqrt{(T-t)},
\end{aligned}
\tag{10}
$$

where $\Phi()$ is the cumulative distribution function of a standard Gaussian random variable, t is the time at which the forecast is issued, and σ_t is the stock return volatility estimated at time t. Remember the assumption $S_t = 1$ which explains the way $d_{1,t}$ is written.

Our approach implies that we do not distinguish between under- and over-achievement. If we consider two forecasts $TP^1_{t,T} = S_t - \Delta$ and $TP^2_{t,T} = S_t + \Delta$, we should obtain the same forecast quality if, at the end of the horizon, we have $\left|S_T - TP^1_{t,T}\right| = \left|S_T - TP^2_{t,T}\right|$. However, because $\ln(\frac{S_t}{S_t+\Delta}) \neq -\ln(\frac{S_t}{S_t-\Delta})$, this is not the case. In order to solve this issue, we apply a simple transformation (see Appendix C).

Appendix C

Consider two forecasts $TP^1_{t,T} = S_t - \Delta$ and $TP^2_{t,T} = S_t + \Delta$. As we do not distinguish between under- and over-achievement, we should have $\left|S_T - TP^1_{t,T}\right| = \left|S_T - TP^2_{t,T}\right| \Longrightarrow TPFQ^1_{t,T} = TPFQ^2_{t,T}$. However, because $\ln(\frac{S_t}{S_t+\Delta}) \neq -\ln(\frac{S_t}{S_t-\Delta})$, we have

$$
\left[C_t\left(TP^1_{t,T}\right) + P_t\left(TP^1_{t,T}\right)\right]e^{r(T-t)} < \left[C_t\left(TP^2_{t,T}\right) + P_t\left(TP^2_{t,T}\right)\right]e^{r(T-t)}. \tag{11}
$$

It follows that $TPFQ^1_{t,T} < TPFQ^2_{t,T}$. Even though the absolute deviation of the target price from the stock price S_t is the same for both target prices $TP^1_{t,T}$ and $TP^2_{t,T}$, and the absolute forecast errors $\left|S_T - TP^1_{t,T}\right|$ and $\left|S_T - TP^2_{t,T}\right|$ are the same at the end of the horizon, we do not obtain the same quality for the two forecasts. We apply a simple transformation to correct this.

As we do not distinguish between under- and over-achievement, the forecast quality of the target price depends only on the size of the deviation of the stock price from the target price. We therefore adopt the following convention. When a target price is below the stock price, we consider the symmetric of the price with respect to the target price. That is, we set the target price equal to 1 and consider the stock price to be equal to $1 + |S_t - TP_{t,T}|$. However, when there is a positive drift $\mu = r > 0$, the probability of reaching a target price of $TP_{t,T} = S_t - \Delta'$ is lower than the one of reaching a target price of $TP_{t,T} = S_t + \Delta'$. Therefore, we need to consider the symmetric of the price with respect to the discounted target price. The consequence of defining the stock price as a function of the discounted target price is that the risk-free rate in the Black-Scholes model is equal to 0.

Definition 5 *We consider the function f which measures the discounted deviation of the stock price from the target price. We write*

$$f(S_{t+\tau}, TP_{t,T}, r) = 1 + \left| S_{t+\tau} - TP_{t,T} e^{-r(T-t-\tau)} \right| . \tag{12}$$

The forecast quality of a target price issued at time t with horizon $T - t$ writes

$$\begin{aligned} TPFQ_{t,T} &= (C_t + P_t)\, e^{r(T-t)} - |f(S_T, TP_{t,T}, r) - 1| \\ &= e^{r(T-t)} \left(f(S_t, TP_{t,T}, r) \left[\Phi\left(d_{1,t}\right) - \Phi\left(-d_{1,t}\right) \right] - \left[\Phi\left(d_{2,t}\right) - \Phi\left(-d_{2,t}\right) \right] \right) \\ &- |f(S_T, TP_{t,T}, r) - 1| , \end{aligned} \tag{13}$$

with

$$\begin{aligned} d_{1,t} &= \frac{\ln\left(f(S_t, TP_{t,T}, r)\right) + \left(\frac{1}{2}\sigma_t^2\right)(T-t)}{\sigma_t\sqrt{T-t}} \\ d_{2,t} &= d_{1,t} - \sigma_t\sqrt{(T-t)}, \end{aligned}$$

where $\Phi()$ is the Gaussian cumulative distribution function. t is the time at which the forecast was issued, C_t is the value of the call option at time t, P_t is the value of the put option at time t, σ_t is the stock return volatility estimated at time t, r is the risk-free rate and $T - t$ is the horizon of the target price.

Appendix D

Proposition 6 *For a given final stock price S_T and a given target price $TP_{t,T}$, the forecast quality $TPFQ_{t,T}$ is an increasing function of the stock return volatility σ_t and of the length of the horizon $T-t$.*

Proof. Proof For a given final stock price S_T and a given target price $TP_{t,T}$, the sensitivity of the forecast quality $TPFQ_{t,T}$ to the volatility σ_t writes

$$\begin{aligned} \frac{\partial TPFQ_{t,T}}{\partial \sigma_t} &= \left[\frac{\partial C_t}{\partial \sigma_t} + \frac{\partial P_t}{\partial \sigma_t}\right] e^{r(T-t)} \\ &= 2e^{r(T-t)} S_t \sqrt{(T-t)} \Phi'(d_{1,t}) > 0, \end{aligned} \tag{14}$$

with $\Phi'(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$. For a given final stock price S_T and a given target price $TP_{t,T}$, the sensitivity of the forecast quality $TPFQ_{t,T}$ to the horizon $T-t$ writes

$$\frac{\partial TPFQ_{t,T}}{\partial (T-t)} = \left[\frac{\partial C_t}{\partial (T-t)} + \frac{\partial P_t}{\partial (T-t)}\right] e^{r(T-t)} + re^{r(T-t)} (C_t + P_t). \tag{15}$$

The sensitivity of a straddle to the maturity $T-t$ writes

$$\begin{aligned} \frac{\partial C_t}{\partial (T-t)} + \frac{\partial P_t}{\partial (T-t)} &= S_t \Phi'(d_{1,t}) \frac{\partial d_{1,t}}{\partial (T-t)} \\ &\quad - e^{-r(T-t)} TP_{t,T} \Phi(d_{2,t}) \frac{\partial d_{2,t}}{\partial (T-t)} \\ &\quad + S_t \Phi'(-d_{1,t}) \frac{\partial d_{1,t}}{\partial (T-t)} \\ &\quad + e^{-r(T-t)} TP_{t,T} \Phi'(d_{2,t}) \frac{\partial d_{2,t}}{\partial (T-t)} \\ &= \frac{\sigma}{\sqrt{T-t}} S_t \Phi'(d_{1,t}) \\ &\quad + r\Phi_{t,T} e^{-r(T-t)} \left[\Phi(d_{2,t}) - \Phi(-d_{2,t})\right], \end{aligned} \tag{16}$$

with $\Phi'(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$.

The sensitivity of a call option to the maturity $T-t$ is always positive. The sensitivity of a put option to the maturity $T-t$ is also positive except when the option is deep in the money. The transformation we apply (see Appendix C) implies that the put option is never in the money. Thus, the sensitivity of the straddle to the horizon $T - t$ is always positive. We then have

$$\frac{\partial TPFQ_{t,T}}{\partial (T-t)} = \left[\frac{\partial C_t}{\partial (T-t)} + \frac{\partial P_t}{\partial (T-t)}\right] e^{r(T-t)} + re^{r(T-t)} (C_t + P_t) > 0. \tag{17}$$

Figure 1
An illustration of the properties of target price forecast quality

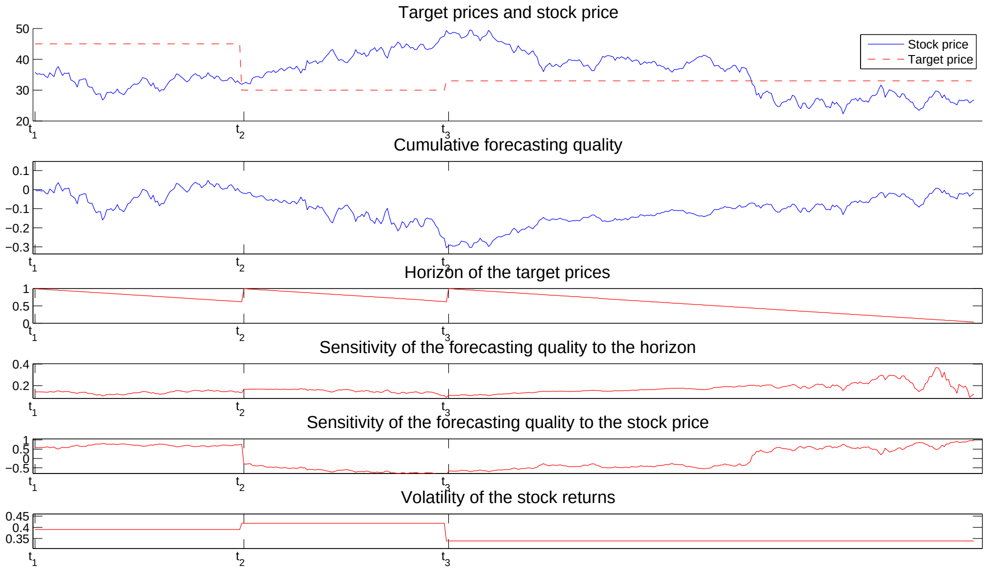

Figure 2
The influence of stock return volatility on target price forecast quality

Figure 3

Average score of analysts' forecast performance (AFP) over the 2001-2010 period

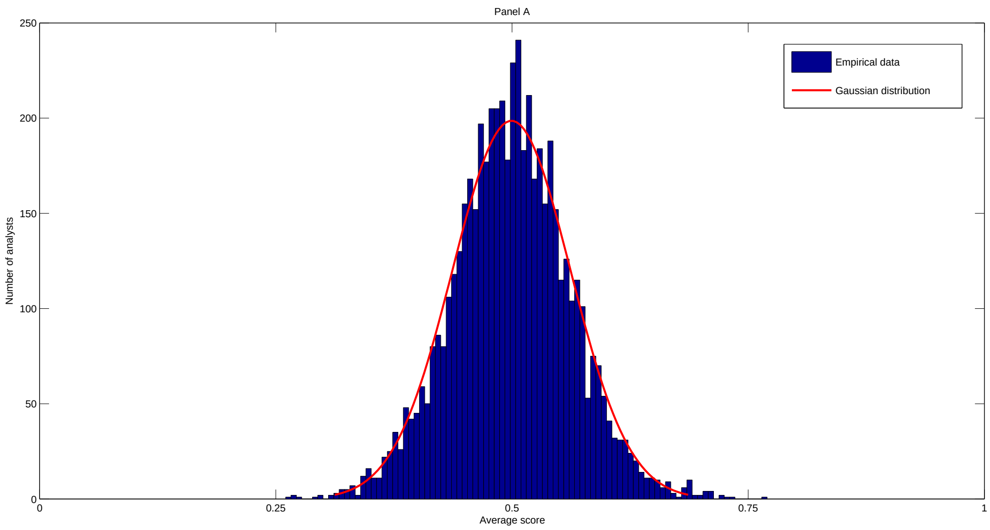

Notes. This figure presents the histogram of the average score of 5,481 analysts over the 2001-2010 period. We include an analyst in the sample only if we can compute at least 24 monthly observations for the analysts' forecast performance (AFP). Each month, the analysts are ranked in five quintiles with respect to their performance AFP. The monthly score is then computed as (quintile-1)/4. The average score corresponds to the mean of the monthly scores over the sample period.

Table 1
Descriptive statistics

	Number of forecasts	Number of active analysts	Number of analysts covering a stock			Number of stocks covered per analyst		
			Mean	Median	Max	Mean	Median	Max
2000	36,825	-	-	-	-	-	-	-
2001	44,178	4,466	6.30	4	50	6.38	4	213
2002	48,756	4,611	7.10	5	50	6.31	4	183
2003	51,263	4,342	7.22	5	55	6.61	4	102
2004	54,863	3,773	6.54	4	44	7.16	5	64
2005	56,291	3,731	6.67	5	51	7.72	6	76
2006	59,952	3,800	6.82	5	45	7.97	6	87
2007	65,377	3,768	7.00	5	46	8.42	7	100
2008	77,281	3,829	7.53	6	46	8.59	7	92
2009	75,275	3,833	8.00	6	50	8.48	7	81
2010	79,410	3,908	8.69	6	59	8.47	7	79

Notes. The sample consists in a total of 649,471 target prices made by 9,367 analysts (583 brokers) on 7,268 U.S. stocks for the 2000-2010 period. The first column indicates the number of target prices issued each year. The second column shows the number of active analysts. The three following columns report the average, median and maximum number of active analysts per stock. The remaining columns indicate the average, median and maximum number of stocks covered per analyst. The statistics for 2000 are not reported as the target prices issued in 1999 and still active in 2000 cannot be observed.

Table 2
Relationship between volatility and absolute forecast errors (AFE)

Panel A: Absolute forecast errors computed using actual target prices											
Volatility quintiles	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010
1 (Low volatility)	0.2813	0.2522	0.2431	0.1994	0.1759	0.1656	0.1698	0.2852	0.4712	0.2265	0.1796
2	0.3821	0.3549	0.3154	0.2558	0.2331	0.2153	0.2197	0.3824	0.5458	0.3091	0.2385
3	0.5274	0.5210	0.4161	0.3193	0.2903	0.2661	0.2865	0.4692	0.5538	0.3844	0.3015
4	0.7947	0.7711	0.5742	0.3953	0.3698	0.3351	0.3654	0.5624	0.6064	0.4952	0.3760
5 (High volatility)	1.2803	1.1467	0.8577	0.6502	0.5284	0.5075	0.4958	0.7327	0.7671	0.7782	0.5092
Diff (5-1)	0.9989	0.8946	0.6146	0.4508	0.3525	0.3419	0.3260	0.4475	0.2959	0.5517	0.3296
Mean t-test	104.0171***	109.6470***	71.6153***	62.3884***	76.3825***	67.8006***	61.2385***	94.0024***	45.3043***	68.8016***	85.0555***

Panel B: Absolute forecast errors computed using naive forecasts											
Volatility quintiles	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010
1 (Low volatility)	0.2534	0.1908	0.1903	0.2213	0.1648	0.1508	0.1590	0.2450	0.3328	0.2365	0.1927
2	0.3210	0.2499	0.2484	0.2898	0.2370	0.2039	0.2058	0.3239	0.4038	0.3201	0.2408
3	0.3963	0.3513	0.3211	0.3545	0.2885	0.2493	0.2600	0.3927	0.4216	0.4052	0.2859
4	0.5342	0.4659	0.4468	0.4200	0.3333	0.3104	0.3210	0.4554	0.4975	0.5374	0.3526
5 (High volatility)	0.6954	0.6052	0.7336	0.6852	0.3729	0.4319	0.4017	0.5264	0.6874	0.8857	0.4160
Diff (5-1)	0.4420	0.4144	0.5433	0.4639	0.2081	0.2811	0.2427	0.2814	0.3546	0.6492	0.2233
Mean t-test	83.8646***	99.4626***	59.7663***	57.8992***	54.8949***	57.6214***	49.7714***	84.4185***	54.5609***	72.9812***	64.2128***

Notes. Each year, the target prices in the sample are assigned to five quintiles with respect to the 6-month historical volatility of the underlying stock estimated at the issue date. The table reports, for each year and each quintile, the average target price accuracy (average absolute forecast errors). Panel A reports the average absolute forecast errors using the actual target prices from our sample. Panel B reports the average absolute forecast errors using naive forecasts (each target price is replaced by a forecast that is equal to the price at the issued date capitalized by the risk-free rate). The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.

Table 3
Nonlinear relationship between volatility and absolute forecast errors (AFE)

	Absolute forecast errors			
k-th decile= $\mathbf{1}_{jt}^{k}\sigma_{jt}$	Coefficient (β_k)	Standard error	t-statistic	p-value
1st decile (Low volatility)	0.2311***	0.060756	3.80	0.000
2nd decile	0.4129***	0.045587	9.06	0.000
3rd decile	0.4824***	0.038026	12.69	0.000
4th decile	0.5361***	0.032872	16.31	0.000
5th decile	0.5850***	0.028770	20.33	0.000
6th decile	0.5979***	0.025255	23.67	0.000
7th decile	0.5990***	0.022002	27.22	0.000
8th decile	0.6059***	0.018644	32.50	0.000
9th decile	0.6136***	0.015021	40.85	0.000
10th decile (High volatility)	0.6395***	0.010952	58.39	0.000
Number of observations		604,677		
R-squared		0.1893		

Notes. This table shows the coefficient estimates (Coefficients) from the following OLS regression

$$AFE_{jt} = \alpha + \sum_{k=1}^{10} \beta_k \mathbf{1}_{jt}^{k} \sigma_{jt} + \epsilon_{jt}, \tag{18}$$

where AFE_{jt} is the absolute forecast error of a target price on firm j issued by any analyst at time t, σ_{jt} is the stock return volatility of stock j measured at time t and $\mathbf{1}_{jt}^{k}$ is an indicator that takes the value 1 if the stock return volatility σ_{jt} belongs to the k-th volatility decile and 0 otherwise. We denote the variable $\mathbf{1}_{jt}^{k}\sigma_{jt}$ as the k-th decile. ***/**/* correspond to 1%/5%/10% significance levels. P-values are computed using robust standard errors.

Table 4

Persistent differences in absolute forecast errors (AFE)

	Panel A: Quarterly period				
		Measurement period $]t-1;t]$		Test period $]t;t+\theta+1]$	
Performance quintile (measurement period)	Number of observations	Analysts' performance (average AFE)	Volatility of stocks covered	Analysts' performance (average AFE)	Volatility of stocks covered
1 (Best)	11,083	0.1551	0.3870	0.3322	0.3720
2	11,079	0.2847	0.4364	0.3654	0.4161
3	11,085	0.4026	0.4919	0.4108	0.4673
4	11,079	0.5627	0.5720	0.4838	0.5284
5 (Worst)	11,083	0.9693	0.6850	0.5779	0.5973
Diff (5-1)		0.8143	0.2980	0.2456	0.2253
Mean t-test		17.7540***	11.0277***	12.5456***	12.6069***

	Panel B: Semiannual period				
		Measurement period $]t-1;t]$		Test period $]t;t+\theta+1]$	
Performance quintile (measurement period)	Number of observations	Analysts' performance (average AFE)	Volatility of stocks covered	Analysts' performance (average AFE)	Volatility of stocks covered
1 (Best)	6,258	0.1713	0.3747	0.3266	0.3653
2	6,260	0.2992	0.4291	0.3672	0.4102
3	6,256	0.4128	0.4924	0.4116	0.4648
4	6,260	0.5662	0.5729	0.4847	0.5217
5 (Worst)	6,258	0.9559	0.6876	0.5758	0.5897
Diff (5-1)		0.7847	0.3128	0.2491	0.2244
Mean t-test		12.6568***	8.1821***	10.5185***	10.6874***

Notes. This table reports the analysts' performance in the test period $]t;t+\theta+1]$ conditional on their performance in the measurement period $]t-1;t]$. θ is a 12-month lag which ensures independence in prices (the measurement and the test periods are not overlapping). The analyst's performance is measured, for a given period, as the average of the absolute forecast errors (AFE) on all the target prices she issued during that period. The measurement periods are quarterly (Panel A) and semiannual (Panel B). We rank analysts in 5 quintiles based on their performance in the measurement period and we obtain the corresponding performance in the test period. We also report the volatility of the stocks covered, computed as the average of the 6-month historical stock return volatilities on all the target prices issued by the analyst during the period under consideration. Conditional on the ranking made during the measurement period $]t-1;t]$, we report, for the test period $]t;t+\theta+1]$, both the analysts' performance and the volatility of the stocks covered. The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.

Table 5
Persistent differences in absolute forecast errors (AFE) using naive forecasts

Panel A: Quarterly period

Performance quintile (measurement period)	Number of observations	Measurement period $]t-1;t]$ Analysts' performance (average AFE)	Measurement period $]t-1;t]$ Volatility of stocks covered	Test period $]t;t+\theta+1]$ Analysts' performance (average AFE)	Test period $]t;t+\theta+1]$ Volatility of stocks covered
1 (Best)	11,083	0.1317	0.4044	0.3057	0.3845
2	11,079	0.2422	0.4445	0.3254	0.4199
3	11,085	0.3357	0.4988	0.3564	0.4678
4	11,079	0.4550	0.5673	0.3975	0.5219
5 (Worst)	11,083	0.7614	0.6573	0.4481	0.5870
Diff (5-1)		0.6297	0.2529	0.1424	0.2025
Mean t-test		16.3544***	9.9624***	6.9959***	11.8540***

Panel B: Semiannual period

Performance quintile (measurement period)	Number of observations	Measurement period $]t-1;t]$ Analysts' performance (average AFE)	Measurement period $]t-1;t]$ Volatility of stocks covered	Test period $]t;t+\theta+1]$ Analysts' performance (average AFE)	Test period $]t;t+\theta+1]$ Volatility of stocks covered
1 (Best)	6,258	0.1465	0.3908	0.3033	0.3756
2	6,260	0.2554	0.4386	0.3319	0.4180
3	6,256	0.3434	0.4944	0.3603	0.4609
4	6,260	0.4567	0.5720	0.3964	0.5192
5 (Worst)	6,258	0.7507	0.6609	0.4475	0.5779
Diff (5-1)		0.6042	0.2701	0.1442	0.2022
Mean t-test		12.2240***	7.3593***	6.2430***	9.4990***

Notes. This table presents the analysts' performance in the test period $]t;t+\theta+1]$ conditional on their performance in the measurement period $]t-1;t]$. θ is a 12-month lag which insures independence in prices (the measurement and the test periods are not overlapping). The analyst's performance is measured, for a given period, as the average of the absolute forecast errors (AFE) on all the target prices she issued during that period. The measurement periods are quarterly (Panel A) and semiannual (Panel B). This analysis uses naive forecasts instead of the actual target prices. That is, analysts systematically issue target prices with an implied return equal to the risk-free interest rate. We rank analysts in 5 quintiles based on their performance in the measurement period and we obtain the corresponding performance in the test period. We also report the volatility of the stocks covered, computed as the average of the 6-month historical stock return volatilities on all the target prices issued by the analyst during the period under consideration. Conditional on the ranking made during the measurement period $]t-1;t]$, we report, for the test period $]t;t+\theta+1]$, both the analysts' performance and the volatility of the stocks covered. The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.

Table 6
Test of forecasting abilities using the *ex-post* measure of target price forecast quality

Panel A: Quarterly period

		Measurement period $]t-1;t]$		Test period $]t;t+\theta+1]$	
Performance quintile (measurement period)	Number of observations	Analysts' *ex-post* forecast performance ($exAFP$)	Volatility of stocks covered	Analysts' *ex-post* forecast performance ($exAFP$)	Volatility of stocks covered
1 (Best)	11,082	0.3536	0.6222	0.0716	0.4967
2	11,079	0.1700	0.5000	0.0637	0.4470
3	11,082	0.0819	0.4588	0.0585	0.4385
4	11,079	−0.0122	0.4647	0.0578	0.4634
5 (Worst)	11,082	−0.2487	0.5148	0.0651	0.5231
Diff (5-1)		−0.6023	−0.1075	−0.0065	0.0264
Mean t-test		−20.2776***	−6.3892***	−0.9256	1.7703*

Panel B: Semiannual period

		Measurement period $]t-1;t]$		Test period $]t;t+\theta+1]$	
Performance quintile (measurement period)	Number of observations	Analysts' *ex-post* forecast performance ($exAFP$)	Volatility of stocks covered	Analysts' *ex-post* forecast performance ($exAFP$)	Volatility of stocks covered
1 (Best)	6,257	0.3268	0.6095	0.0651	0.4905
2	6,260	0.1560	0.4961	0.0614	0.4402
3	6,254	0.0762	0.4613	0.0551	0.4358
4	6,260	−0.0086	0.4660	0.0577	0.4585
5 (Worst)	6,257	−0.2292	0.5131	0.0633	0.5149
Diff (5-1)		−0.5561	−0.0963	−0.0018	0.0244
Mean t-test		−14.5024***	−4.1288***	−0.2123	1.2637

Notes. This table presents the analysts' *ex-post* forecast performance $exAFP$ in the test period $]t;t+\theta+1]$ conditional on their forecast performance in the measurement period $]t-1;t]$. θ is a 12-month lag which ensures independence in prices (the measurement and the test periods are not overlapping). The analyst's *ex-post* forecast performance $exAFP$ is measured, for a given period, as the average of the *ex-post* target price forecast quality on all the target prices she issued during that period. The *ex-post* target price forecast quality is measured as the expected value of the absolute forecast error estimated at the time the target price is issued minus the realized absolute forecast error measured at the end of the 12-month horizon. In this set-up, we consider a revision as a new and independent forecast. The measurement periods are quarterly (Panel A) and semiannual (Panel B). We rank analysts in 5 quintiles based on their *ex-post* forecast performance $exAFP$ in the measurement period and we obtain the corresponding forecast performance in the test period. We also report the volatility of the stocks covered, computed as the average of the 6-month historical stock return volatilities on all the target prices issued by the analyst during the period under consideration. Conditional on the ranking made during the measurement period $]t-1;t]$, we report, for the test period $]t;t+\theta+1]$, both the analysts' *ex-post* forecast performance $exAFP$ and the volatility of the stocks covered. The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.

Table 7
Test of forecasting abilities in a dynamic setting

		Panel A: Quarterly period	
		Measurement period $]t-1;t]$	Test period $]t;t+\theta+1]$
Performance quintile (measurement period)	Number of observations	Analysts' forecast performance (AFP)	Analysts' forecast performance (AFP)
1 (Best)	21,684	0.1312	0.0207
2	21,687	0.0551	0.0200
3	21,680	0.0211	0.0188
4	21,687	−0.0135	0.0196
5 (Worst)	21,684	−0.1043	0.0149
Diff (5-1)		−0.2355	−0.0059
Mean t-test		−17.0704***	−0.5806
		Panel B: Semiannual period	
		Measurement period $]t-1;t]$	Test period $]t;t+\theta+1]$
Performance quintile (measurement period)	Number of observations	Analysts' forecast performance (AFP)	Analysts' forecast performance (AFP)
1 (Best)	9,251	0.1904	0.0414
2	9,249	0.0853	0.0387
3	9,250	0.0395	0.0334
4	9,249	−0.0071	0.0306
5 (Worst)	9,251	−0.1326	0.0223
Diff (5-1)		−0.3230	−0.0191
Mean t-test		−12.6962***	−1.3708

Notes. This table reports the analysts' forecast performance AFP in the test period $]t;t+1]$ conditional on their forecast performance in the measurement period $]t-1;t]$. The analyst's forecast performance (AFP), for a given period, is defined as the sum of the target price forecast quality $TPFQ$ on all her outstanding target prices during that period. The measurement periods are quarterly (Panel A) and semiannual (Panel B). We rank the analysts in 5 quintiles based on their forecast performance in the measurement period and we obtain the corresponding forecast performance in the test period. In order to be included in the analysis, an analyst must have an active target price for at least 80 percent of the days in the sample period. Conditional on the ranking made during the measurement period $]t-1;t]$, we report the analyst's forecast performance AFP for the test period $]t;t+1]$. The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.

Table 8
Impact of learning and forecasting abilities

	Panel A: Quarterly period		
		Measurement period $]t-1;t]$	Test period $]t;t+\theta+1]$
Performance quintile (measurement period)	Number of observations	Analysts' forecast performance (AFP)	Analysts' forecast performance (AFP)
1 (Best)	4,246	0.1136	0.0155
2	4,241	0.0451	0.0167
3	4,245	0.0168	0.0134
4	4,241	−0.0124	0.0128
5 (Worst)	4,246	−0.0953	0.0147
Diff (5-1)		−0.2090	−0.0008
Mean t-test		−14.9959***	−0.3498
	Panel B: Semiannual period		
		Measurement period $]t-1;t]$	Test period $]t;t+\theta+1]$
Performance quintile (measurement period)	Number of observations	Analysts' forecast performance (AFP)	Analysts' forecast performance (AFP)
1 (Best)	1,761	0.1602	0.0293
2	1,763	0.0678	0.0267
3	1,764	0.0246	0.0249
4	1,763	−0.0173	0.0299
5 (Worst)	1,761	−0.1426	0.0267
Diff (5-1)		−0.3028	−0.0026
Mean t-test		−10.4057***	−0.5724

Notes. This table reports the analysts' forecast performance AFP in the test period $]t;t+1]$ conditional their forecast performance in the measurement period $]t-1;t]$. The analyst's forecast performance (AFP), for a given period, is defined as the sum of the target price forecast quality $TPFQ$ on all her outstanding target prices during that period. We restrict the sample to the analysts who have at least two years of experience. The sample period is 2003-2010. The measurement periods are quarterly (Panel A) and semiannual (Panel B). We rank the analysts in 5 quintiles based on their forecast performance in the measurement period and we obtain the corresponding forecast performance in the test period. In order to be included in the analysis, an analyst must have an active target price for at least 80 percent of the days in the sample period. Conditional on the ranking made during the measurement period $]t-1;t]$, we report the analyst's forecast performance AFP for the test period $]t;t+1]$. The statistical significance of the difference across top and bottom quintiles is computed using a t-test. ***/**/* represent significance at the 0.01, 0.05 and 0.1 level.