

No, you can't predict the outcome of Canada's 42nd federal election

Benjamin Gregory Carlisle

2015-10-25

Abstract

INTRODUCTION. We report a summary of lay-forecasts of Canada's 42nd federal election in 2015. **METHODS.** Subjective probability distributions were collected from volunteer forecasters using an online tool for collecting predictions of Canada's 42nd federal election. Predictors were instructed to construct subjective probability distributions by placing "chips" representing certainty into "bins" corresponding to ranges of seats that they predict will be won by each of the major five political parties. **RESULTS.** Seventy-six forecasts were collected. We report an overall Brier score of 0.0545, (95% confidence interval 0.0527–0.0562). The expectation for NDP seats was 88 (actual 44); for LPC seats was 134 (actual 184); for CPC seats was 116 (actual 99). **DISCUSSION.** The Liberal majority result was very surprising, as was the NDP's final seat count. The CPC result was significantly easier to forecast. Neither level of education nor level of party involvement nor time of forecast conferred any statistically significant advantage in forecasting the result. Supporters of a party were no better or worse at predicting their party's outcome than those who did not identify as such. Spending more time making a prediction was correlated with a better Brier score, but not significantly so.

Introduction

The ability to forecast future events and plan for the future is a uniquely human skill, and a difficult one to master. Through a reframing of one's memory known as Hindsight Bias,¹ an event that was actually very surprising *a priori* may be perceived to have been predictable by *a posteriori*. Hence, I have attempted to quantify explicitly a sample of lay-forecasts of seat counts for the major political parties in the 42nd Canadian federal election.

Methods

Subjective probability distributions were collected using an online collection tool originally developed by Carlisle et al for the STREAM forecasting project.² Predictors were shown a 20x20 grid, and allowed 40 "chips," each representing 2.5% certainty, and instructed to place them in bins corresponding to ranges of seats that they predict will be won by each of the major five political parties. For example, if a forecaster placed 10 "chips" in the column labelled "0–16," that would represent a forecast with 25% certainty that the party in question would receive between 0 and 16 seats in the House of Commons.

The forecast collection tool was available starting 10 days prior to the election, and was closed to new predictions at the time the first polls closed to voting. Forecasters were encouraged to make use of any external polls, articles or predictive models to which they had access, as this was meant to be a test of a person's ability to produce a well-calibrated forecast of the election, not a test of the person's memory.

Exclusions

Forecasts were excluded if they appeared to not represent a good-faith attempt to provide a forecast, or if there appeared to be a fundamental misunderstanding of the forecasting task.

¹https://en.wikipedia.org/wiki/Hindsight_bias

²<http://www.translationalethics.com/projects/forecast-study/>

Examples of excluded forecasts include:

- All cells left blank
- All cells filled in
- Alternating filled in and blank columns

In the case of predictions where less than 40 chips were used, the used chips were re-weighted such that the sum of the probabilities of the used chips always equals 1.

Forecasts of the Green Party and the Bloc Quebecois were not included in this analysis. These predictions were excluded because a grid of 20 ranges of approximately 17 seats per column does not allow for a sufficiently granular prediction space. Many forecasters placed 20 chips in the first column (the maximum allowed in a single column) and left the second column blank, while others filled the first two columns with 20 chips each, to use the allotted 40 chips. It was impossible to discern between those who would have used only the first column, had the option been explicitly given, and those who honestly meant to give a prediction that is twice as wide, and so these were excluded for simplicity.

Statistical methods

Summary probability distributions, g over discrete variable x (number of seats in the House of Commons) were calculated using the following formula, where N is the total number of forecasts, w_t is a weighting factor for the forecast in question and $f_t(x)$ is the t -th subjective probability function over x . Unless otherwise specified, $w_t = \frac{1}{N}$

$$g(x) = \sum_{t=1}^N w_t f_t(x)$$

The expectation of a subjective probability distribution was calculated as follows, where E is the expectation of some discrete variable x (here, the number of seats in the House of Commons), N is the number of forecasts, $f_t(x)$ is the value of the probability function at x for the forecast in question.

$$E = \frac{1}{N} \sum_{t=1}^N x f_t(x)$$

The calibration of a forecast was assessed using Brier scores³ (S_B), which were calculated using the following formula, in which N is the number of forecasts, f_t is the value of the probability function for the forecast in question and o_t is the outcome for the forecast (0 or 1).

$$S_B = \frac{1}{N} \sum_{t=1}^N (f_t - o_t)^2$$

Differences between groups and hypothesis-testing was done using a two-tailed t -test using R version 3.2.2⁴ and the *scoring* package.⁵

³https://en.wikipedia.org/wiki/Brier_score

⁴<https://www.r-project.org/>

⁵<https://cran.r-project.org/web/packages/scoring/scoring.pdf>

Results


In total, 87 forecasts were collected between 2015 October 10 and 2015 October 19 at 20:30, Newfoundland time. After applying the above exclusion criteria, 76 were eligible for inclusion.

Sample demographics		Total n=76
Preferred party		
	NDP	26 (34%)
	LPC	14 (18%)
	GPC	3 (4%)
	CPC	2 (3%)
	Other	1 (1%)
	Not stated	30 (39%)
Party affiliation		
	NDP	9 (12%)
	LPC	3 (4%)
	GPC	2 (3%)
	CPC	2 (3%)
	Other	2 (3%)
	Not stated	58 (76%)
Religious affiliation		
	Christian	10 (13%)
	Atheist	13 (17%)
	Agnostic	7 (9%)
	Other	3 (4%)
	Not stated	43 (57%)
Primary news source		
	CBC	24 (32%)
	Internet/Twitter/Reddit	7 (9%)
	Globe and Mail	5 (7%)
	La Presse	2 (3%)
	Mixed	3 (4%)
	CTV	1 (1%)
	Global	1 (1%)
	Le Devoir	1 (1%)
	Toronto Star	1 (1%)
	Not stated	31 (41%)


Combined subjective probability distributions


The following are the combined subjective probability distributions, weighted equally among all forecasters. The coloured column represents the range of possible outcomes that includes the actual number of seats won by the party in question, namely 44, 184 and 99 seats for the NDP, LPC and CPC respectively.

NDP



LPC





The expectation as calculated from our forecaster’s subjective probability distributions and actual results for the combined subjective probability distributions for the NDP, LPC and CPC are reported in the table below. The final seat projection from CBC’s Poll Tracker and ThreeHundredEight.com are also provided for reference.

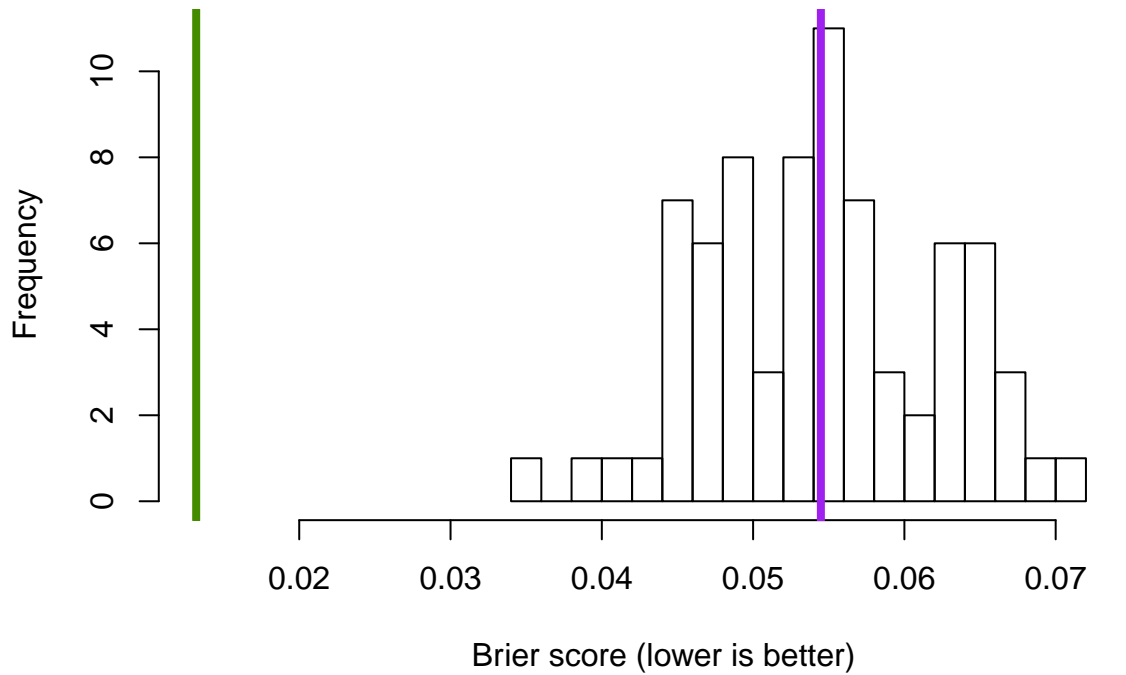
Party	Expectation	CBC ⁶	308 ⁷	Actual outcome	Difference
NDP	88	66	56	44	-44
LPC	134	146	154	184	50
CPC	116	118	120	99	-17

The LPC result was the most surprising, followed closely by the NDP result. The number of seats won by the CPC was the least surprising.

Brier scores

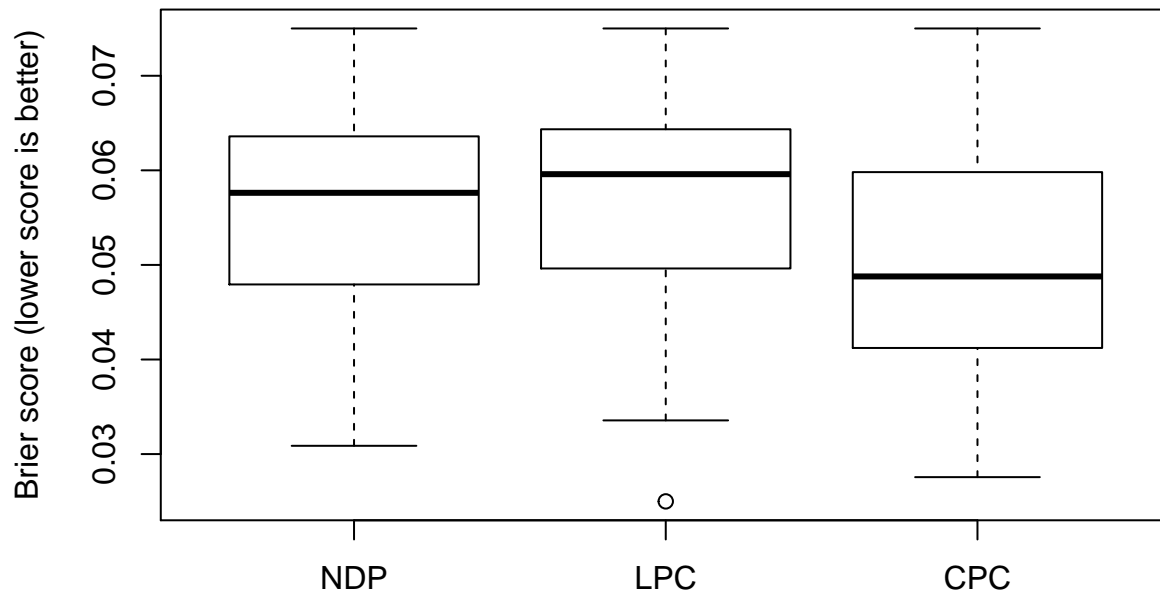
Brier scores were calculated for provided subjective probability distributions for forecasts of the number of seats to be won by the NDP, LPC and CPC, and an overall Brier score was calculated for each predictor as well, by taking the mean of the scores for each party’s seat prediction. The mean Brier score for all predictors was 0.0545, (95% confidence interval 0.0527–0.0562). Forecaster’s Brier scores are plotted in a histogram below, with the mean as a vertical purple line. For reference, because of the constraints placed on the forecasters (20 bins with a maximum of 50% certainty in one bin and a requirement that all chips be used), the worst possible prediction would receive a score of 0.075 (marked with a red vertical line) and the best possible prediction would receive a score of 0.013 (marked with a green vertical line).

Overall mean Brier scores for all predictions



The following is a box plot for Brier scores of forecasts of the number of seats in the House of Commons for the NDP, LPC and CPC. The mean Brier score for forecasts of the NDP was 0.0558 (95% confidence interval 0.0536–0.0581), 0.0571 (0.0547–0.0594) for the LPC and 0.0506 (0.0479–0.0532) for the CPC. The number of seats won by the LPC was more surprising than the second most surprising result, the seat count for the NDP, but not significantly so. The result for the CPC was least surprising, and significantly so.

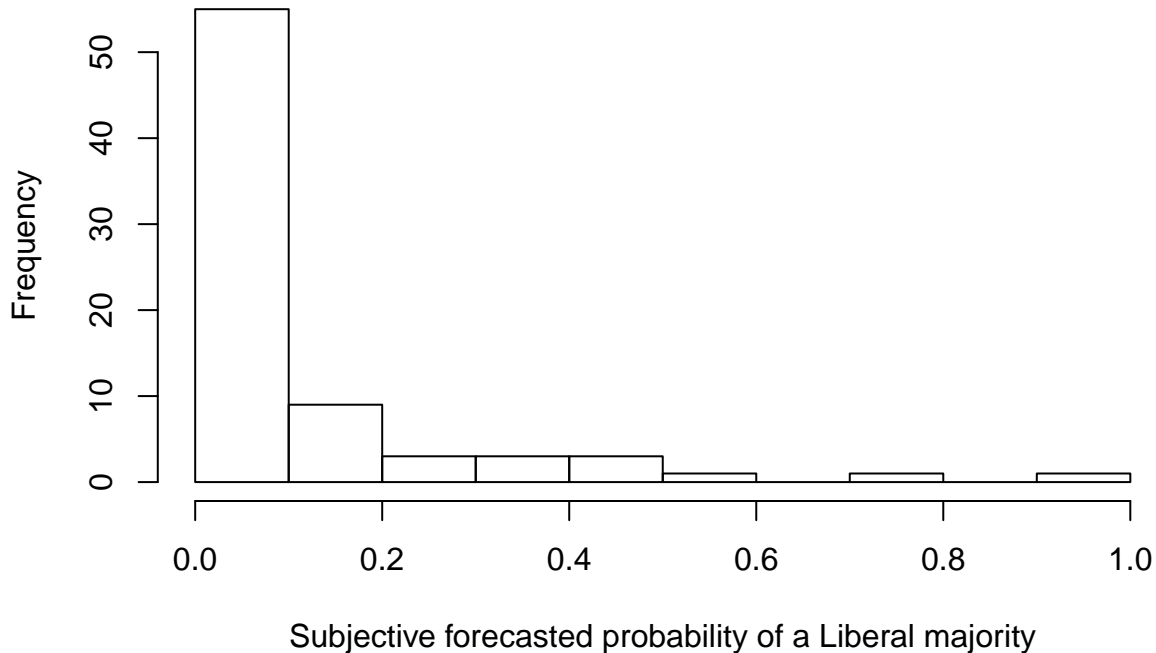
Brier scores for forecasts of seat numbers by party



Forecasts of a Liberal majority

We summed the subjective probabilities assigned to the ranges of seat outcomes that corresponded to a LPC majority (namely, all columns representing outcomes of 170 seats or more). A histogram of the combined subjective probability of a Liberal majority outcome has been plotted below. Only 3 forecasters out of 76 forecasted Liberal majority with a certainty greater than 50%. In total, 38 out of 76 forecasters judged that there was a 0% probability of electoral outcomes that included a Liberal majority.

Combined forecasted probability of a LPC majority



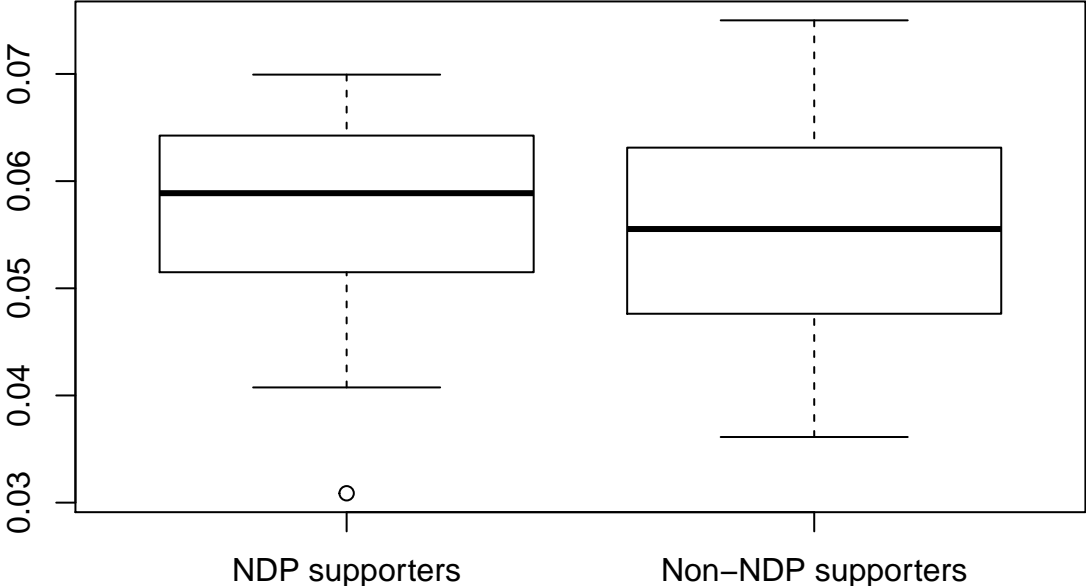
Comparison of Brier scores for supporters and non-supporters of each of the political parties

The mean Brier score for predictions of the NDP seat count among NDP supporters was 0.0568 (95% confidence interval 0.0529–0.0607), while the mean Brier score for predictions of the NDP seat count among those who did not identify as NDP supporters was 0.0553 (0.0525–0.0582).

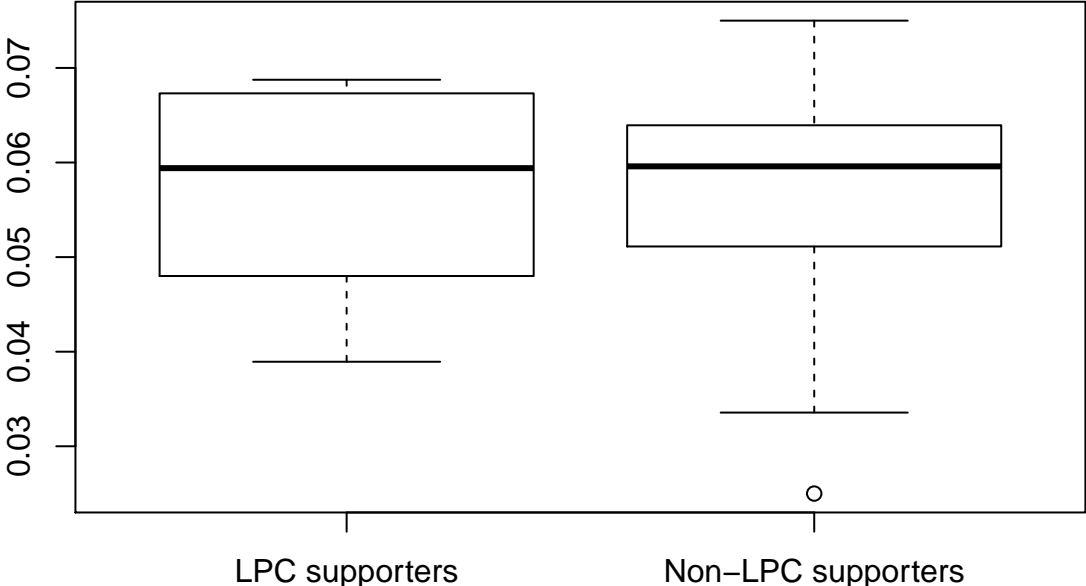
The mean Brier score for predictions of the LPC seat count among LPC supporters was 0.0567 (0.0505–0.0629), while the mean Brier score for predictions of the LPC seat count among those who did not identify as LPC supporters was 0.0571 (0.0545–0.0597).

The mean Brier score for predictions of the CPC seat count among CPC supporters was 0.063 (0.000–0.127), while the mean Brier score for predictions of the CPC seat count among those who did not identify as CPC supporters was 0.0502 (0.0476–0.0529).


Brier scores for predictions of seat counts for the NDP



Brier scores for predictions of seat counts for the LPC



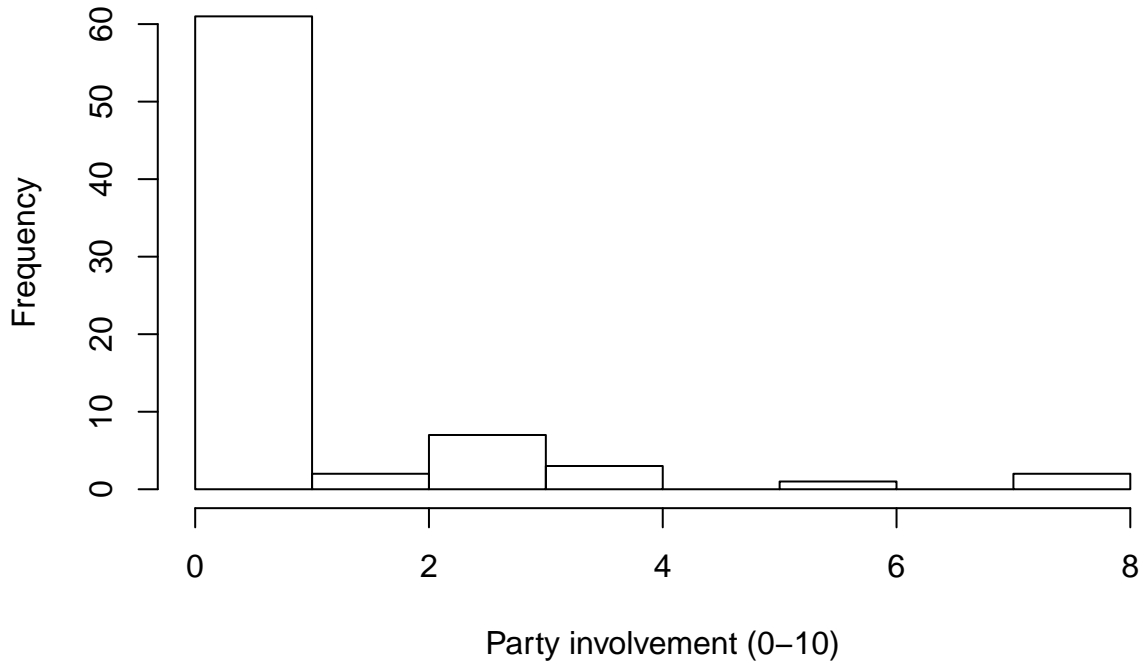
Brier scores for predictions of seat counts for the CPC



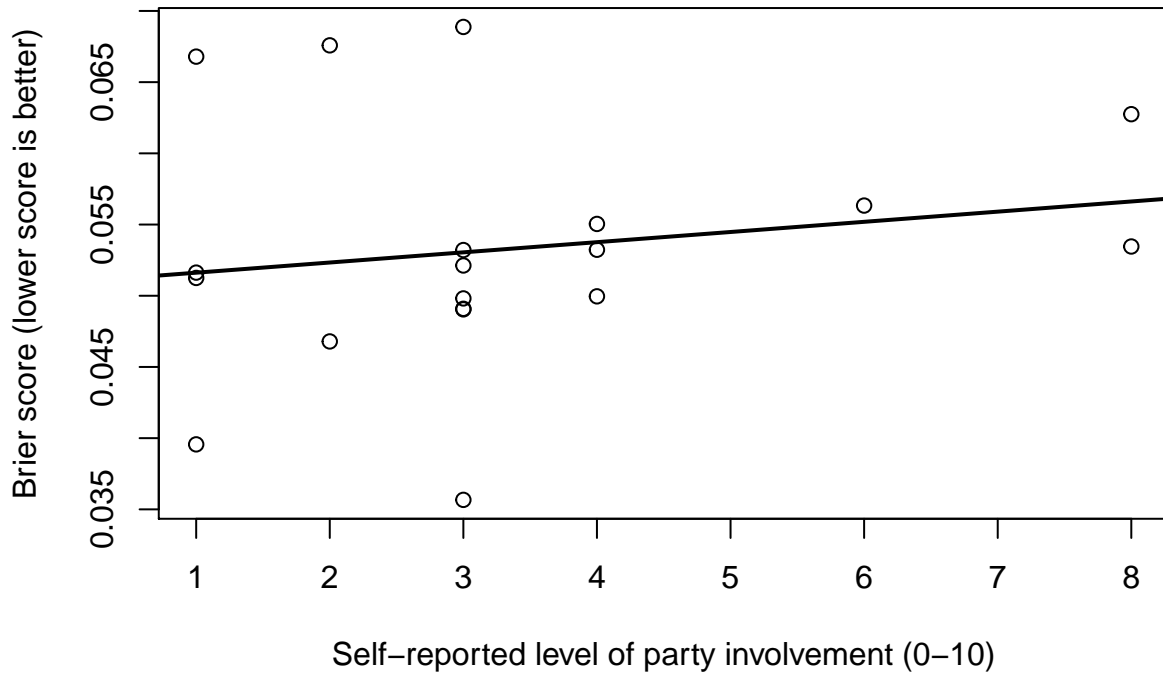
Party involvement

To test the relationship between self-reported level of party involvement (0–10, where 0 is “none” and 10 is “party leader”) and Brier score using a linear regression model. We considered only forecasters whose level of party involvement was greater than 0 ($n = 19$), since a level of 0 could mean either “no involvement” or “decline to answer.” There was a trend toward worse predictions among forecasters with higher levels of party involvement, but this was not statistically significant ($p = 0.48$).

Self-reported level of party involvement



Brier score vs party involvement

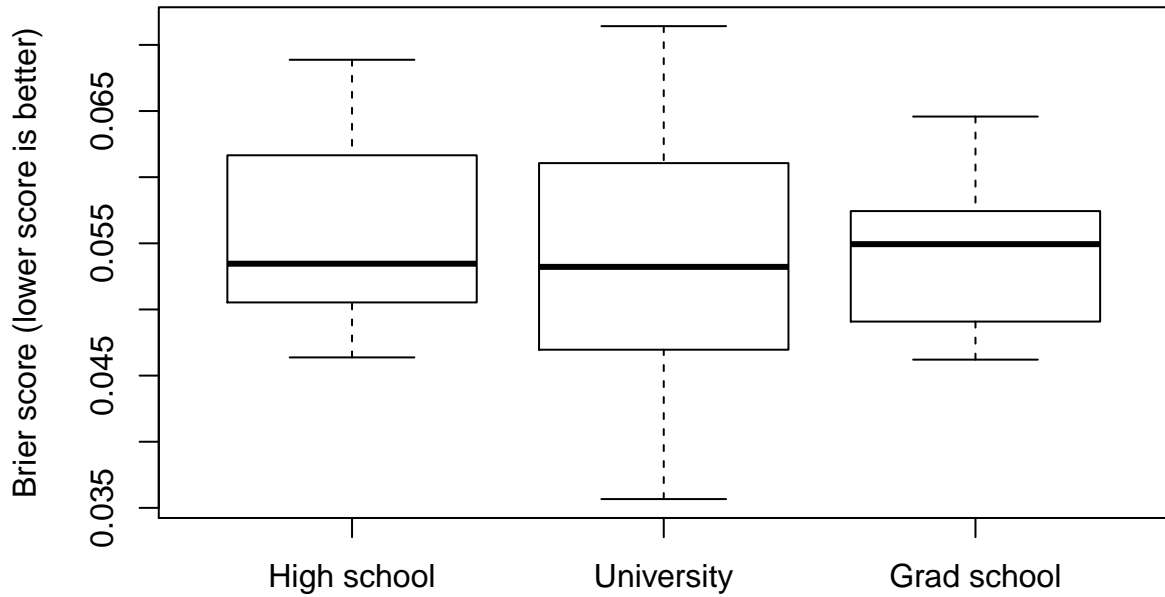


Level of education

Among the forecasters who reported their level of education ($n = 55$), there was no significant difference between the mean Brier scores for those whose level of education was high school (0.056, 95% confidence

interval 0.0507–0.0613), university (0.0534, 0.0502–0.0566) or grad school (0.0541, 0.0507–0.0574).

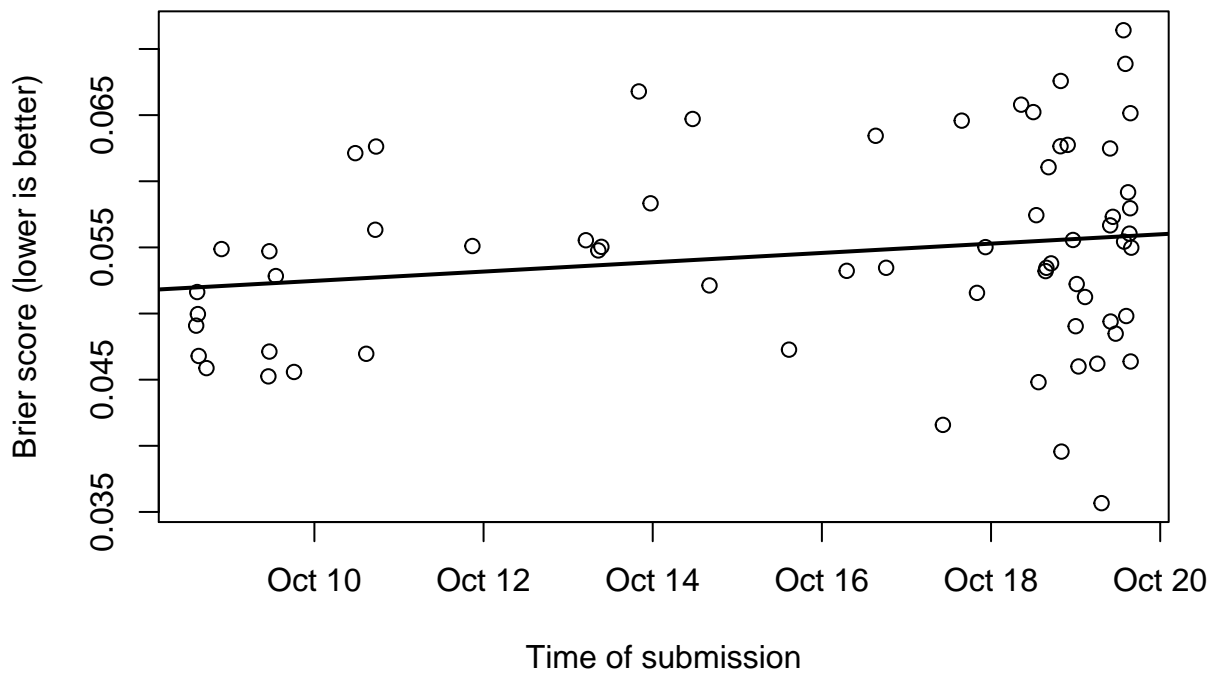
Brier score vs level of education



Time of submission

We regressed time of forecast submission against overall Brier score. There was a small and statistically non-significant trend toward worse predictions over time ($p = 0.13$).

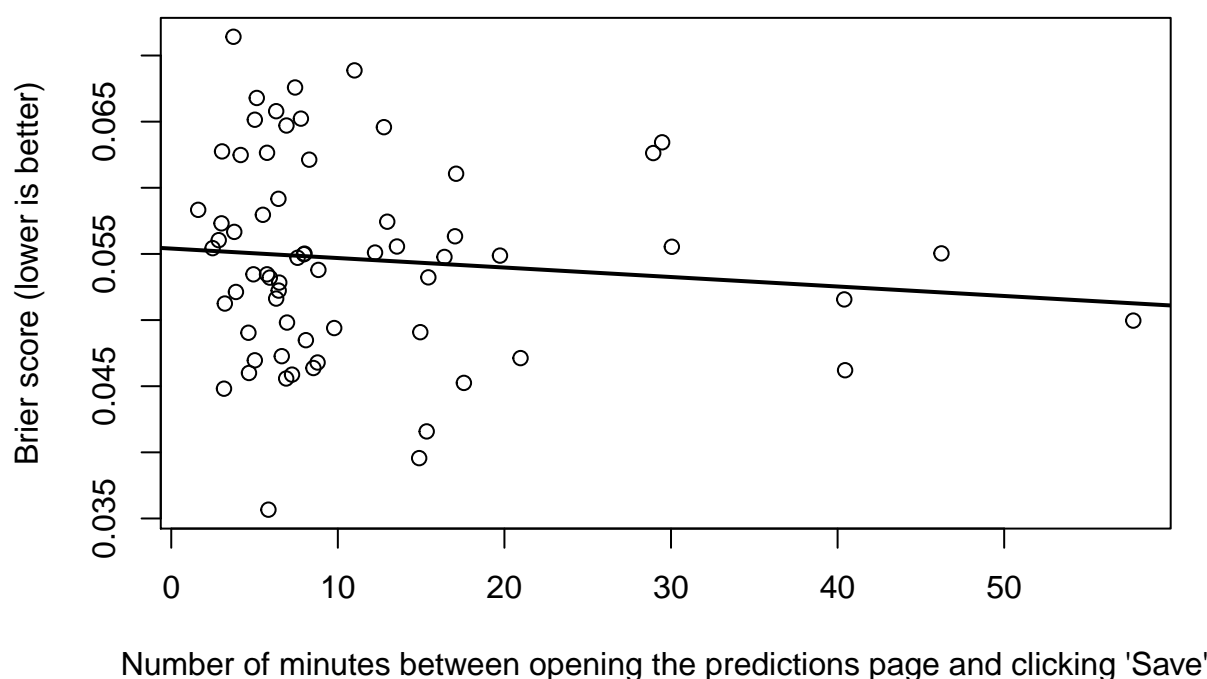
Time of submission vs Brier score



Time spent making a prediction

We measured the amount of time spent making a prediction by subtracting the time at which a forecaster first opened the forecasting page, and when that forecaster clicked “Save.” This measurement of time spent making a prediction was regressed the forecaster’s Brier score. There was a trend toward making better predictions, the longer a predictor spent, however this trend was also not significant ($p = 0.4$).

Time spent making a prediction vs Brier score



Discussion

The question I posed to potential forecasters was, “Can you predict the results of Canada’s 42nd federal election?” The answer appears to be: “No.”

In brief, the result of this election was very surprising, and it was consistently surprising across a number of variables that we considered. Our sample underestimated the LPC’s result by 50, overestimated the NDP’s result by 44 and overestimated the CPC’s result by 17.

The LPC majority government outcome was especially surprising. Nearly all (96%) of our forecasters bet against the outcome of a Liberal majority, giving it a probability of 50% or less. Half (50%) of those sampled ruled that outcome out entirely, assigning it a 0% chance of occurring. The number of seats won by the NDP was also difficult to predict, and Brier scores of LPC vs NDP seat-forecasts were not significantly different ($p > 0.05$). The CPC result was significantly less surprising ($p < 0.05$), although still difficult to predict.

The outcome of this election seems to be equally surprising across different demographics. Our data give us no reason to think that those with higher involvement in a party are better able to predict the outcome of an election. Indeed, the prediction with the best score in our sample came from a forecaster who self-reported a level of 3 on a scale of party involvement from 0 (“none”) to 10 (party leader). Similarly, level of educational achievement was no predictor of forecasting success.

Contrary to expectation, making a prediction closer to the time when the voting booths closed gave no advantage in one’s Brier scores. Spending more time making a prediction trended toward better predictions,

but neither of these trends were statistically significant ($p > 0.05$).

For the NDP, LPC and CPC, neither those who identified as preferring a party, nor those who did not identify as preferring that party were significantly different in their ability to forecast the outcome of that party. The mean Brier score for CPC supporters, 0.063, falls outside the 95% confidence interval for the mean Brier score for non-CPC supporters, 0.0502 (0.0476–0.0529), suggesting that CPC supporters were significantly worse at predicting the seat count for the CPC, but this is an artefact of the comparatively small sample of forecasters who identified as preferring the CPC ($n = 2$). The Brier scores of the two samples are not significantly different ($p = 0.43$).

Polling numbers prior to the election were representative of the proportions of votes in the general population, but the result was not foreseen by CBC's Poll Tracker, which listed a low of 51 for the NDP, a high of 161 for the LPC and a low of 100 for the CPC.⁸

Going into the election, most forecasters expected a seat count that would represent a LPC minority government; some even projected a minority CPC government. It is possible that the spectre of the LPC not being certain to defeat another CPC government, combined with the projected poor performance of the NDP caused some NDP supporters who otherwise would not have otherwise voted for the LPC to do so, giving the LPC an unexpected majority government.

⁸<http://www.cbc.ca/news/politics/canada-election-2015-grenier-projections-oct22-1.3282596>