

On-line Appendix

Learning Together Slowly: Bayesian Learning About Political Facts

Journal of Politics, 2017

Seth J. Hill
University of California, San Diego*

March 8, 2017

Contents

A	Crossover scoring method	2
B	Derivation of logit specification of Bayesian learning	2
C	Details on post-stratification weights	3
D	Example individual learning	4
E	Details of second experiment for learning on abstract fact	4
F	Additional tables and figures	6
G	Comparison to tipping point model	6
H	Main tables without post-stratification weighting	13
I	Consequences of measurement error	13
J	Experiment instructions	18

*Department of Political Science, 9500 Gilman Drive #0521, La Jolla, CA 92093-0521; sjhill@ucsd.edu,
<http://www.sethjhill.com>.

A Crossover scoring method

The crossover scoring method elicits probabilistic beliefs from participants with incentives aligned for truthful reporting. The design asks participants for what probability p they would be indifferent between receiving a payment with probability p and receiving a payment if their answer is correct. To elicit this probability with incentives, the subject is told that after the probability p is reported, a number y will be drawn at random from the uniform distribution on $[0, 1]$. If $y > p$, the subject will enter a lottery which pays incentive v with probability y and 0 with probability $1 - y$. If $y < p$, the subject is paid incentive v if the statement is true, and 0 if it is false.

To see that a subject maximizes their chance of receiving incentive v by accurately reporting their true belief, consider a subject with true belief p^* who reports a belief \hat{p} . With the uniform draw of y and the mechanism above, reporting belief \hat{p} means that the subject is paid based upon the truth of the statement with probability \hat{p} and enters the lottery with probability $1 - \hat{p}$ (because $\Pr(y < \hat{p}) = \hat{p}$ for a uniform random variate y). Their expected payout under the truth mechanism given true beliefs p^* is vp^* . Expected payout under the lottery is $v[(1 - \hat{p})/2 + \hat{p}]$, the midpoint of the uniform distribution of y conditional on $y > \hat{p}$ (i.e., ranging from \hat{p} to 1). Then, the expected value of giving report \hat{p} given true belief p^* is

$$EV[\hat{p}] = v\hat{p}p^* + v(1 - \hat{p})\left(\frac{1 - \hat{p}}{2} + \hat{p}\right). \quad (\text{A1})$$

To see that setting $\hat{p} = p^*$ maximizes expected payout, take the derivative and solve for the F.O.C.:

$$dEV/d\hat{p} = vp^* + v(1 - \hat{p})\left(-\frac{1}{2} + 1\right) + v\left(\frac{1 - \hat{p}}{2} + \hat{p}\right)(-1) \quad (\text{A2})$$

$$= vp^* + v(1 - \hat{p})\frac{1}{2} - v(1 - \hat{p})\frac{1}{2} - v\hat{p}.$$

$$0 = vp^* - v\hat{p}$$

$$\hat{p} = p^*. \quad (\text{A3})$$

Thus, subjects maximize incentives when reporting their true beliefs, $\hat{p} = p^*$.

B Derivation of logit specification of Bayesian learning

I show here how Bayes' Rule from Equation 1 can be transformed to the regression model of learning in Equations 2 and 3. As before, consider a factual statement T with a probabilistic prior belief that it is true $\Pr(T)$ [and corresponding prior belief the statement is false $1 - \Pr(T) = \Pr(F)$] and a probabilistic posterior belief $\Pr(T|S = s)$ after receiving a stochastic signal $s \in \{t, f\}$, with f indicating false and t indicating true. The regression specification with dependent variable the logit of the posterior beliefs, $\log[\Pr(T|S = s)/\Pr(F|S = s)]$ can be derived by letting

$$\Pr(S = s) = \Pr(S = s|T)\Pr(T) + \Pr(S = s|F)\Pr(F)$$

be the probability of the data, and the two Bayes' Rule specifications of posterior beliefs be

$$\begin{aligned}\Pr(T|S = s) &= \Pr(T) \frac{\Pr(S = s|T)}{\Pr(S = s)} \\ \Pr(F|S = s) &= \Pr(F) \frac{\Pr(S = s|F)}{\Pr(S = s)}.\end{aligned}$$

Then, the posterior odds are

$$\begin{aligned}\frac{\Pr(T|S = s)}{\Pr(F|S = s)} &= \frac{\Pr(T)\Pr(S = s|T)/\Pr(S = s)}{\Pr(F)\Pr(S = s|F)/\Pr(S = s)} \\ &= \frac{\Pr(T)}{\Pr(F)} \times \frac{\Pr(S = s|T)}{\Pr(S = s|F)}.\end{aligned}$$

Taking logs of both sides,

$$\text{logit}[\Pr(T|S = s)] = \text{logit}[\Pr(T)] + \log[\Pr(S = s|T)/\Pr(S = s|F)].$$

Noting that the signals $S = t$ and $S = f$ have similar forms but with different likelihood ratios, we can construct the combined logit specification of Bayesian learning in round t for subject i having observed signal $S_{it} = s \in \{t, f\}$ after prior beliefs $\Pr_{i,t-1}(T)$

$$\begin{aligned}\text{logit}[\Pr_{it}(T|S_{it} = s)] &= \text{logit}[\Pr_{i,t-1}(T)] + \mathbf{1}[S_{it} = t] \times \log[\Pr(S = t|T)/\Pr(S = t|F)] \\ &\quad + \mathbf{1}[S_{it} = f] \times \log[\Pr(S = f|T)/\Pr(S = f|F)],\end{aligned}$$

where $\mathbf{1}[\cdot]$ returns a 1 when its argument is true, and 0 otherwise.

C Details on post-stratification weights

To help ameliorate potential non-representativeness of Mechanical Turk subjects, I asked survey questions exactly as they were asked on the 2014 Pew Polarization Survey that allow me to construct post-stratification weights using the `rake` function from the R library `survey` (R Development Core Team, 2015; Lumley, 2011). The weights rake to the marginal distributions of respondents to the Pew Survey, which was a nationally-representative telephone-based sample of 10,013 respondents surveyed January to March 2014. I rake to questions related to political confirmation bias and the MTurk sample composition: Census region, age, gender, education, marital status, party identification, ideology, favorability to the two parties, and three ideological policy questions.¹ The weighted distribution of subjects is more representative: 52 percent female (versus 55 percent unweighted), average age of 44 (36 unweighted), 34 percent four year college degree or more (48), 70 percent voting in 2012 (69), 46 percent Democrat (53), 39 percent Republican (28), and 35 percent conservative or very conservative (19). All aggregate statistics (regression coefficients, means, medians, etc.) in the main text use these stratification weights, although results are quite similar with unweighted analysis, presented in Appendix Section H.

¹ I trim the resulting weights to range from 1/8 to 8 to limit variance. The case with the largest pre-trimmed weight was a 65+ year old Northeastern male with some college or less education who reported being a conservative Republican. The case with the smallest pre-trimmed weight was a 18-29 year old Northeastern female with a 2-year college degree or more and a liberal Republican. Pew survey data accessed from <http://www.people-press.org/2014/03/16/2014-political-polarization-survey/>, and I used the Pew weights to construct the Pew target distributions.

D Example individual learning

In Figure A1, I present examples of learning at the level of the individual subject. Each frame presents that subject's elicited belief that the statement is true over the five rounds of each contest. Elicited beliefs are plotted with black circles and connected by the black line. The gray circles and lines present how a perfect Bayesian would respond to the signals received given the prior belief elicited from the subject in the first round and the set of the signals actually delivered. Along the x-axis I present the signal presented to the respondent in each round.

The upper left three frames plot the behavior of subject 300. In the first round of the first contest, the subject evaluated that the statement was true with a probability of 75. In the second round, they received a signal True and revised their beliefs up to 80. A perfect Bayesian with a prior of 75 would have updated after one signal to closer to 90, as indicated by the gray line – the smaller updating here is what I call cautious learning. In the following three rounds, the subject received three more True signals and responded in the appropriate direction in each case. This subject ended the five rounds almost exactly where Bayes' Rule suggests given initial beliefs and this set of signals. In the second contest, subject 300 again appears to be learning in a fashion similar to Bayesian but cautiously. In the IQ contest, the subject updated beliefs almost exactly as indicated by Bayes' Rule.


The other frames present the behavior of three other subjects in the experiment. Subjects 721 and 508 in the upper right and bottom left also appear to be responding to signals in an imperfect but Bayesian fashion. Subject 721 is also rather cautious in response to receiving four True signals in the first contest and to receiving four False signals about their performance on the IQ quiz. In each case, the subject revises their beliefs in the correct direction but much less than would be indicated by Bayes' Rule. The behavior of the fourth subject in the bottom right is an example of the set of subjects who did not respond to signals and effectively did not participate in the experiment. This participant did not change their beliefs in any round of any contest.

E Details of second experiment for learning on abstract fact

Between September 8 and 12, 2016, I recruited 395 participants aged 18 and older and U.S. citizens from Amazon.com's Mechanical Turk (MTurk) worker platform to participate in the second experiment. The design was mostly similar to the first experiment described in the main text. Here I note differences. Participants were paid a \$0.60 flat show-up fee rather than \$0.50 in the first experiment. Subjects participated in an IQ quiz but did not evaluate a fact about their IQ performance. Subjects also answered questions for a separate research study during the same time.

After the IQ-like quiz, subjects were again taught about the experiment and informed that for each round won, they would be paid a \$0.10 bonus, \$0.00 otherwise. As before, they were told that the signal from the computer would indicate that the correct answer was true or false, and that this signal would be correct three out of four times on average. In the second experiment, subjects evaluated each of two (rather than three) statements, and beliefs were again elicited for five rounds. Each subject evaluated one of the two political statements, drawn at random, from rows one and four of Table 1, the questions on household income change under President Obama and federal debt change under President Reagan. In contrast to the first experiment, each statement of fact had both a true and a false version, which was drawn at random for each subject. To make a false version of the Obama fact, the words "fell by more" were replaced with "fell by less." To make a false version of the Reagan statement, the words "more than doubled" were replaced with "was cut

Figure A1: Examples of individual-level updating of beliefs



Note: Each frame plots the respondent's beliefs over the five rounds of each of three contests (black circles and lines) compared to what a perfect Bayesian learning model would predict (gray circles and lines), given the prior belief from round 1 and the signals delivered in rounds 2 through 5 (presented along the x-axis).

by more than 50%.” Signals were delivered conditional on the truth of the statement, and after the experiment, responses and signals were recoded so that all subject responses were in the direction of true. The goal of this randomization was to control for any overall bias towards true or false from respondents, given that only two facts were queried about. As caution is estimated similarly in the two experiments, this change does not appear to have had much effect.

Each subject also evaluated a non-political ego-irrelevant fact meant to abstract away from any self-interest of the individual. The statement of fact had a true and a false version: “On January 8, 2012, the length of the day from sunrise to sunset in the city of Doha, Qatar was [less/more] than 11 hours” [true/false]. The order of the abstract and the political fact were randomized at the subject level. In between the two facts for this experiment, subjects had beliefs elicited about other statements of fact for the separate research study.

After completing the contests, participants again answered a series of survey questions about their demographics, political attitudes, and political behaviors. Payments via MTurk bonuses were calculated and delivered as in the first experiment.

F Additional tables and figures

This section presents additional Tables referenced in the main body, [A1](#), [A2](#), [A3](#), [A4](#), [A5](#), and [A6](#).

I present in Appendix Table [A3](#) variation in learning about political facts varies by individual characteristics updating by characteristic of the individual on partisan facts in the first experiment (excludes IQ rounds). Columns two and three compare primary voters to non-primary voters, with point estimates suggesting primary voters exhibit more bias. In columns four, five, and six, I find that moderates learn much closer to the Bayesian ideal and with less bias than liberals or conservatives. In columns seven and eight, I find that those who like politicians who compromise and work with others learn more from consistent and inconsistent signals and those who do not like compromise. Finally, in columns nine and ten, I find little difference in learning between those who donate or contact elected officials than those who do not.

G Comparison to tipping point model

In this section, I present a nonparametric evaluation of whether an alternative “tipping point” model better characterizes learning of political information than the Bayesian model in the first experiment. A tipping point model of learning suggests that the largest revisions of beliefs should be for subjects who receive a consistent set of signals, say TTT, FFF, TTTT, or FFFF, thus “tipping” them over into finally updating their beliefs. The Bayesian model of learning, in contrast, is memoryless. For Bayesian learning at any value of prior belief, a true or false signal has the same meaning regardless of the prior pattern of signals.

For each pattern of signals received by participants, I tabulate the average and median revision in beliefs to the most recent signal. For example, for subjects who received true signals in rounds two and three, I tabulate the average and median change in their beliefs from round three beliefs to round four beliefs for those whose round four signal was true versus false. Table [A7](#) presents mean and median revisions for each pattern, sorted descending by largest absolute revision in belief.² For both partisan and IQ contests, the largest average revision to beliefs comes with a round four signal of true following earlier round signals of one true and two false (row one in each frame). The patterns with the second largest revisions are T,F,F,F for round five in the partisan contest, and

² Limited to movers, subjects who changed their beliefs at least once in the contest.

Table A1: Models of learning political facts, by Round

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Partisan 1-2	Partisan 1-3	Partisan 1-4	Partisan 1-5	Partisan 2-2	Partisan 2-3	Partisan 2-4	Partisan 2-5	IQ-2	IQ-3	IQ-4	IQ-5
Logit prior (δ)	0.46** (0.07)	0.51** (0.05)	0.62** (0.05)	0.58** (0.05)	0.65** (0.05)	0.65** (0.04)	0.69** (0.05)	0.71** (0.04)	0.61** (0.06)	0.58** (0.06)	0.67** (0.04)	0.65** (0.05)
Signal (β)	0.71** (0.15)	1.11** (0.12)	0.64** (0.13)	0.69** (0.13)	0.68** (0.11)	0.66** (0.11)	0.67** (0.13)	0.61** (0.12)	0.70** (0.12)	0.80** (0.13)	0.67** (0.11)	0.36** (0.14)
Observations	941	954	953	948	962	972	969	965	956	968	970	969
R-squared	0.253	0.428	0.436	0.420	0.471	0.465	0.547	0.572	0.387	0.389	0.522	0.456

Robust standard errors in parentheses

** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on both variables.

Table A2: Models of learning political facts, by Question and Partisanship

VARIABLES	(1)	(2)		(3)		(4)		(5)		(6)		(7)		(8)		(9)		(10)		(11)		(12)		(13)	
	IQ	Kerry/Share Dems	Kerry/Share Reps	Obama/Income Dems	Obama/Income Reps	Abortion Dems	Abortion Reps	Bush/Unemploy Dems	Bush/Unemploy Reps	Reagan/Debt Dems	Reagan/Debt Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps	Romney/Share Dems	Romney/Share Reps
Logit prior (δ)	0.62** (0.03)	0.50** (0.08)	0.82** (0.05)	0.51** (0.08)	0.48** (0.11)	0.52** (0.05)	0.63** (0.08)	0.60** (0.05)	0.72** (0.08)	0.51** (0.08)	0.56** (0.08)	0.71** (0.04)	0.29* (0.14)												
Signal TRUE	0.80** (0.09)	0.96** (0.31)	0.41* (0.18)	1.05** (0.23)	1.09** (0.15)	0.047 (0.36)	0.93** (0.31)	1.19** (0.19)	1.00** (0.24)	1.08** (0.16)	1.28** (0.23)	0.54* (0.21)	0.85 (0.46)												
Signal FALSE	0.50** (0.10)	0.92** (0.14)	0.68** (0.17)	0.24 (0.27)	-0.70 (0.50)	0.73** (0.15)	0.72* (0.27)	0.55* (0.23)	0.54* (0.24)	0.32 (0.38)	0.62 (0.36)	0.45** (0.17)	0.21 (0.30)												
Observations	3,863	718	357	642	412	699	298	733	402	636	335	687	325												
R-squared	0.440	0.367	0.662	0.432	0.478	0.350	0.436	0.630	0.764	0.436	0.550	0.477	0.119												
Std. error of regression	2.33	2.53	1.83	2.30	2.34	2.45	2.30	1.98	1.56	2.35	2.11	2.43	3.16												
N subjects	988	184	94	164	106	180	77	187	103	162	86	179	87												

Robust standard errors in parentheses
** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on signals and priors. Sample sizes vary due to subjects who failed to enter a probability in individual rounds. Standard errors clustered on the subject-game. The first three partisan facts (columns 2 through 7) have True favoring the Republicans, the second three have True favoring the Democrats.

Table A3: Heterogeneity in updating on partisan-relevant facts

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Primary voter	Not primary	Liberal	Moderate	Conservative	Likes compromise	Dislikes compromise	Active	Not active
Logit prior (δ)	0.60** (0.03)	0.63** (0.04)	0.56** (0.04)	0.67** (0.03)	0.55** (0.04)	0.59** (0.06)	0.61** (0.04)	0.59** (0.04)	0.69** (0.04)	0.57** (0.03)
Signal (β)	0.59** (0.07)	0.58** (0.09)	0.63** (0.11)	0.54** (0.08)	0.80** (0.10)	0.47** (0.13)	0.72** (0.09)	0.50** (0.10)	0.59** (0.14)	0.60** (0.08)
Signal*Signal consistent (β_2)	0.45** (0.14)	0.54** (0.19)	0.28 (0.17)	0.50** (0.19)	0.25 (0.19)	0.60* (0.29)	0.41* (0.19)	0.49* (0.19)	0.23 (0.30)	0.50** (0.15)
Logit prior*Signal consistent (δ_2)	-0.078 (0.06)	-0.19* (0.09)	0.082 (0.06)	-0.11 (0.08)	0.00051 (0.07)	-0.12 (0.12)	-0.11 (0.08)	-0.054 (0.08)	-0.063 (0.08)	-0.073 (0.07)
Observations	6,138	3,158	2,964	3,147	1,643	1,340	3,089	3,033	1,284	4,846
R-squared	0.439	0.410	0.486	0.491	0.435	0.416	0.444	0.437	0.515	0.419
Std. error of regression	2.38	2.42	2.30	2.27	2.33	2.47	2.36	2.39	2.23	2.41
N subjects	804	414	388	407	217	179	402	400	170	633
Wald test on null $\delta = 1$	0	0	0	0	0	0	0	0	0	0
Wald test on null $\beta = 1$	2.7e-09	2.3e-06	0.00094	1.5e-07	0.11	0.00021	0.0025	1.1e-06	0.0053	1.9e-07

Robust standard errors in parentheses

** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on signals and priors. All columns limited to partisans (including leaners). Columns *Active* and *Not active* classify participants as those who report contacting or donating to political candidates or officials in the last two years, or report neither. Question wording: "In 2014, did you vote in your state's primary election to nominate candidates for Congress or state office?" ("I did not vote in the primary election in 2014"; "Yes, voted in a Democratic primary"; "Yes, voted in a Republican primary"; "Yes, voted in a nonpartisan or other primary"; "I usually vote in primary elections, but not in 2014."); "In general, would you describe your political views as" ("Very liberal"; "Liberal"; "Moderate"; "Conservative"; "Very conservative"). "Which statement comes closer to your view?" ("I like elected officials who make compromises with people they disagree with"; "I like elected officials who stick to their positions"). "During the past two years, have you contacted any political official or representative for any reason?" "During the past two years, did you donate money to a political candidate, campaign, party, or political organization?" Question wording from 2014 Pew Polarization Survey. Standard errors clustered at the subject-game.

Table A4: Tests for post-treatment bias in covariate measurement

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PID:	Democrat	Pure Ind	Republican	Ideo: Conservative	Ideo: Moderate	Ideo: Liberal	Likes Compromise	Active in Politics
Observations	989	989	989	989	989	989	987	988
R-squared	0.051	0.067	0.059	0.051	0.049	0.040	0.061	0.043
F-test p-value	0.57	0.11	0.29	0.58	0.66	0.91	0.23	0.86

Standard errors in parentheses

** p<0.01, * p<0.05

Note: Dependent variables are indicators for each covariate value used as a moderator in regression models in the paper. Because each was measured after the experiment, they may be influenced by treatment assignment. Each model regresses the indicator on treatment assignment: an indicator for which fact was assigned in the first partisan contest, an indicator for which fact was assigned in the second partisan contest, and the interaction of each of these indicators with the cumulative signals observed by the subject in that contest (coefficient estimates suppressed from the table). The F-test statistics evaluate whether the treatment assigned predicts the covariate response given beyond chance. None of the p-values reject a null hypothesis of no relationship between treatment assignment and covariate response at standard levels.

Table A5: Bayesian learning about political facts, Low scores on IQ quiz

VARIABLES	(1) Pooled	(2) Signal consistent	(3) Not consistent	(4) Pooled	(5) Dems/Reps only
Logit prior (δ)	0.53** (0.03)	0.46** (0.06)	0.52** (0.04)	0.52** (0.04)	0.51** (0.05)
Signal (β)	0.67** (0.07)	0.96** (0.15)	0.48** (0.10)	0.48** (0.10)	0.50** (0.12)
Signal*Signal consistent (β_2)				0.48* (0.19)	0.42 (0.23)
Logit prior*Signal consistent (δ_2)				-0.057 (0.07)	-0.039 (0.08)
Observations	2,875	1,196	1,620	2,816	2,183
R-squared	0.347	0.487	0.223	0.347	0.332
Std. error of regression	2.55	2.35	2.69	2.55	2.59
N subjects	376	337	368	374	290
Wald test on null $\delta = 1$	0	0	0	0	0
Wald test on null $\beta = 1$	9.2e-06	1	7.3e-07	7.0e-07	0.000038

Robust standard errors in parentheses
** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on both variables. Standard errors clustered on the subject-game.

Table A6: Bayesian learning about political facts, High scores on IQ quiz

VARIABLES	(1) Pooled	(2) Signal consistent	(3) Not consistent	(4) Pooled	(5) Dems/Reps only
Logit prior (δ)	0.73** (0.02)	0.66** (0.04)	0.72** (0.03)	0.72** (0.03)	0.73** (0.03)
Signal (β)	0.82** (0.06)	1.05** (0.12)	0.74** (0.06)	0.74** (0.06)	0.75** (0.07)
Signal*Signal consistent (β_2)				0.31* (0.13)	0.31* (0.14)
Logit prior*Signal consistent (δ_2)				-0.061 (0.05)	-0.078 (0.05)
Observations	3,809	1,692	2,069	3,761	3,199
R-squared	0.611	0.799	0.420	0.610	0.619
Std. error of regression	1.95	1.49	2.26	1.95	1.94
N subjects	487	454	470	487	413
Wald test on null $\delta = 1$	0	0	0	0	0
Wald test on null $\beta = 1$	0.0018	1	0.00011	0.00010	0.00037

Robust standard errors in parentheses

** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on both variables. Standard errors clustered on the subject-game.

T,F,T for round four in the IQ contest. Reading down the rows, the largest revisions almost always occur in cases with mixed rather than consistent signals. The first tipping point pattern for partisan contests is F,F with a mean revision of -21.7, but the remaining tipping point patterns all occur in the final rows of the table with the smallest revisions. For IQ contests, the first tipping point pattern is about one third down the table, T,T with a mean revision of 16.4, with the remaining in the bottom third of the table.

Because most of the tipping point patterns occur in the final rows of Table A7 while the largest revisions occur in response to patterns of mixed signals, a tipping point model of learning does not appear to be an effective explanation for learning behavior.

H Main tables without post-stratification weighting

In this section, I reproduce Tables 1, 2, and 3 without the Pew post-stratification weights (Tables A8, A9, and A10). The unweighted results are broadly consistent with the results using the post-stratification weights.

I Consequences of measurement error

One concern is that the estimate of cautious learning is due to measurement error in the instrument used to elicit beliefs. This would attenuate observed estimates and might lead to an estimate of caution for citizens who are actually learning as Bayesians. For example, I observe heaping of beliefs at integers that end in 0 or 5, suggesting that participants may be rounding beliefs. To assess the influence of this rounding, I took the set of responses and signals as observed and recalculated beliefs under the following rule: subjects learned as perfect Bayesians, but rounded posterior beliefs to the nearest integer ending in 0 or 5. For each contest, I took the first round beliefs as given (I did not apply the rounding rule), applied Bayes' Rule given the signal received in round 2 to generate posterior beliefs, rounded the round 2 beliefs to the nearest 0 or 5, and then used these rounded beliefs as the prior to round 3 beliefs. No additive noise other than the rounding rule was part of this simulation. I then ran the same model as in Table 2 on these alternative observations.³

When perfect Bayesian subjects apply a rounding rule, the model does estimate caution and learning that departs from Bayes' Rule, but not as much as with the observed data. In the observed data pooling both IQ and partisan contests from the first experiment, the estimates of δ and β are 0.62 [0.60, 0.63] and 0.70 [0.66, 0.74], 95 percent confidence intervals in brackets. With perfect Bayesians and a rounding rule, the estimates are 0.82 [0.81, 0.83] and 0.77 [0.75, 0.79]. These results do suggest some but not all of the caution in learning is due to a rounding rule applied by the subjects. Because in practice many subjects did give beliefs that were not rounded to 0 and 5, the simulation here where all apply the rule is likely overstating the influence of rounding. Note also that this rounding heuristic might also be applied in the real-world settings of learning, and so may also influence real-world parameters.

³ I kept the pattern of missing responses as observed in the data, though this leads to some additional missingness in the alternative rounding model due to respondents not bound by a missing prior belief.

Table A7: Revisions in beliefs by signal pattern and question type

Partisan

Round number	Lagged signals	Current signal	Mean revision	Median revision	Count
5	T,F,F	T	32.4	16.9	48
5	T,F,F	F	-24.0	-5.7	97
5	F,T,F	T	22.8	5.8	62
3	F	F	-21.7	-5.0	521
5	F,T,T	F	-20.8	0.0	61
4	F,T	T	20.5	5.0	141
5	F,T,F	F	-19.1	-9.0	89
4	T,T	F	-16.3	0.0	124
4	F,T	F	-16.0	-3.9	151
5	T,F,T	F	-15.6	-10.0	49
2	F	F	-15.5	-4.3	813
5	F,F,T	F	-15.1	0.0	93
5	T,T,F	F	14.0	0.0	38
5	F,F,F	T	13.4	0.0	94
3	F	T	12.7	0.0	292
2	T	T	12.3	5.0	694
3	F,F,T	T	11.8	0.0	425
5	F,F,T	T	11.7	0.0	56
3	T	F	-11.5	0.0	269
5	T,T,T	F	-10.7	0.0	73
4	F,F	T	8.7	0.0	149
5	T,F,T	T	8.2	0.0	75
4	T,F	T	8.2	5.0	124
5	F,T,T	T	8.0	0.0	80
4	T,F	F	-7.8	0.0	145
4	T,T	T	7.6	0.0	301
5	T,T,T	T	-6.7	0.0	228
4	F,F	F	-6.4	0.0	372
5	F,F,F	F	-3.3	0.0	278
5	T,T,F	T	0.8	0.0	86

IQ

Round number	Lagged signals	Current signal	Mean revision	Median revision	Count
5	F,T,F	T	50.0	49.0	28
4	T,F	F	-32.6	-5.0	57
5	F,T,T	F	-30.7	-10.0	33
3	F	T	26.0	17.7	134
5	T,F,F	T	25.3	8.6	24
4	F,T	F	-23.9	-3.4	68
5	T,T,F	F	-20.9	-15.0	19
5	F,F,T	T	17.1	0.5	31
3	T	T	16.4	9.0	207
2	F	F	-16.1	-9.0	372
3	F	F	-13.7	0.0	238
2	T	T	13.7	5.0	336
4	T,F	T	11.6	4.9	72
4	F,T	T	11.4	0.6	66
5	F,F,F	T	11.2	0.0	38
5	T,F,T	F	-10.9	0.0	22
4	T,T	F	-10.3	0.0	67
5	T,F,T	T	10.1	2.4	50
4	F,F	F	-9.1	0.0	157
5	F,F,F	F	9.0	0.0	119
5	T,T,T	F	-8.8	0.0	45
5	F,F,T	F	8.4	0.0	50
5	F,T,T	T	-7.9	0.0	33
4	T,T	T	7.1	0.0	140
5	T,T,T	T	-4.6	0.0	95
5	T,T,F	T	4.1	0.0	48
5	F,T,F	F	-3.8	0.0	40
5	T,F,F	F	-3.1	0.0	33
4	F,F	T	-2.0	0.0	81
3	T	F	-1.3	0.0	129

Note: Cells present the average and median revisions to subject beliefs in response to the current signal by lagged signal and type. Rows are sorted descending on the absolute value of the mean revision. Limited to subjects who changed their beliefs at least once in the contest.

Table A8: Prior and posterior beliefs by question and partisanship (unweighted)

Question	All respondent means		Democrat means		Republican means	
	Prior	Posterior	Prior	Posterior	Prior	Posterior
The rate at which American women aged 15-44 had legal abortions fell more between 1980 and 1988, while Ronald Reagan was president, than between 1992 and 2000, while Bill Clinton was president. (FALSE)	53.1	33.3	49.5	30.1	65.5	39.2
From 2009, when President Obama took office, to 2012, median household income adjusted for inflation in the United States fell by more than 4 percent. (TRUE)	57.9	76.2	52.7	75.4	63.3	75.8
The total public debt of the United States federal government more than doubled from quarter 2 in 1981 to quarter 1 in 1989 while Ronald Reagan was president. (TRUE)	59.9	71.4	61.3	70.7	53.1	68.4
In the 2004 Presidential Election, John Kerry was defeated by George W. Bush. In the nation as a whole, of all the votes cast for Kerry and Bush, Kerry won less than 48 percent. (FALSE)	64.2	36.2	63.4	34.5	65.0	37.1
From January 2001, when President Bush first took office, to January 2005, when President Bush started his second term in office, the civilian unemployment rate increased by more than 1 percentage point. (TRUE)	74.0	76.4	77.4	81.6	66.7	72.1
In the 2012 Presidential Election, Barack Obama defeated the Republican Mitt Romney. In the nation as a whole, of all the votes cast for Obama and Romney, Romney won less than 48 percent. (FALSE)	80.1	45.0	82.6	48.2	76.8	42.8
Your IQ quiz score is in the top half (respondents for which TRUE).	59.9	76.2	60.3	77.9	61.7	71.4
Your IQ quiz score is in the top half (respondents for which FALSE).	53.5	39.9	54.0	38.5	52.6	43.4

Note: Each cell presents the average elicited probabilistic belief that the statement is true among the participants who evaluated that fact. Prior beliefs are the beliefs offered in the first round prior to any signals. Posterior beliefs are the beliefs offered in the fifth round after receiving four signals. All subjects evaluated their score on the IQ quiz and evaluated two of the partisan statements drawn at random.

Table A9: Bayesian learning about political facts (unweighted)

VARIABLES	(1) Pooled	(2) Signal consistent	(3) Not consistent	(4) Pooled	(5) Dems/Reps only
Logit prior (δ)	0.65** (0.01)	0.57** (0.02)	0.63** (0.02)	0.63** (0.02)	0.64** (0.02)
Signal (β)	0.78** (0.03)	1.10** (0.06)	0.63** (0.04)	0.63** (0.04)	0.62** (0.04)
Signal*Signal consistent (β_2)				0.47** (0.07)	0.50** (0.07)
Logit prior*Signal consistent (δ_2)				-0.060* (0.03)	-0.070* (0.03)
Observations	7,664	3,294	4,227	7,521	6,138
R-squared	0.500	0.691	0.328	0.503	0.507
Std. error of regression	2.22	1.83	2.47	2.21	2.20
N subjects	990	902	958	988	804
Wald test on null $\delta = 1$	0	0	0	0	0
Wald test on null $\beta = 1$	0	0.21	0	0	0

Robust standard errors in parentheses

** p<0.01, * p<0.05

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on both variables. *Movers* excludes rounds from contests where the respondent did not change their beliefs in any round in response to signals. Standard errors clustered on the subject-game.

Table A10: Learning about relative quiz performance as benchmark (unweighted)

VARIABLES	(1) Pooled	(2) Pooled	(3) Dems/Reps only	(4) All contests
Logit prior (δ)	0.70** (0.02)	0.68** (0.02)	0.68** (0.03)	0.68** (0.02)
Signal (β)	0.65** (0.04)	0.52** (0.05)	0.52** (0.05)	0.52** (0.05)
Signal*Signal consistent (β_2)		0.40** (0.09)	0.43** (0.10)	0.40** (0.09)
Logit prior*Signal consistent (δ_2)		-0.043 (0.03)	-0.060 (0.04)	-0.043 (0.03)
Logit prior*Partisan fact				-0.047 (0.03)
Signal*Partisan fact				0.11 (0.06)
Partisan*Signal*Signal consistent				0.070 (0.11)
Partisan*Logit prior*Signal consistent				-0.016 (0.04)
Observations	3,863	3,808	3,104	11,329
R-squared	0.533	0.535	0.535	0.514
Std. error of regression	2.11	2.11	2.10	2.18
N subjects	988	969	791	990
Wald test on null $\delta = 1$	0	0	0	0
Wald test on null $\beta = 1$	0	0	0	0
Robust standard errors in parentheses				
** p<0.01, * p<0.05				

Note: Dependent variable is logit-beliefs that the statement is correct in that round for rounds 2 through 5. A perfect Bayesian would have coefficients of 1 on both variables. *Movers* excludes rounds from contests where the respondent did not change their beliefs in any round in response to signals. Standard errors clustered on the subject-game.

References

Lumley, Thomas. 2011. *survey: analysis of complex survey samples*. R package version 3.24-1.

R Development Core Team. 2015. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

J Experiment instructions

The final five pages present screen shots of the experimental instructions presented to participants along with the practice rounds each played.

This and the four pages that follow present the instructions for the crossover scoring method given to the subjects, along with the practice rounds of the contest they played.

Contest instructions

Instructions for the study

In the next part of this study, you are invited to participate in a game. We are going to ask you about three different statements of fact over the course of 15 rounds. These statements of fact may be true or may be false. In each round, you have the opportunity to win \$0.10, paid to you as a bonus. The contest works as follows:

We will present to you a statement of fact that may be true or false. You will indicate how

likely you believe the statement is true. Specifically, you will give a number from 0 to 100 that indicates how likely you believe the statement is to be true with 0 meaning false beyond any doubt and 100 meaning true beyond any doubt. For example, if you were almost entirely certain the statement is true, you might enter 99. If you were almost entirely certain the statement is false, you might enter 1. If you were totally uncertain about the truth of the statement, you should enter 50. You might believe it likely to be true but not be fully certain and enter 70. In each round, please enter how likely you believe the statement to be true.

We ask that you please not look up the answer to the question during the contest.

On the next page, we'll present how your your response determines whether or not you win that round.

Instructions for the study

Winning in each round of the game depends upon your response.

At the most basic level, in each round your task is to give your best guess about whether or not the statement is true. The contest is designed so that your chances of winning are highest if your response is an accurate reflection of how likely you believe the statement is true.

You will maximize your chance of the highest possible bonus by being as accurate as possible in each round.

Here is how your response generates a bonus in the game. You can skip these details if you are not interested in the underlying process. In each round, the computer will draw a random number from 0 to 100. Each number from 0 to 100 is equally likely to be drawn by the computer. We'll call this number Draw 1. How you win or lose that round of the contest depends on what number the computer draws for Draw 1 and your response:

1. If Draw 1 is less than your response, you win if the statement is true and do not win if the statement is false. For example, if you enter a response of 99, you are very likely to win if the statement is true and very likely to not win if the statement is false. The higher your response, the more likely you win if the statement is true. Similarly, the lower your response, the more likely you win if the statement is false.

2. If Draw 1 is greater than your response, then the computer will draw a second random number from 0 to 100. As before, each number from 0 to 100 is equally likely to be drawn by the computer. We'll call this random number Draw 2. If Draw 2 is less than Draw 1, then you win the round. If Draw 2 is greater than Draw 1, then you do not win the round.

The contest is designed so that you have the best chance for earning a bonus by being as accurate as possible with your response. The random numbers and payment calculations happen behind the scenes. You will not see the draws in any round.

Finally, you will have 20 seconds to submit your response on each screen.

Instructions for the study

We will ask your belief about whether the statement is true for each of 15 rounds of the contest. We will present the same statement more than once.

When we repeat a statement, the computer will provide you with a signal about the correct answer. The computer will present you a signal "TRUE" or "FALSE." Part of the contest is that three out of every four signals are correct, on average. That is, if the statement is true, the computer will signal "TRUE" three out of four times and "FALSE" one out of four times. If the statement is false, the computer will signal "FALSE" three out of four times and "TRUE" one out of four times. You will not know, however, whether or not each signal you see is correct.

We again emphasize that this is a NO DECEPTION study. The signals you receive will be correct three out of four times, on average.

You may use the information from the signal to change your response in that round from what you had said earlier.

When we give you more than one signal about the same question, we will store and present the signals for you so that you do not have to keep track in your head.

After you have completed the survey, we will calculate how many rounds you won and pay you your total bonus payment.

[new page]

Here is an example of what the contest will look like. Note: you are not being paid for these practice responses.

Factual statement:

* Please tell us how likely you believe this statement is true:

It rained (more than 0.00 inches of precipitation) in Santa Fe, New Mexico on July 7, 2004.

How likely you believe that the statement is true (for example, 1 if you believe it almost certainly false, 99 if you believe it almost certainly true, 50 if totally unsure):

Here is an example of what the contest will look like WHEN YOU RECEIVE A SIGNAL (timer not used here, but will be used in actual contest):

Factual statement:

* Please tell us how likely you believe this statement is true:

It rained (more than 0.00 inches of precipitation) in Santa Fe, New Mexico on July 7, 2004.

Last response:

* Your last response was ZZ.

Computer signal:

* The computer has produced a signal for you. Remember, three out of four times this signal will be accurate and one out of four times it will be inaccurate.

*** THE SIGNAL FROM THE COMPUTER IS "FALSE."**

How likely you believe that the statement is true (for example, 1 if you believe it almost certainly false, 99 if you believe it almost certainly true, 50 if totally unsure):

Here is another example of what the contest will look like WHEN YOU RECEIVE A SIGNAL (timer not used here, but will be used in actual contest):

Factual statement:

* Please tell us how likely you believe this statement is true:

It rained (more than 0.00 inches of precipitation) in Santa Fe, New Mexico on July 7, 2004.

Last response:

* Your last response was ZZ.

Previous signals:

* Your previous signal on this question was "FALSE."

Computer signal:

* The computer has produced a new signal for you. Remember, three out of four times this signal will be accurate and one out of four times it will be inaccurate.

*** THE SIGNAL FROM THE COMPUTER IS "TRUE."**

How likely you believe that the statement is true (for example, 1 if you believe it almost certainly false, 99 if you believe it almost certainly true, 50 if totally unsure):

Here is what will be going on "behind the scenes" after you submit your response in each round.

Your last belief that the statement, "It rained in Santa Fe, New Mexico on July 7, 2004" was true was ZZ.

According to Weather Underground (www.wunderground.com), there were 0.00 inches of precipitation in Santa Fe, New Mexico on July 7, 2004.

The correct answer is that the statement is FALSE.

If the random number drawn by the computer (Draw 1) was less than your response ZZ, because the statement is FALSE, you would have LOST.

If Draw 1 was greater than your response ZZ, the computer would draw a second number at random from 0 to 100 (Draw 2). If Draw 2 is less than Draw 1, you win, if Draw 2 is greater than Draw 1, you do not win. Again, **you will be most likely to win each round when you accurately report your belief.**

In the rest of the survey, you will not see the outcome of the random draws or your wins and losses. After the study, we will calculate your winnings based on your responses and random numbers drawn by the computer, and pay these to you as a bonus.

Now that you have seen an example, it is time to begin the contest. For this set of 15 questions, you will be paid \$0.10 as a bonus for each round you win, and \$0.00 for each round you lose. This bonus is in addition to your show-up fee, which you will be paid no matter the outcome.

Click here for a popup that briefly reviews contest instructions: [Review contest instructions](#)



We again ask that you please not look up any answers.

When you are ready to begin, please press "Next."