Managerial Recalibration: Do CFOs Learn?

Michael Boutros

Duke University, Durham, NC 27708 USA

John R. Graham

Duke University, Durham, NC 27708 USA
National Bureau of Economic Research, Cambridge, MA 02138 USA

Campbell R. Harvey*

Duke University, Durham, NC 27708 USA
National Bureau of Economic Research, Cambridge, MA 02138 USA

First Draft (Incomplete and Preliminary)
Current Version: July 10, 2018

Abstract

Using almost 14,000 forecasts of one-year S&P 500 returns made by Chief Financial Officers over a 12-year period, we track the individual executives that provide multiple forecasts to evaluate how they adapt and recalibrate their forecasts in response to return realizations. We focus on the confidence intervals that each CFO provides for their forecast. A simple model of Bayesian learning suggests that confidence intervals should recalibrate as a result of return realizations. We find that the confidence intervals significantly widen when CFOs miss their confidence ranges and narrow when realizations fall within their ex ante intervals. We also detail interesting asymmetries. For example, missing the confidence interval on the downside leads to a much larger adjustment of the downside component of the interval. The results are consistent with CFOs learning.

Keywords: Learning, Information, Behavioral Economics, Volatility Forecasts, Market Forecasts, Behavioral Finance, Bayesian Updating.

---

*Send correspondence to: Campbell R. Harvey (cam.harvey@duke.edu). We appreciate the comments of . . . and seminar participants at . . . . We are especially indebted to our colleague John Payne who suggested we pursue the learning research.

1. Introduction

It is now widely understood that decisions are impacted by the circumstances of the decision maker's environment. Children of the Great Depression are likely to act differently than people that did not experience such hardship (Malmendier and Nagel, 2011). Vissing-Jorgensen (2003) investigates the difference in behavior between wealthy and non-wealthy investors. Less is known about the way people may change beliefs in reaction to immediate experiences.

The goal of our paper is to test whether Chief Financial Officers (CFOs) learn through time. Ben-David et al. (2013), in a multi-year survey of executives, show that CFOs are, on average, miscalibrated: the confidence intervals for their forecasts of S&P 500 returns are far too narrow. We augment their data in two ways: first, we update the data and second, and more importantly, we focus on the individual forecaster. We have thousands of CFOs that make multiple forecasts. This allows us to test whether the CFOs learn from their mistakes and adjust forecasts if they do well.

Our database consists of 13,757 forecasts of the one-year return on the S&P 500 and 4,249 are identified as CFOs making multiple forecasts. There is a CFO in our dataset that has made 43 forecasts. The respondents to the survey also provide 80% confidence intervals for their forecasts. We focus on the interval given that it is very difficult to learn what the right S&P 500 forecast should be. Our empirical results focus on how the width of the confidence interval reacts to return realizations.

Our framework is a simple model of Bayesian learning that focuses on the confidence intervals. While Bayes' Rule is well known and tested in many laboratory settings, there is no practical way to engage senior financial executives in a multi-year laboratory experiment. Instead, we rely on the results of quarterly survey that has been ongoing for 20 years and administered electronically over the last 12 years.

Essentially, we track individual CFOs and determine if their confidence intervals react to whether they hit or miss their prespecified interval. Conveniently, all CFOs are trying to forecast the same item: next year's return on the S&P 500. The overall market is something that CFOs pay attention to given that investors routinely ask them about the performance of the company's stock price and it is unlikely they can explain their firm's stock performance without reference to overall market movements.

Our results are striking. There is strong evidence that CFOs "recalibrate" their forecast confidence intervals given the feedback of realized market returns. Intervals increase when CFO forecasts miss their ex ante interval and intervals decrease when CFOs hit their range. Interestingly, the intervals are not necessarily symmetric around the forecast. We find that CFOs after missing the interval from the downside, for example, tend to lower the downside of the interval - not the upside.

Given the evidence in Ben-David et al. (2013) of confidence intervals being too narrow, missing the confidence interval tends to reduce the degree of miscalibration. However, for those that hit the interval, the intervals become more miscalibrated as the (lucky) forecasts fuels additional overconfidence. Nevertheless, the recalibration after missing (learning) is much greater than the narrowing after hitting the interval.

We build a model of learning where forecasters use Bayes' rule to update their beliefs of the return process given new observations. The model predicts that when forecasters miss their confidence intervals, they should widen their next forecast's confidence intervals, and that when they hit their confidence intervals, they should narrow their next confidence interval. Further, the degree to which the forecaster narrows or widens their confidence interval depends on the degree to which they hit or missed.

Our paper is related to the broad literature on overprecision and, in particular, overconfidence, defined as "excessive faith that [the agent] knows the truth" (Moore and Schatz, 2017). Overconfidence is well documented in many settings, notably amongst CFOs in Ben-David et al. (2013) but also in many other smaller laboratory settings (Moore et al., 2016). This literature documents the presence of overconfidence but stops short of testing whether overconfident agents re-evaluate their confidence and adjust in the presence of making mistakes.

On the theoretical side, Gervais and Odean (2001) build a model of traders who learn about their abilities over time through experience. Empirically, the model is evaluated by looking at cross-sectional moments of market data, and not by examining the empirical beliefs of traders. Our paper is also broadly related to the reenforcement literature, summarized in Barber and Odean (2013) as stating that "the simplest form of learning may be to repeat behaviors that previously coincided with pleasure and avoid those that coincided with pain." Consistent with our results of asymmetric learning, Kuhnen (2014) finds that people in the "negative domain" of returns often form relatively more pessimistic beliefs about the investment options available to them. This negative bias is driven by agents responding more to outcomes in the negative versus position domain. We construct a model of Bayesian learning, tapping the vast literature built using the well-known Bayes' rule.

Finally, our paper is related to the growing literature on the beliefs held by firms and their financial officers. Bloom et al. (2017) use data from the Census Bureau's Management and Organizational Practices Survey on firms' reported subjective probability distributions of important future outcomes such as employment and input costs. They find that firms' subjective expectations are generally coherent probability distributions and similar to historical data, indicating that firms are able to successfully generate subjective distributions based on previous observations of data.

In addition to Ben-David et al. (2013), Greenwood and Shleifer (2014) and Gennaioli et al. (2016) use the same CFO survey as in this paper. Greenwood and Shleifer (2014) analyze a long time series of investor expectations and their relationship to expected returns in standard finance models. Gennaioli et al. (2016) look at a sample of firms and analyze the relationship between expectations of growth and investment. In contrast to the learning mechanism developed in this paper (albeit for learning about a different stochastic process), they find that firms often use simple extrapolation in forming their next-period beliefs of growth. Similarly, Kuchler and Zafar (2017) find that US households often extrapolate local home price changes when surveyed about beliefs of national home price changes over the next year.

Our paper is organized as follows. In the second section, we present a simple model of Bayesian learning that focuses on confidence intervals. The model makes predictions about how the posterior changes as a result of return realizations. The third section details the data that we use and provides some summary statistics. The fourth part provides our main empirical analysis. Some concluding remarks are offered in the final section.

2. A Model of Bayesian Learning

We present a simple Bayesian learning model that shows how confidence intervals react to forecast realizations. The model is simplified in two ways. First, we assume that the confidence intervals are symmetric when we know that 40% of the confidence intervals presented by the CFOs are not. We believe that allowing for asymmetry is unlikely to substantially change our results. Second, and more importantly, we initially focus on the unconditional distribution of stock returns. In this framework, CFOs learn about the variance of the returns process.

Examining the unconditional distribution has the disadvantage of implying the the volatility of the stock market is constant when it is well documented that volatility is time varying. As a result, a change in the confidence interval could be driven by the CFO's perception of changing volatility. For example, the observed return might fall within the CFO's confidence interval but the CFO might increase the interval next period because of a perception of heightened market volatility. Indeed, CFOs might have a good idea of the unconditional distribution of returns (based on both their experiences and over a century of data) but may need to learn about the evolution of the conditional distribution. Nevertheless, a starting point is the unconditional distribution. We discuss the implications of conditional volatility later.

Consider a discrete-time environment with a single forecaster representing a CFO. In each period, she must generate a forecast the one-year-ahead return and a corresponding confidence interval. We assume the forecaster believes the true return process, conditional on all observations, is normally distributed according to $R_{t+1} \sim N(\mu, \sigma)$. The forecaster does not know the exact values of μ or σ but has in mind a joint distribution of possible values for each parameter. In each period, the forecaster observes new data and learns about distribution of the return process. We focus on learning in the context of recalibrating the size of the confidence interval around the point forecast provided by each CFO.

Denote μ_s the belief of the mean formed using all observations $\{r_1, r_2, \ldots, r_s\}$. At the start of period t, before observing r_t, the forecaster's belief of the mean is μ_{t-1}. This is called the prior belief.[1] Once r_t is observed, the forecaster uses Bayes' rule to combine the new observation with the prior belief and form the posterior belief, μ_t. In each period, then the forecaster's belief of the return process is distributed according to $N(\mu_t, \sigma_t)$, and using this belief, she generates a point estimate of the future return. Similarly, she begins the period with a belief of the standard deviation, σ_{t-1}, and uses Bayes' rule and the observed return to form the posterior belief σ_t. Given the point forecast, the forecaster constructs a confidence interval using the posterior belief. In the next period, $t+1$, the prior beliefs are μ_t and σ_t, and the learning process continues after r_{t+1} is observed.

[1] At time $t = 1$, the forecaster cannot use Bayes' rule since there have been no realizations of the return process. Instead, she is endowed with some exogenous prior beliefs μ_0 and σ_0.

We assume that the forecaster's beliefs of the mean and variance, μ_t and σ_t, are jointly distributed according to the Gaussian-Inverse-Gamma (GIG) distribution

$$(\hat{\mu}_t, \hat{\sigma}_t^2) \sim N(\phi_t, \kappa_t^{-1}\sigma_t)(\Gamma(\alpha_t, \beta_t))^{-1}.$$

The joint distribution is described by four hyperparameters, ϕ, κ, α, and β, which in turn govern beliefs of the two parameters of interest, μ and σ. The hyperparameters α and κ exist primarily to adjust for the number of observations used in constructing beliefs, while ϕ and β are driven by the observations and map directly into the beliefs. Below and in the appendix, we describe the evolution of the hyperparameters, and, in turn, the evolution of the beliefs.

We choose to use this prior distribution for several reasons. First, this distribution is physically plausible; the mean can take on any finite value while the variance must be strictly positive. Second, when this prior is combined with an observed return using Bayes' rule, the posterior distribution is also in the GIG family of distributions. This property makes the GIG a *conjugate prior distribution*, and by assuming the prior is a conjugate prior, we ensure that all subsequent beliefs are also in the GIG family. Conjugate priors are widely used in the Bayesian literature because they provide analytical tractability, allowing us to derive our main results in closed form.

Given this distribution of beliefs, the forecaster uses the expected value of each marginal distribution as her belief of the mean and variance:

$$\mu_t = E[\hat{\mu}_t],$$
$$\sigma_t^2 = E[\hat{\sigma}_t^2].$$

The evolution of the belief of the mean is described in the appendix. Simply put, it tracks the sample mean, updating with each new observation.

The marginal prior distribution of the variance belief is an Inverse-Gamma distribution with two hyperparameters, α and β, governing the scale and shape. [2] The forecaster's prior belief of the unknown variance is summarized by the expected value of this distribution:

$$\sigma_t^2 = \frac{\beta_t}{\alpha_t}. \tag{1}$$

The hyperparameter β_t is the sum of squared errors in the sample:

$$\beta_t = \beta_{t-1} + \frac{\kappa_{t-1}}{\kappa_t}\frac{(r_{t-1} - \mu_{t-2})^2}{2} \tag{2}$$

Note that this is a recursive definition, initialized by some prior sum of squared errors, β_0. This expression combines the previous sum of squares, β_t, with the newly observed squared deviation weighted by $\frac{\kappa_{t-1}}{\kappa_t}$. Thus with each new observation, the sum of squares is updated to incorporate the new squared deviation, and in turn this feeds into the belief of the variance.

The hyperparameters α and κ serve to adjust for the number of observations used in constructing the belief.[3] At time t, the values of these hyperparameters are given by:

$$\kappa_t = \kappa_{t-1} + 1,$$
$$\alpha_t = \alpha_{t-1} + \frac{1}{2}.$$

Define $m_t = \frac{\kappa_t}{\kappa_{t+1}}$ as the weight on the new squared deviation in (2). Note that since $\kappa_{t+1} = \kappa_t + 1$, $m_t > 1$ for small t and $m_t \approx 1$ for large t. Intuitively, as more realizations of the return are observed, the relative weight of the new squared deviation decreases. In what follows, we assume that the forecaster

[2]The shape parameter essentially determines the mass surrounding the peak of the probability density while the scale parameter governs how "spread out" it is. As the value of the shape parameter increases, the peak of the PDF increases and more mass surrounds out. As the scale increases, the tails widen, drawing mass from the peak.

[3]See Appendix A for more details.

has observed sufficiently many returns such that $m_t \approx 1$.

Proposition 1 (Evolution of the Variance Belief)**.**

(a) *If the squared error of the previous forecast is greater than the prior belief of the variance, or* $(r_t - \mu_{t-1})^2 > \sigma_{t-1}^2$*, then the posterior belief of the variance is greater than the prior belief:*

$$\sigma_t^2 > \sigma_{t-1}^2.$$

(b) *If the squared error of the previous forecast is less than the prior belief of the variance, or* $(r_t - \mu_{t-1})^2 < \sigma_{t-1}^2$*, then the posterior belief of the variance is less than the prior belief:*

$$\sigma_t^2 < \sigma_{t-1}^2.$$

(c) *If the squared error of the previous forecast is equal to the prior belief of the variance, or* $(r_t - \mu_{t-1})^2 = \sigma_{t-1}^2$*, then the posterior belief of the variance is equal to the prior belief, or* $\sigma_t^2 = \sigma_{t-1}^2$.

The condition in part (a) can be rewritten as

$$r_t \notin [\mu_{t-1} - \sigma_{t-1}, \mu_{t-1} + \sigma_{t-1}].$$

This corresponds to a one-standard deviation interval around the mean belief. Proposition 1 states that if the observed return falls outside of this interval, the forecaster will increase her belief of the variance. If instead the observed return falls outside of this interval, the forecaster decreases her belief of the variance.

Further, given a confidence level, the width of the confidence interval depends only on the magnitude of the standard deviation. Therefore, there is a direct relationship between the belief of the variance and the width of the confidence interval, and as the belief of the variance changes, so too does the width of the forecaster's confidence interval.

Corollary 1. *If the condition in part (a) of proposition 1 is satisfied, then the forecaster's confidence interval widens.*

Corollary 1 states that the width of confidence interval for the next forecast is, through proposition 1, directly related to a one-standard deviation confidence interval around the previous forecast. Specifically, if the observed return does not fall inside a one-standard deviation confidence interval of the forecast, the confidence interval around the next forecast will be wider. This is true regardless of the new point forecast, which depends on the belief of the mean, since the width of the confidence interval depends only on the belief of the variance.

In the appendix, we derive a version of the model with imperfect observation. In that model, the width of the confidence interval for the next forecast is again directly related to a confidence interval around the previous forecast. It's length is a function of the degree of imperfect observation instead of simply one standard deviation.

3. Data and Summary Statistics

We utilize a set of stock market predictions made by financial executives in the quarterly Duke CFO Survey. Detailed information on the dataset is available in Ben-David et al. (2013). The primary survey question we are interested in asks:

> *Over the next year, I expect the annual S&P 500 return will be:*
>
> *- There is a 1-in-10 chance the actual return will be less than __%.*
>
> *- I expect the return to be: __%.*
>
> *- There is a 1-in-10 chance the actual return will be greater than __%.*

This question gives both a point estimate and a confidence interval around that estimate. In total, we have 13,757 responses over the 2005 – 2017 period. Figure 1 illustrates that we are able to construct a very large panel of data by connecting survey responses over time. Almost 1,000 executives have responded to the survey exactly twice while for many we are able to construct much longer panels. For example, over 400 respondents have responded to the survey at least nine times.

Figure 1: Overview of Panel Length

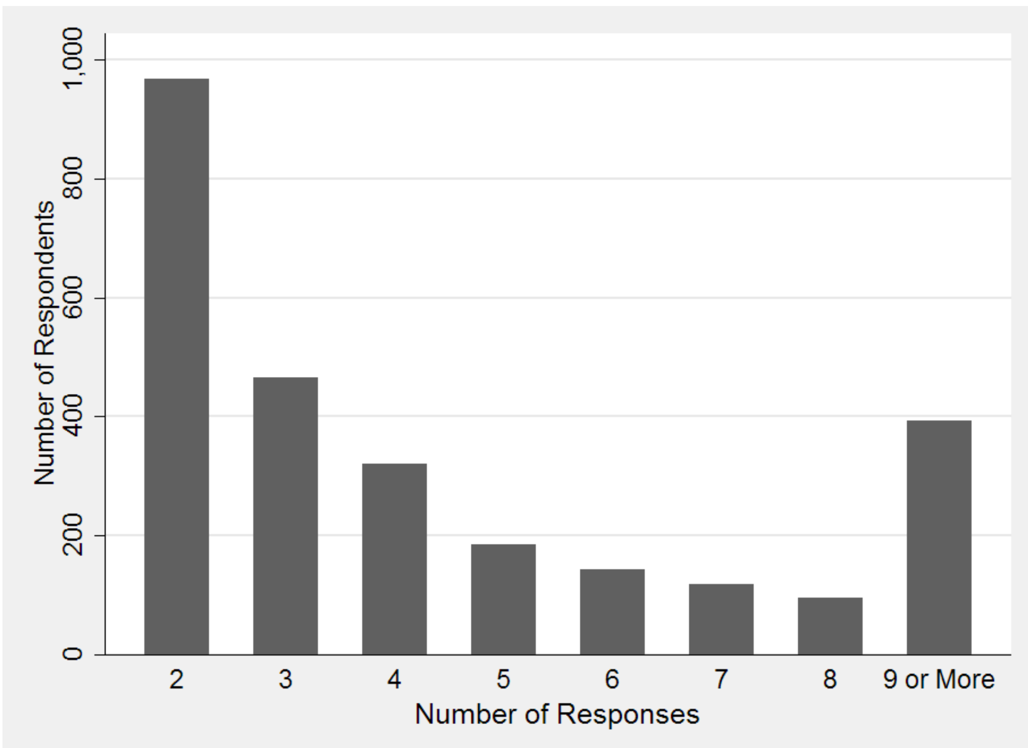

The figure presents the number of individual forecasters who responded to the survey a given amount. For example, almost 1,000 forecasters responded exactly once, and approximately 450 responded exactly three times. These amounts include all responses and not necessarily time-sequential responses.

In the context of learning, we focus on pairs of responses from the same respondent that are four quarters apart. Because respondents are asked to forecast the one-year-ahead S&P 500 return, they need four quarters after their response to observe the realized return and make a new one-year-ahead

Table 1: Confidence Interval Width Summary Statistics

	All	Miss	Hit
Mean	15.8	12.9	23.1
Median	12.0	10.0	20.0
25th Percentile	7.0	5.0	12.0
50th Percentile	20.0	17.0	30.0
N	4,249	3,031	1,218

Summary statics of the width of the confidence interval in the second of each four-quarters-apart pair of forecasts. The first column includes observations where the realized return fell within the forecast interval and the second column includes observations where the realized return missed the confidence interval.

forecast. The sample used in the bulk of our analysis is thus composed of 4,249 pairs of observations. For each observation, we record the respondent's forecast and confidence interval, the realized return, and the respondent's new forecast and confidence interval generated after the realized return.

Table 1 summarizes the widths of the confidence intervals for each of the second responses in the 4,249 pairs of responses. In the appendix, we present results for the entire sample of responses and note that they are very similar.

Among all forecasts, the mean width of the confidence interval is 15.8 percentage points. We divide the sample into the those whose confidence intervals contained the realized return, or those who hit the confidence interval, and those who did not. Among those who hit the confidence interval, the mean width is a much higher 23.1 percentage points. Note that since these are the second responses, we are conditioning on the second forecast being correct in another four quarters, and thus this is *not* evidence of learning. The mean width of the responses that miss the confidence interval is a smaller 12.9 percentage points. In general, respondents who hit the confidence interval had wider confidence intervals.

The model presented above assumes that the forecaster's beliefs are symmetrically centered around their point forecast. We construct a simple measure of asymmetry as:

$$A = \frac{r_U + r_L}{2} - r_P, \tag{3}$$

where r_P is the point forecast, r_U is the upper bound, and r_L is the lower bound. When $A = 0$, the interval is symmetric; when $A > 0$, the interval is skewed to the right, and when $A < 0$, the interval is skewed to the left. Figure 2 plots the distribution of this measure of asymmetry for the first response in the main sample of observations. The appendix contains a similar figure for the entire sample and, again, the shape of that distribution is largely the same.

Figure 2: Distribution of Asymmetry in Main Sample

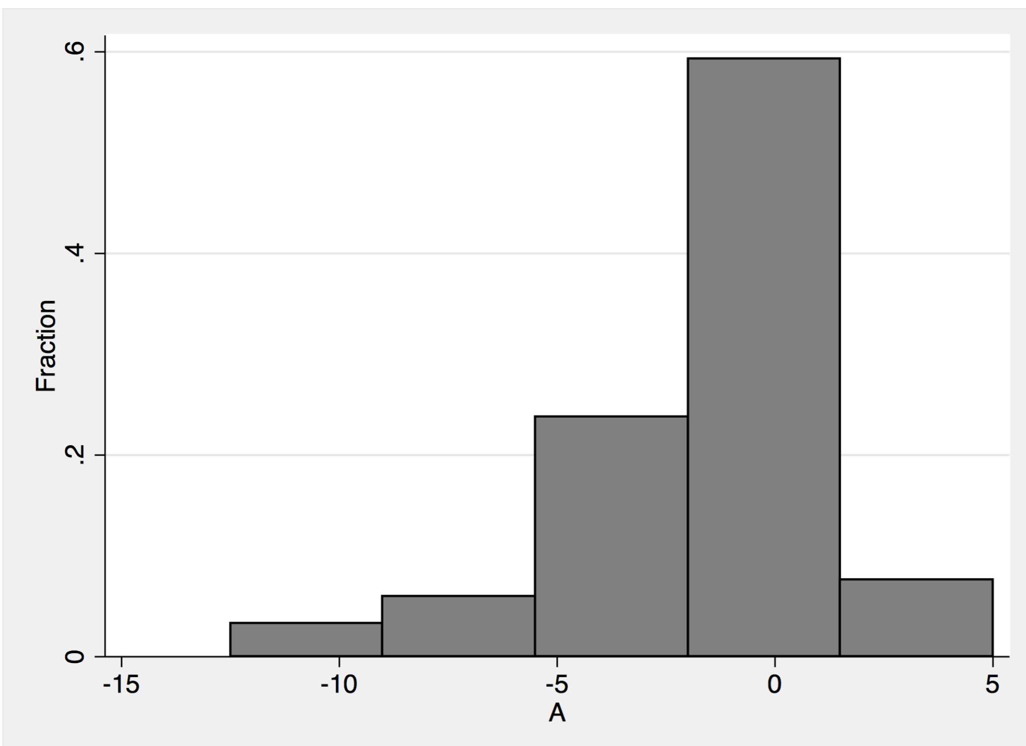

Th distribution of the measure of asymmetry in equation (3). We exclude the top and bottom one-percentile of observations.

The mean asymmetry is approximately -1.9 percentage points and there are many more left-skewed confidence intervals than right-skewed. However, approximately 20% of the observations in the main sample have exactly zero asymmetry, and just over 60% are symmetric to within two percentage points. Given this, it is not unreasonable to initially examine a model that imposes symmetry.

4. Results

Motivated by our model, we look to the behavior of confidence intervals as evidence of learning. In particular, does a respondent's confidence interval widen after they observe that the return was outside their confidence interval four quarters ago? Does it narrow if the return was inside their previous confidence interval?

We find that the answers to both questions are affirmative. Looking deeper, we investigate whether respondents adjust their confidence intervals differently if they miss high versus low. Finally, departing from the asymmetric view of the confidence interval, we document whether changes to the confidence interval are driven by adjusting the lower bound, upper bound, or both.

4.1. Changes in CI Width

The first row of Table 2 shows the mean change in the confidence interval between the response today and the response four quarters ago after observing the return today. We also report in parentheses the t-statistics from a two-sided t-test that the mean change is equal to zero.[4]

On average, respondents who hit the confidence interval in the first of the two forecasts respond by narrowing their next confidence interval by approximately 2.8 percentage points. On the other hand, respondents who

[4] In unreported results, we also verify that the differences are statistically different from one another.

Table 2: Changes in Confidence Interval Widths

	Hit	Miss	Miss High	Miss Low
Δ CI Width	−2.81	1.73	0.45	4.15
	(−8.34)	(9.40)	(2.03)	(13.36)
Δ Upper Bound	−2.81	0.92	1.43	−0.04
	(−13.42)	(7.50)	(11.02)	(−0.15)
Δ Lower Bound	0.00	−0.81	0.98	−4.18
	(−0.01)	(−4.74)	(4.81)	(−15.16)
Observations	1,350	2,899	1,897	1,002

Changes between pairs of observations four quarters apart. Reported in parentheses are t-statistics from a two-sided test that the difference is equal to zero. The first column includes observations where the realized return fell within the forecasted interval and the second column includes observations where the realized return missed the confidence interval. The third column includes observations where the realized return was greater than the upper bound of the confidence interval, while the fourth column includes observations where the realized return was less than the lower bound of the confidence interval.

miss the confidence interval respond by widening their confidence interval by approximately 1.7 percentage points. This is in line with the predictions of our model and is consistent with learning.

4.2. Asymmetries in Learning

We turn to analyzing the asymmetries evident in our data in two ways.

First, the second and third rows decompose the changes in the confidence interval to changes in the upper and lower bounds. Note that $\Delta W_{CI} = \Delta UB - \Delta LB$ where W_{CI} is the width of the confidence interval, UB is the upper bound, and LB is the lower bound.

Respondents who hit the CI appear to adjust by decreasing their upper bound by the entire amount. On the other hand, respondents who miss the CI appear to increase the width of their confidence intervals somewhat more symmetrically, increasing the upper bound by 0.9 pp while simultaneously decreasing the lower bound by 0.8 pp.

The second asymmetry we analyze is whether respondents adjust differently if they miss their confidence intervals on the up- versus down-side. In particular, the third column analyzes the changes in the confidence intervals for respondents who missed the confidence interval and whose point forecast was higher than the observed return. The fourth column does the same for those who predicted too low of a return.

Respondents who miss high appear to adjust very little, increasing the width of their confidence intervals by only 0.5 pp. However, inspecting the changes in the upper and lower bounds gives credence to a different form of learning: instead of increasing the width of their confidence interval, these respondents shift the entire interval upwards. After increasing the upper bounds of their confidence intervals by approximately 1.4 pp, respondents who missed high also increase the lower bounds by approximately 1.0 pp. Since both the lower and upper bounds increase, the total width changes very little.

Respondents who miss low adjust by increasing the width of their confidence interval by approximately 4.2 percentage points. Essentially all of this change is driven by decreasing the lower bound.

Figure 3: Relationship Between Size of Miss and Change in CI Width

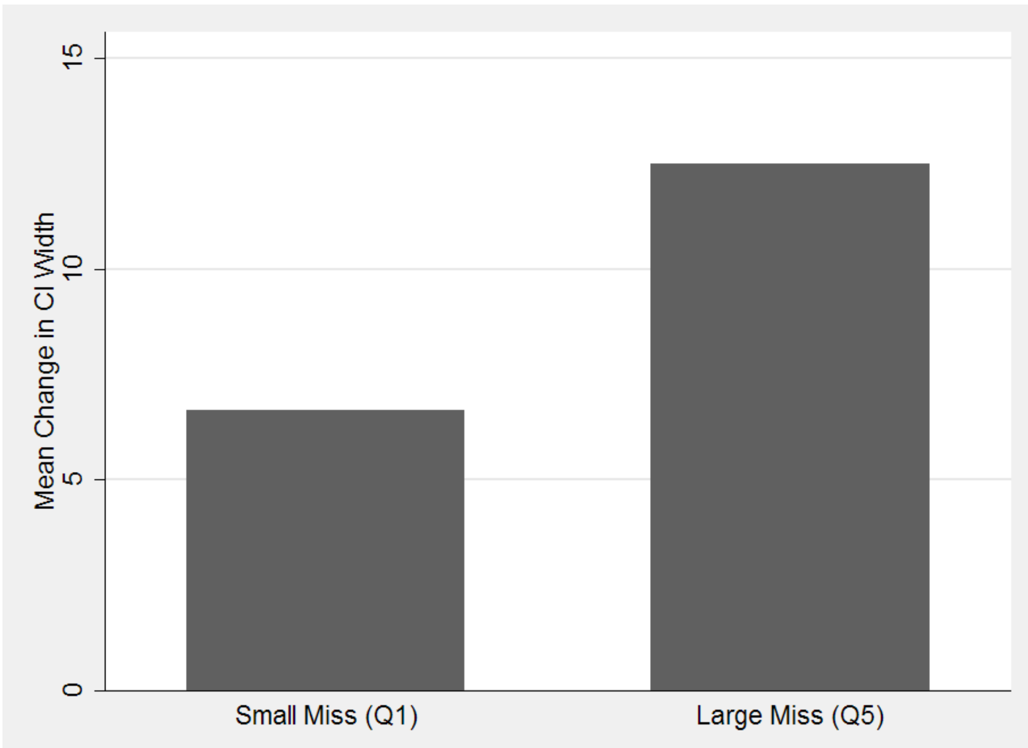

The relationship between the size of the miss, measured as the absolute difference between the forecast and realized return, and the change in the CI width. Observations are sorted by size of miss and divided into five quintiles. Within each quintile, the average change in the CI width is calculated. Presented in the figure is the average change in the CI width for the smallest (Q1) and largest (Q5) quintiles.

4.3. Size of Miss and Change in CI Width

We have thus far provided evidence that upon hitting or missing the CI, respondents adjust by narrowing or widening their confidence intervals. However, our model further predicts that the size of the miss should directly affect the size of the confidence interval width change. In particular, if the forecaster misses the CI, larger misses should imply more widening of the confidence interval.

Figure 3 shows the relationship between the size of the miss and the size of the change in the confidence interval. Conditioning on those who missed the confidence interval and responded by widening their next confidence interval, we further divide forecasters into five quintiles according to the absolute difference between their point forecasts and the realized return. Respondents in the first quintile widened their confidence intervals by an average 6.7 percentage points, while respondents in the fifth quintile widened by an average 12.5 percentage points.

5. Conclusion

[To be completed.]

References

Barber, B.M., Odean, T., 2013. The behavior of individual investors, in: Constantinides, G.M., Harris, M., Stulz, R.M. (Eds.), Handbook of the Economics of Finance. Elsevier. volume 2. chapter 22, pp. 1533–1570.

Ben-David, I., Graham, J.R., Harvey, C.R., 2013. Managerial miscalibration. The Quarterly Journal of Economics 128, 1547–1584.

Bloom, N., Davis, S.J., Foster, L., Lucking, B., Ohlmacher, S., Saporta-Eksten, I., 2017. Business-level expectations and uncertainty. Working paper .

Gennaioli, N., Ma, Y., Shleifer, A., 2016. Expectations and investment, in: Eichenbaum, M., Parker, J. (Eds.), NBER Macroeconomics Annual. volume 30.

Gervais, S., Odean, T., 2001. Learning to be overconfident. The Review of Financial Studies 14, 1–27.

Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. Review of Financial Studies 27, 714–746.

Kuchler, T., Zafar, B., 2017. Personal experiences and expectations about aggregate outcomes. FRB of NY Staff Reports .

Kuhnen, C.M., 2014. Asymmetric learning from financial information. Journal of Finance 70, 2029–2062.

Malmendier, U., Nagel, S., 2011. Depression babies: Do macroeconomic experiences affect risk taking? The Quarterly Journal of Economics 126, 373–416.

Moore, D.A., Schatz, D., 2017. The three faces of overconfidence. Social and Personality Psychology Compass 11.

Moore, D.A., Tenney, E.R., Haran, U., 2016. Overprecision in judgment, in: Keren, G., Wu, G. (Eds.), Handbook of Judgement and Decision Making. Wiley Blackwell. chapter 6.

Vissing-Jorgensen, A., 2003. Perspectives on behavioral finance: Does "irrationality" disappear with wealth? evidence from expectations and actions, in: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual. volume 18.

Appendix A. Details on Bayesian Learning

The marginal prior distribution of the mean belief is normally distributed with mean given by the hyperparameter ϕ:

$$\mu_{t-1} = E[\hat{\mu}_{t-1}] = \phi_{t-1}, \tag{A.1}$$

The hyperparameter is defined as:

$$\phi_s = \frac{\kappa_0 \phi_0 + s\bar{r}_s}{\kappa_0 + s} = \frac{\kappa_0 \phi_0 + (s-1)\bar{r}_{s-1} + r_s}{\kappa_0 + s},$$

where $\bar{r}_s = \frac{1}{s}\sum_{i=1}^{s} r_i$ is the sample average of all observations up to time s. The first equality shows that the belief of the mean is a weighted average between the initial prior ϕ_0 and the sample mean $\bar{r}_s$, where, as discussed above, κ_0 is the number of observations used in constructing the prior hyperparameter ϕ_0. The second equality illustrates that the new sample average, $\bar{r}_s$, can be expressed as a weighted average between the newly observed return and the previous sample average, highlighting that the evolution of the hyperparameter is driven by the newly observed return.

Proposition 2 (Evolution of the Mean Belief)**.**

(a) *If the observed return is greater than the prior belief of the mean, or $r_t > \mu_{t-1}$, the posterior belief of the mean is greater than the prior belief:*

$$\mu_t > \mu_{t-1}.$$

(b) *If the observed return is less than the prior belief of the mean, or $r_t < \mu_{t|t-1}$, the posterior belief of the mean is less than than the prior belief:*

$$\mu_t < \mu_{t-1}.$$

(c) *If the observed return is equal to the prior belief of the mean, or $r_t = \mu_{t-1}$, then the posterior belief is equal to the prior belief, or $\mu_t = \mu_{t-1}$.*

In (A.1), the belief of the mean perfectly tracks the evolution of the hyperparameter ϕ, which is simply the weighted sample average. If the newly observed return is larger (in absolute terms) than the previous sample average, the new sample average increases; by extension, if the newly observed return is higher than the previous belief of the mean, the new belief of the mean is larger. The same logic applies if instead the newly observed return is smaller than the previous belief of the mean.

Appendix B. Bayesian Learning with Imperfect Observation

In the last section, the evolution of the forecaster's belief of the variance is connected to confidence interval around their forecast that is roughly one-standard deviation wide. We now introduce a layer of heterogeneity in our model with the goal of varying the size of this confidence interval, motivated by the observation that in our data, forecasters respond to the same realized returns in different ways.

Assume that the forecaster is unable to perfectly observe the return in each period. In particular, the forecaster observes a constant factor $\theta > 0$ of the return. Denote the observations of the return as $\tilde{r}_t = \theta r_t$. In this section, we analyze how the forecaster's beliefs will evolve relative to the perfect observation case above where $\theta = 1$. Let $\tilde{\mu}_s$ and $\tilde{\sigma}_s$ denote the forecaster's beliefs at time s given the imperfectly observed returns, and similarly for the hyperparameters $\tilde{\beta}_s$ and $\tilde{\phi}_s$.[5]

The relationships between the hyperparameters and the beliefs remain the same:

$$\tilde{\sigma}^2_{s-1} = \frac{\tilde{\beta}_{s-1}}{\alpha_{s-1}} \tag{B.1}$$

$$\tilde{\mu}_{s-1} = \tilde{\phi}_{s-1}, \tag{B.2}$$

[5] As noted above, α and κ evolve mechanically and independently of any observed data, so $\tilde{\alpha}_s = \alpha_s$ and $\tilde{\kappa}_s = \kappa_s$ for all s.

The evolution of the hyperparameters is now affected by the fact that the observations are imperfectly observed. Let $\tilde{\bar{r}}_s$ denote the sample average of all imperfectly observed returns up to time s. Note that:

$$\tilde{\bar{r}}_s = \frac{1}{s}\sum_{i=1}^{s} \tilde{r}_s = \theta\frac{1}{s}\sum_{i=1}^{s} r_s = \theta\bar{r}_s.$$

We assume that under imperfect information, the agent's initial priors are equal to imperfect observations of the initial priors above. Specifically, $\tilde{\phi}_0 = \theta\phi_0$ and $\tilde{\beta}_0 = \theta^2\beta_0$. These assumptions simplify the derivations below but are not necessary for our primary qualitative results, especially since as the forecaster observes more returns, less weight is placed on the initial priors.

We begin by analyzing the hyperparameter controlling the mean belief. At any time s,

$$\tilde{\phi}_s = \frac{\kappa_0\tilde{\phi}_0 + s\tilde{\bar{r}}_s}{\kappa_0 + s} = \frac{\kappa_0\theta\phi_0 + s\theta\bar{r}_s}{\kappa_0 + s} = \theta\frac{\kappa_0\phi_0 + s\bar{r}_s}{\kappa_0 + s} = \theta\phi_s.$$

If both the prior and every observation are imperfectly observed by a constant factor, it follows that the imperfectly observed sample average is the actual sample average multiplied by that constant factor, and therefore

$$\tilde{\mu}_{s-1} = \theta\mu_{s-1}. \tag{B.3}$$

The imperfectly observed hyperparameter controlling the variance belief can be expressed at any time s as:

$$
\begin{aligned}
\tilde{\beta}_s &= \tilde{\beta}_0 + \frac{1}{2}\sum_{i=1}^{s}(\tilde{r}_i - \bar{\tilde{r}}_s)^2 + \frac{\kappa_0 s}{\kappa_0 + s}\frac{(\bar{\tilde{r}}_s - \tilde{\mu}_{0|0})^2}{2} \\
&= \theta^2\beta_0 + \frac{1}{2}\sum_{i=1}^{s}(\theta r_i - \theta\bar{r}_s)^2 + \frac{\kappa_0 s}{\kappa_0 + s}\frac{(\theta\bar{r}_s - \theta\mu_{0|0})^2}{2} \\
&= \theta^2\beta_0 + \theta^2\frac{1}{2}\sum_{i=1}^{s}(r_i - \bar{r}_s)^2 + \theta^2\frac{\kappa_0 s}{\kappa_0 + s}\frac{(\bar{r}_s - \mu_{0|0})^2}{2} \\
&= \theta^2\beta_s
\end{aligned}
$$

It again follows that if every observation, including the initial prior, is imperfectly observed by a constant factor θ^2, then the imperfectly observed sum of squares is the the actual sum of squares multiplied by that constant factor.

As above, if the newly (imperfectly) observed return is larger than the mean belief, the mean belief increases. Similarly, if the newly (imperfectly) observed return is outside of a one-standard deviation interval of the previous mean belief, the variance belief increases. It is interesting, however, to inspect this mechanism more closely and understand how imperfectly observing the return affects the mean variance evolution relative to the perfect observation case.

Proposition 3 (Evolution of the Variance Belief With Imperfect Observation)**.** *If*

$$(\tilde{r}_t - \tilde{\mu}_{t-1})^2 > m_t\tilde{\sigma}^2_{t-1},$$

then $\tilde{\sigma}^2_t > \tilde{\sigma}^2_{t-1}$.

Corollary 2. *With imperfect observation, if the condition in proposition 3 is satisfied, then the forecaster's confidence interval widens.*

Proposition 3 and corollary 2 are comparable to proposition 1 and corollary 1. Again assuming $m_t \approx 1$, the condition in proposition 3 can be rewritten as

$$(r_t - \mu_{t-1})^2 > \theta^{-2}\tilde{\sigma}_{t-1}^2.$$

On the left-hand side of this inequality is the squared difference between the return and the mean belief under perfect observation. On the right hand side is the forecaster's belief of the variance under imperfect observation multiplied by the inverse of the observation factor. This inequality can be rewritten as:

$$r_t \notin [\mu_{t|t-1} - \theta\tilde{\sigma}_{t-1}, \mu_{t-1} + \theta\tilde{\sigma}_{t-1}]$$

With imperfect observation, the variance belief evolves as-if the forecaster perfectly observes the return and, in determining whether this new observation should increase or decrease her belief of the variance, again constructs an interval using her beliefs of the mean and variance. Recall that with perfect observation above, this interval is centered around the belief of the mean and one-standard deviation wide (using the belief of the variance).

With imperfect observation, the interval is again centered around the same belief of the mean, but the width is controlled by the variance of the belief under imperfect observation, $\tilde{\sigma}_{t-1}$, and the observation factor, θ. Specifically, the interval is (θ^{-1})-standard deviations wide, where the standard deviation is constructed from the belief of the variance under imperfect observation.

For example, assume that $\theta < 1$ such that the forecaster consistently under-observes the return. This means that in every period, the forecaster's beliefs of the mean and variance are smaller than they otherwise would have been. In turn, this generates a narrower confidence interval relative to the perfect observation case, which implies that since the forecaster is using the smaller variance belief, her confidence level must be higher than one standard deviation. In particular, the new confidence interval is (θ^{-1})-standard deviations wide, and since $\theta^{-1} > 1$, this confidence interval is indeed wider.

Appendix C. Proofs