Journal of Behavioral and Experimental Economics 57 (2015) 55–63

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Economics

journal homepage: www.elsevier.com/locate/socec

How do experts update beliefs? Lessons from a non-market environment

Michael Sinkey*

University of West Georgia, Department of Economics, Richards College of Business, 1601 Maple Street, Carrollton, GA 30118, United States

ARTICLE INFO

Article history:
Received 14 August 2013
Revised 14 April 2015
Accepted 14 April 2015
Available online 17 April 2015

Keywords:
Cognitive biases
Confirmatory bias
Experts
Bayesian updating
Belief formation

ABSTRACT

Experts are regularly relied upon to provide their professional assessments in a wide array of markets (e.g., asset pricing, stock and bond ratings, expert witnesses, forecasting), which frequently have characteristics that may generate incentives for experts to provide biased analyses. I ask how experts update beliefs in a relatively simple environment with minimal market incentives. Using data from the Associated Press (AP) Top 25 Poll for college football I find that many standard sets of Bayesian beliefs are rejected by the data, and that experts, while using Bayes' rule, may still be subject to similar biases as non-experts, including confirmatory bias and lagged signal response, which may be symptomatic of inattention, voter heterogeneity, and signal reassessment. In more complex environments, experts may have strong incentives to substantially deviate from Bayes' rule, biasing expert predictions in unknown directions.

1. Introduction

Experts are frequently courted for their ability to give unbiased, knowledgeable assessments about future events or unknown parameters. There are large markets for their services as analysts, forecasters and expert witnesses. While economists typically assume these experts engage in Bayesian updating (Ottaviani and Sorensen, 2006b), the economics and finance literature also reports gaps between experts' true beliefs and their announced beliefs. Theoretical explanations of this gap include preference differences (Crawford and Sobel, 1982; Krishna and Morgan, 2001), herding toward a unified prediction (Banerjee, 1992; Scharfstein and Stein, 1990; Trueman, 1994), and paternalism. Morris (2001). Additionally, because experts benefit from being perceived as smart, reputational concerns may play a role. Advisors may try to mirror their advisee's best estimate (Ottaviani and Sorensen, 2006a) or avoid contradicting the biases of a decision-maker (Morris, 2001). Empirically, explicit monetary incentives may sway predictions. For example, the mutual funds from fund families which advertise in a publication tend to receive more favorable ratings from the publication's analysts (Reuter and Zitzewitz, 2006). In a similar vein, Malmendier and Shantikumar (2007) found that financial analysts frequently rate securities more favorably when the analyst's firm underwrites the security.

Taken as a whole, while experts are asked to provide unbiased opinions, there are many reasons why they cannot be relied upon to do so. Because of these confounds, it is difficult to distinguish rationally-altered, strategic predictions from predictions impaired by behavioral or cognitive limitations. Empirically comparing such contaminated responses to theoretical predictions based on Bayes' Rule would thus yield only limited insight into whether experts' reported predictions reflect Bayesian updating absent these confounding influences.

This study compares voting patterns in the Associated Press College Football Poll (AP Poll) to patterns consistent with Bayesian updating. In the AP Poll, between 60 and 70 college football writers and analysts individually submit weekly team rankings. Individual rankings are aggregated and not shared with other voters while voting, and these votes determine the week's rankings. Experts are unpaid for their opinions and, while voter ballots were available during some portions of the sample period, they were not widely disseminated or analyzed, alleviating reputation concerns. These data are matched with informative signals about performance generated weekly from college football games, such as winning and losing, and betting lines – which proxy for the predicted margin of victory in a given game – to estimate how experts respond to new data when updating their beliefs.

Because the AP Poll is an attractive testing ground, other researchers have also used the AP Poll to examine belief updating. This paper expands on efforts to study Bayesian updating using data from the AP Poll, and is most closely related to two studies that use the poll: Andrews et al. (2012) and Stone (2013). Both this work and Andrews et al. (2012) find evidence for confirmatory bias. Andrews et al. (2012) use a regression-discontinuity approach to identify the magnitude of confirmation bias by looking how voters process small differences between the score differential and the betting line. They find that voters substantially upgrade teams that are already in the poll based on these small differences, which is striking since these teams are otherwise

* Tel.: +1 6788395166.
E-mail address: msinkey@westga.edu

http://dx.doi.org/10.1016/j.socec.2015.04.003
2214-8043/

exchangeable. Stone (2013) finds evidence that pollsters may overreact and underreact to new information—in particular, they overreact to losses by highly-ranked teams and underreact to road victories.[1]

In this paper, I describe how college football voters should respond to signals about a particular team's quality. In college football, information related to a team's quality varies in intensity and is received weekly. Teams vary in quality over time and expert benchmark beliefs are derived from Bayes' rule. I then describe specific deviations from Bayesian updating. The first deviation is a lagged response to signals. Bayesian updating stipulates that individuals respond to signals when they are received. I describe instances when college football voters incorporate old information in the formation of posterior beliefs despite the fact that those signals should be fully incorporated into prior beliefs. This lagged response may demonstrate a rational inattention by experts to certain signals, only to examine those ignored signals later, or signal reassessment, which I define as occurring when individuals use information contained in previous signals to update their current beliefs outside of and in addition to their Bayesian inclusion in prior beliefs.

The second deviation is confirmatory bias, where one "interprets ambiguous signals as confirming his current hypotheses about the world (Rabin and Schrag, 1999)." Rabin and Schrag describe confirmatory bias as arising from three distinct "information-processing problems": interpreting ambiguous signals, interpreting statistical evidence to assess correlation between "phenomena that are separated by time," and "hypothesis-based filtering", where ambiguous information that has been interpreted to be consistent with a previous hypothesis is then used as further evidence for the same hypothesis. This paper concentrates on the second and third type of processing problems. In particular, I analyze how upgrades and downgrades in the poll – changes in prior beliefs – affect whether later signals are interpreted more or less favorably and how old information is included in the formation of new posterior beliefs. In turn, this provides a proxy for the extent to which hypothesis-based filtering may exist in a relative simple environment.

The data and analysis yield evidence for both lagged responses to previous data and confirmatory bias. Using the differences in announced beliefs before and after a game (i.e., from their prior belief to their posterior after viewing a game) I find that experts reevaluate the quality of some signals (performance against betting lines from the previous week), but not others (wins and losses from the previous week) despite the fact that these signals are both well correlated and observed at the same time. Both signals should be completely included in prior beliefs; because they are not, I interpret this finding as a lagged signal response, which is inconsistent with Bayesian updating. Experts use changes in previous beliefs when evaluating new information; when experts believe that team quality is declining, they interpret new signals more pessimistically and have less-optimistic incorporations of old signals. When experts believe that team quality is improving, they interpret the *same* new signals more positively, suggesting the presence of confirmatory bias. Consequently, if experts engage in Bayesian updating, they may also be subject to several of the psychological biases that affect non-experts.

2. Lagged signal response and confirmatory bias in economic settings

There is evidence for Bayesian deviations in laboratory settings. Much of this literature deals with relatively simple situations where subjects are asked to report probabilities of certain events occurring based on observed data and underlying base rates (Tversky and Kahneman, 1972; 1974; Grether, 1992; El-Gamal and Grether, 1995; Holt and Smith, 2009, and others). Those probabilities are then compared to the actual probabilities that would be generated under Bayes' rule.

2.1. Lagged signal response

This paper coincides with previous studies that find lagged responses to signals. A compelling and oft-cited explanation for why individual beliefs do not respond immediately to signals is inattention (rational or irrational), which has been well-documented in the finance literature. Dellavigna and Pollet (2009) find that investors react more slowly to information disseminated on Fridays but that prices drift toward their full-information counterparts, indicating that inattention is a potential source of lagged signal response. Hirshliefer et al. (2009) find that extraneous news causes individuals to react more slowly to pertinent signals, indicating potential investor distraction. Inattention may be present in college football as well, since voters are tasked with following a large number of teams but are not compensated. Heterogeneity in voter attention may cause a lagged signal response, as voters who did not pay attention to signals in one week may collect both past signals and present signals in a current week.

Lagged signal response may coincide with overreaction and underreaction to information. Stone (2013) finds that, for the Associated Press Top 25 Poll, voters may overreact and underreact to information, and are more likely to correctly assess signals that are more salient. Griffin and Tversky (1992) indicate that "the tendency to focus on the strength of the evidence leads people to neglect or underweight other variables." previous researchers in finance have found evidence that individuals both respond initially to signals and continue to respond to those same signals at later dates, despite the fact that the information contained in those signals should be fully contained in the prior.

In this paper I also introduce and propose an alternative explanation for lagged response to signals, which I define as "signal reassessment." Signal reassessment occurs whenever an individual uses information contained in previous signals to update current beliefs outside of and in addition to the inclusion of those signals in prior beliefs. Essentially, this is when an individual double-counts data that should be already included in prior beliefs. Signal reassessment may occur if an individual discards his or her prior belief. Ortoleva (2012) models the possibility that an individual may receive a signal that is so unexpected that he or she discards an old prior belief to select a new prior belief that is consistent with the evidence that is received. However, the process of discarding an old model for a new model is non-Bayesian and may manifest as double-counting of evidence.

Reassessment may also be driven by an existing cognitive bias, such as confirmatory bias. If an individual believes strongly in a hypothesis, then he or she may choose to bolster the hypothesis by continuing to update beliefs using previously-acquired data rather than new. To illustrate, a manager may continue to use a past accomplishment as evidence of an employee's continued productivity if the manager believed the employee was productive before. I consider the process of double-counting that same accomplishment to be signal reassessment.

2.2. Confirmatory bias

The sharpest examples of confirmatory bias are found in the psychology literature (Nickerson, 1998). Hypothesis-based evaluations are inconsistent with Bayesian updating, but are nonetheless consistently seen in experimental settings. Notable examples include instances where subjects attempt to evaluate academic performance in the light of socioeconomic class (Darley and Gross, 1983), where individuals evaluate the validity of academic research (Mahoney, 1977),

[1] Other studies examining belief changes in the AP Poll include (Stone and Zafar, 2014), who find that voters may engage in social learning, Coleman et al. (2010) who find that individuals may exhibit biases toward their home state, and Logan (2010) who finds that voters treat losses later in the season differently from losses earlier in the season.