CJPS   RCSP

Canadian Journal of Political Science

Revue canadienne de science politique

CAMBRIDGE
UNIVERSITY PRESS

46:1   March 2013 mars

Articles

The *Canadian Journal of Political Science*, whose policy is set by the Editorial Board, is committed to the publication of scholarly research in all fields of political science. Manuscripts are assessed on the basis of their contribution to the discipline and by the authors' adherence to the highest scholarly standards.

La *Revue canadienne de science politique*, dont la politique est mise en place par le comité de rédaction, s'engage à publier des travaux de recherche académique dans toutes les orientations de la science politique. Les manuscrits sont évalués sur la base de leur contribution à l'avancement des connaissances dans la discipline et selon la conformité des textes aux exigences académiques les plus élevées.

Canadian Journal of Political Science / Revue canadienne de science politique

Membership in either the Canadian Political Science Association and/or the Société québécoise de science politique includes a subscription to the *Journal*. Institutional subscriptions (which do not include membership privileges) are also available. The *Journal* may be obtained in paper and online. / La cotisation versée à l'Association canadienne de science politique et/ou à la Société québécoise de science politique donne droit à un abonnement à *La Revue*. Des abonnements déstinés aux institutions (qui ne donnent pas droit aux privilèges réservés aux membres) sont également offerts. Vous pouvez recevoir *La Revue* sous la forme d'une revue et en ligne..

Choice of Associations/ Choix d'associations	Students and any other person earning under CDN $40,000 annually/ Étudiants et toute autre personne dont le revenu annuel ne dépasse pas 40 000 $ CDN	All other members/ Tous les autres membres
CPSA only / ACSP seulement	CDN $50	CDN $150
SQSP seulement / only*	CDN $50	CDN $140
Joint/conjoint CPSA/ACSP-SQSP*	CDN $60	CDN $170

* These categories of membership include a subscription to the SQSP's journal *Politique et Sociétés*. / Chaque catégorie d'adhésion comprend un abonnement à la revue de la SQSP, *Politique et Sociétés*.

Return postage guaranteed / Port de retour garanti

Articles

Follow the Pollsters: Inaccuracies in Media Coverage of the Horse-race during the 2008 Canadian Election

FRANÇOIS PÉTRY   *Université Laval*
FRÉDÉRICK BASTIEN   *Université de Montréal*

The media's faith in polls ... is symbolized by a little number that accompanies every major poll... It encapsulates all the uncertainty about a poll's veracity and rolls it into a tiny ball that can be ignored. This number is arguably the most misunderstood and abused mathematical concept that journalists have gotten their fingers on: the margin of error.

Charles Seife, *Proofiness* (2010)

There has been a marked increase in the use of media-sponsored polls during recent Canadian federal election campaigns. More than 200 national surveys on voter intentions were issued by major polling organizations during the 37-day federal election campaign of 2008 (including daily tracking polls). By comparison, Johnston and colleagues (1992: 121) mention 22 national polls on voter intentions reported by the media during the 51-day federal election campaign of 1988. The recent proliferation of media-sponsored election polls has contributed to exacerbate

François Pétry, Département de science politique, Université Laval, Cité Universitaire, Québec QC, G1K 7P4. Email: Francois.petry@pol.ulval.ca
Frédérick Bastien, Département de science politique, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal QC, H3C 3J7.

horse-race reporting—the tendency to interpret elections in terms of how the candidates are faring in the polls relative to one another—to the point of ignoring the issues (Andersen, 2000, Farnsworth et al., 2007, Trimble and Sampert, 2004).

Considering the tight focus on the horse race during Canadian election campaigns, an important question to ask is whether pollsters and journalists interpret the horse race accurately. Most citizens do not have the expertise to understand the impact of statistical details on the meaning of election poll results. They rely on pollsters and journalists to interpret the horse race, and their interpretation will be accurate only if the interpretation by pollsters and journalists is itself accurate. We know that American journalists fail sometimes to interpret the horse race accurately. Using anecdotal evidence from recent US presidential campaigns, Patterson claims that journalists tend "to report small changes from one poll to the next as a manifestation of actual change in voters' preferences rather than as a reflection of survey error" (2005: 717). Larson (2003) assessed directly the accuracy of TV journalists' interpretations of the horse race during the 2000 US presidential campaign and found almost half of these interpretations were inaccurate. The main mistake was journalists' declaring one candidate "ahead" when his advantage was within the margin of error.

Patterson's "images of the game" depiction (1993, 2005) offers a convincing argument as to why journalists overestimate the extent by which one party leads another and the change in party support over time. He claims that journalists frame election coverage within four narratives: a party is leading or trailing or gaining ground or losing ground. The "leading" and "trailing" narratives prompt journalists to overestimate the gap between parties in order to identify clearly a "frontrunner" and a "likely loser" (1993: 120–23). The "gaining ground" and "losing ground" narratives lead journalists to overestimate actual changes in vote intentions. There is considerable drama to write about when a party suddenly increases its share of voter intention, or faces a severe decline in its support. Thus, journalists may try to find a new story among poll figures despite the stability of public opinion.

This paper follows the path opened by Larson (2003) and investigates the accuracy of media interpretations of the horse race during the 2008 Canadian election. We adapt Larson's methodology to a multiparty system and widen the scope of her analysis in three ways. First, we use the standard method of hypothesis testing for proportions which gives a truer picture of the horse race than the method used previously in the literature. Second, we do not simply replicate previous analyses of synchronous situations (who is ahead? who is behind?), we also break new empirical ground in analyzing diachronic situations (who is up? who is down?). This allows us to show an even higher level of inaccuracy in

Abstract. We identify frequent inaccuracies in journalistic interpretations of the horse race ("who is ahead?") and of change over time ("who is gaining?") in poll reports during the Canadian election of 2008. We test two explanations. The "mistaken mindset" hypothesis holds that journalists exaggerate the horse race because they systematically miscalculate the margin of error. The "follow-the-pollster" hypothesis holds that journalists follow the horse-race interpretations that they find in pollsters' reports. We find strong support for the "follow-the-pollster" hypothesis in the data and in interviews with pollsters and journalists and conclude that pollsters' reports should be a key element to consider in any attempt to improve the level of accuracy in media reports of the horse race.

Résumé. Les journalistes ont souvent commis des erreurs d'interprétation de la marge d'erreur dans les résultats de sondages pendant la campagne électorale canadienne de 2008. Cela les a conduits à surestimer l'avance du parti gagnant et le changement dans le score d'un parti entre deux sondages successifs. Comment expliquer ces erreurs fréquentes? Une première hypothèse affirme que cette surestimation provient du fait que les journalistes se trompent systématiquement dans le calcul de la marge d'erreur. Selon une deuxième hypothèse, les journalistes connaissent tellement mal la marge d'erreur qu'ils se fient à l'interprétation qu'en font les maisons de sondage. Les données empiriques et les réponses aux questions d'entretiens soutiennent la deuxième hypothèse. Nous en concluons que pour mieux porter fruits, les efforts pour améliorer l'interprétation de la marge d'erreur devraient cibler les maisons de sondage autant que les journalistes.

diachronic interpretations. Third, and most important perhaps, a comparison of horse-race interpretations by journalists and by pollsters shows that inaccurate interpretations by the media are causally related to inaccurate interpretations in pollsters' reports.

Hypotheses

Assuming that journalists don't knowingly tell lies to the public, how can we explain the frequent occurrence of inaccurate horse-race statements in the media? This paper examines two testable hypotheses that make distinct but compatible predictions of when inaccuracies occur and when they don't. The "mistaken mindset" hypothesis holds that journalists systematically underestimate the margin of error when they interpret the horse race, and this leads them to inaccurately state that one party is leading when this is not supported statistically. Larson (2003) proposes a clever operational definition of the "mistaken mindset" in synchronous horse races (see also Zukin, 2004). She claims that American broadcast journalists "think that the margin of error is the gap that needs to be between the candidates' results for there to be a significant difference rather than a number that needs to be added and subtracted to each candidate's percentage."[1] Accordingly, they accurately state that no party is leading as long as the gap between party percentages is smaller than the margin of error that they report, but they switch to declaring that one party leads the other as soon as the gap becomes larger than the reported

margin of error, and this is statistically inaccurate. The statement that one party leads the other becomes accurate once the gap clears the "true" margin of error. In formal terms, the relationship between the margin of error overlap and the likelihood of accuracy is expected to take the form of a rectangular step function.

Since journalists get their information about the margin of error from pollsters' reports, the "mistaken mindset" hypothesis poses the question whether pollsters interpret the margin of error accurately. And if so, why don't journalists adjust their interpretations accordingly? The "follow-the-poster" hypothesis addresses this question. This hypothesis holds that in order to be able to use polls so extensively in their horse-race cover-age of election campaigns, journalists must have faith in the poll results they report. That faith is symbolized by the margin of error (Seife, 2010). But because they are unsure about correctly interpreting the margin of error on their own, journalists will defer to the interpretation that they find in pollsters' reports. The hypothesis also holds that, unbeknownst to journalists, pollsters sometimes make inaccurate interpretations of the horse race in their reports to journalists. The accuracy of media reports tends to co-vary with the accuracy of pollsters' interpretations.

Whether polling organizations accurately report the margin of error has never been directly investigated in the scholarly literature on election polls. But, there are theoretical reasons to expect that pollsters may over-state the horse race in their reports to journalists in a manner consistent with Patterson's "image of the game," in spite of efforts by some mem-bers of the polling community to maintain high technical standards. The relationship between the media and the polling industries creates the motive for some pollsters to be at variance with best practices in poll qual-ity control. As Rosenstiel (2005) argues, recent changes in the use of polls by the media—increasingly unreliable polling practices, more outlets com-peting for fewer resources, growing reliance on daily tracking polls, to name a few—have provoked similar changes in the polling industry. To reinforce his point, Rosenstiel quotes Cliff Zukin, past president of the American Association for Public Opinion Research (AAPOR) as saying, "A lot of pollsters are not following best practices and they know it, includ-ing pollsters at some of the biggest news organizations, because there is a demand for it (Rosenstiel, 2005: 714)" Furthermore, the margin of error is a complex mathematical notion which few people clearly understand; this may create the opportunity for some pollsters to overstate the horse race without risking a penalty for doing so. If, as Rosenstiel (2005) argues, pollsters do not always follow best practices with respect to easy-to-understand standards of poll quality (for example, correct sample size, neu-trally worded questions, accurate estimate of undecided voters), it should come to no surprise that they also fail to follow best practices with respect to the hard-to-understand margin of error.

Methodology

To test the hypotheses, we analyze the content of election poll reports released by polling and media organizations during the Canadian federal election campaign of 2008. We are interested in two basic pieces of information: First, we calculate the difference between party scores and whether it is statistically significant at 95 per cent. Second, we assess whether pollsters and journalists provide an accurate interpretation of these data. We analyze the content of reports published by Angus Reid, Compas, CROP, EKOS, Environics, GPS, Harris Decima, Ipsos Reid, Léger Marketing, Nanos, Segma, and Strategic Counsel from September 6, the day the writ of election was issued, through October 13, 2008, the day before polling day.[2] Most reports were downloaded from polling organizations' websites; some were obtained by special request.[3] Next we look at how pollsters' reports translate in media stories. Our media sample includes all election poll reports broadcast on the national evening news of four television networks (CBC's *The National*, CTV *National News*, Radio-Canada's *Le Téléjournal* and *Le TVA, édition 22h*), and all those published in two English-language newspapers distributed across Canada (*The Globe and Mail* and *National Post*) and two French-language newspapers available all over Quebec (*La Presse* and *Le Devoir*) from September 7 to October 14. This allows us to take into consideration the media which have the widest geographical diffusion—and whose audiences are among the largest (*Le Devoir* excepted) across the country in both official languages. To avoid data repetition (for example, the same result being published more than once in the same medium) we only consider reports that are published for the first time by each news outlet. In most cases, the media report consists of one story featuring results from one election poll. Two distinct election polls reported within the same story give two separate media reports.

Assessing Accuracy of Horse-Race Interpretations

Synchronous horse-race interpretations are statements indicating either that a party leads or trails another or that there is no party leading. The parties are the Liberal Party (Lib), the Conservative Party (Con), the New Democratic Party (NDP), and the Bloc Québécois (Bloc). In a media report featuring vote intentions for all these parties, we get a maximum of six possible interpretations (Con-Lib, Con-NDP, Con-Bloc, Lib-NDP, Lib-Bloc, and NDP-Bloc). To be included in the analysis, an interpretation must provide information enabling research analysts to verify whether a party is leading or trailing another (or that no party is leading) and by what magnitude. This narrow definition is intended to avoid sports metaphors or war analogies which are often considered in the horse-race lit-

erature but the accuracy of which is sometime difficult to determine (Trimble and Sampert, 2004).

To assess the accuracy of a synchronous horse-race interpretation, we need to calculate whether two party scores in the same poll are statistically different from each other. This can be done using the confidence interval for a single party score in a poll, given by the equation:

$$\pm 1.96 \sqrt{\frac{p(1-p)}{n}} \qquad (1)$$

where p is the party percentage score and n is the sample size. Although equation (1) is really intended to calculate the confidence interval for a single party score, pollsters use it to report the confidence interval for the difference between two party scores in the same poll, which is equal to twice the margin of error they report (Gawiser and Witt, 1994).[4] A fixed value (.5) is given to party p's score so the score of the other party $(1 - p)$ is also .5. With $n = 1,000$, the reported margin of error is $\pm.031$; the difference between two party scores must therefore be .062 to clear the 95 per cent margin of error.

From a statistician's viewpoint, the requirement that the gap between party scores be equal to or larger than twice the reported margin of error is too "conservative" (Schenker and Gentleman, 2001). First, the margin of error reported by pollsters is the "maximum" one, the margin of error that would materialize if the two party scores were equal. When party scores are not equal, as is usually the case, the observed margin of error is smaller than the theoretical maximum, and the more the scores differ from .5, the smaller the margin of error. Second, the margin of error reported by pollsters is calculated as if there were only two parties and no undecided voters. Taking into account third parties and undecided voters necessarily make the numerator smaller relative to the unchanged denominator, hence, the true margin of error is smaller than the reported margin of error.

The "true" 95 per cent margin of error for *the difference between two party* scores, $p1$ and $p2$, in a single poll is given by the equation:

$$\pm 1.96 \sqrt{\frac{(p1 + p2) - (p1 - p2)^2}{n}} \qquad (2)$$

We use equation (2) in this study to calculate the "true" margin of error in synchronous horse races rather than the "reported" margin of error of equation (1). In other words, we expect that journalists and pollsters will calculate the "true" margin of error. This expectation conforms with Meyer's concept of "precision journalism," which emphasizes the application of social and behavioural science research methods—statistics

included—to the practice of journalism (Meyer, 2002). When $p1 + p2$ is near 1.0, the true margin of error is close to twice the reported margin of error. But when $p1 + p2$ is less than 1.0, the true margin of error is less than twice the reported margin of error, sometimes substantially less. To see this, suppose a survey with $n = 1,000$ where $p_1 = 40\%$ and $p_2 = 20\%$. Applying equation (2) produces an margin of error of $\pm.0235$, which means that the difference between party scores needs to be equal to or larger than .047 to be statistically significant at 95 per cent. This is substantially less than the .062 difference between party scores that would be required with the $\pm.031$ margin of error of equation (1).

Diachronic horse-race interpretations are statements indicating that the support for a given party is rising, declining or stable across two successive polls. There are four possible diachronic horse-race combinations at the federal level in the 2008 election ($\text{Con}_t\text{-Con}_{t-1}$, $\text{Lib}_t\text{-Lib}_{t-1}$, $\text{NDP}_t\text{-NDP}_{t-1}$, $\text{Bloc}_t\text{-Bloc}_{t-1}$). Reports that also include results for subsamples contain more combinations. To remain in the analysis, a diachronic interpretation must provide enough information to enable research analysts to verify that a party has either gained, lost or kept its support over time and by what magnitude.[5]

To determine the accuracy of a diachronic interpretation, it is necessary to identify the previous poll from which the comparison can be drawn. When the previous poll was not clearly identified, we used all available information to find the right survey. Cases for which the previous poll could not be identified with certainty were dropped from the analysis. Unlike the synchronous horse race, the diachronic horse race compares party scores in two separate polls. The equation to calculate the margin of error when the party scores are from two separate samples is:

$$\pm 1.96 \sqrt{\left(\frac{p1q1}{n1} + \frac{p2q2}{n2} \right)} \tag{3}$$

where subscripts 1 and 2 indicate polls 1 and 2, and where $q_1 = 1 - p_1$ and $q_2 = 1 - p_2$. We use equation (3) in this study to calculate the margin of error in diachronic horse races. Note that if the reported margin of error is the same for both polls, then the margin of error for the difference (change over time) is not twice, but 1.41 times the reported margin of error for each poll (Franklin, 2007). Suppose, for example, that the support for a party drops from 36 per cent to 32 per cent in two successive polls of 1,000 adults. Applying equation (3) gives a margin of error of $\pm.043$, which is 1.41 times the $\pm.031$ reported margin of error obtained from equation (1).

Once we know whether the difference between two scores is statistically significant, we can assess the accuracy of a horse-race statement about them, using the following decision rule. A horse-race interpreta-

tion stating that a party is ahead of another or that a party has gained support over time is coded "accurate" only if this is supported by statistical evidence. If no party is ahead or if support has not changed over time statistically speaking, an interpretation must state that no party is ahead or that support has not changed to be considered accurate. An interpretation to the contrary is coded "inaccurate."[6] Conflicting statements from pollsters or journalists as to whether a party leads another and whether there is change over time are coded "ambiguous." Ambiguous cases are excluded from the analysis.

Reliability Tests

Trained research assistants independently coded the data and their interpretations as found in media and pollsters' reports. A preliminary assessment of inter-coder reliability was conducted based on 50 randomly selected synchronous cases and 50 randomly selected diachronic cases. As far as synchronous cases are concerned, there was almost unanimous agreement about the size of the difference between party scores in pollsters' and in media reports, and on whether this difference passed the 95 per cent test of statistical significance or not. There was 96 per cent agreement about the accuracy of synchronous horse-race interpretations in pollsters' and in media reports. The corresponding Cohen's kappa coefficient—a more conservative measure of inter-coder reliability for nominal variables than per cent agreement (Lombard et al., 2002)—was .86. Based on these very robust results, it was decided that no further inter-coder reliability tests of synchronous cases were needed beyond the preliminary test.

The preliminary assessment of diachronic cases produced almost unanimous agreement about the size of the change over time and whether this difference passed the 95 per cent test of statistical significance. Inter-coder agreement about the accuracy of diachronic interpretations was somewhat lower (Cohen's Kappa = .74). To ensure maximum reliability, it was decided that all diachronic interpretations would be coded by two independent coders, with a resulting Cohen's kappa coefficient of .79. Cases where coders disagreed were submitted to the authors for resolution. Additional characteristics of horse-race interpretations were coded, including the language and the names of the polling and media organizations and so forth. Inter-coder agreement about these additional characteristics was always near unanimous in the preliminary assessment phase.

Interviews

Semi-directive interviews were conducted between 2009 and 2011 with five journalists, one for each newspaper and one for Radio-Canada, and

six pollsters from five polling firms that had conducted polls during the 2008 election (Compas, CROP, Environics, Ipsos Reid, and Léger Marketing). All the journalists we interviewed had reported on the 2008 election and used polling in their stories. Four interviews were conducted in person, six by telephone and one by e-mail. The verbal interviews were all tape recorded. The interviews started with questions about the specific experience of the respondent with polling immediately before and during the 2008 election campaign. Most interviews covered other aspects of polling aside of margin-of-error issues. The portion devoted to the margin of error lasted between ten and twenty minutes. There were questions designed to uncover respondents' knowledge of the margin of error, their perception of its importance in election polls and how this was conveyed to the public. Respondents' perceptions of the relation between pollsters and the media in general, and during the 2008 election campaign in particular, were addressed in some interviews. One of our goals was to go beyond the objective measure of how pollsters and journalists report the horse race, and to try to ascertain their subjective perceptions of the horse race and the margin of error. Most sources agreed to answer our questions on the condition that the interview would be conducted on a not-for-attribution basis.

Findings

To assess whether a horse-race interpretation is accurate, we need to know party scores, or the size of the gap that separates party scores, and the sample size. We identified 120 pollsters' reports containing this information. Cases in which the sample size was missing and could not be estimated by other means were dropped from the analysis (this happened sometime with regional subsamples). In the case of media reports, we needed to locate the relevant pollster's report in addition to the party scores and the sample size. We identified 230 media stories for which we could obtain this information (the sample size was retrieved from pollsters' reports when it was missing from media stories).

How accurate are horse-race interpretations? Table 1 reports the result. There were 172 synchronous horse-race interpretations in pollsters' reports, and 401 in media reports. Out of the 401 media interpretations, 354 (88.3%) are accurate; there are 156 accurate pollsters' interpretations out of 172 (90.7%).[7] Table 2 reports the data for the accuracy of diachronic horse-race interpretations. There are 32 diachronic horse-race interpretations in pollsters' reports, 22 of which are accurate (68.7%). There are 119 diachronic horse-race interpretations in media reports, 66 of which are accurate (55.5%). The proportion of accurate diachronic interpretations is lower than the proportion of accurate syn-

chronic interpretations, due in large part to the relatively low number of pairs of diachronic observations in which there is a statistical difference between two successive sample observations. With a smaller number of occurrences in which the difference is "obviously" outside the sample margin of error, there are fewer opportunities in which pollsters and journalists can automatically avoid making inaccurate interpretations.

Journalists make accurate interpretations of synchronous and diachronic horse-races less frequently than pollsters. This meets theoretical expectation, considering that pollsters are more knowledgeable about sample statistics than journalists. Note however the small size of the accuracy differential between synchronous interpretations by pollsters and by journalists (90.7 − 88.3 = 2.4 percentage points). The accuracy differential in diachronic interpretations is 68.7 − 55.5 = 12.2 percentage points.

Almost all interpretations occurring when there is a statistically significant difference between party scores are accurate. But when the difference is not statistically significant, the frequencies of accurate interpretations drop sharply. From Table 1 we see that only 57 out of 102 synchronous media interpretations (55.9%) and 22 out of 33 synchronic pollsters' interpretations (66.7%) are accurate in cases where the difference is not statistically significant. Similarly, from Table 2 we see that when the difference is not statistically significant, the number of accurate diachronic interpretations decreases sharply (from 100% to 33.3% in pollsters' reports, from 94.3% to 26.2% in media reports).[8] Clearly, a horse-race interpretation stands a very high chance of being accurate when the difference between party scores is statistically significant. Including cases of statistically significant different party scores

TABLE 1

Accuracy of Synchronous Horse-race Interpretations in Pollsters' and Media Reports

	Pollsters' Reports			Media Reports		
	No statistical difference	Statistical difference	Total	No statistical difference	Statistical difference	Total
Accurate Interpretations	22 (66.7)	134 (96.4)	156 (90.7)	57 (55.9)	297 (99.3)	354 (88.3)
Inaccurate Interpretations	11 (33.3)	5 (3.6)	16 (9.3)	45 (44.1)	2 (0.7)	47 (11.7)
Total	33 (100)	139 (100)	172 (100)	102 (100)	299 (100)	401 (100)

Notes: Column percentages in parentheses.
$\chi^2 = 27.95$ (pollsters) $p \leq .01$; $\chi^2 = 138.76$ (media) $p \leq .01$
Ambiguous interpretations and media interpretations with no identifiable pollsters' source were dropped from the analysis.

TABLE 2

Accuracy of Diachronic Horse-race Interpretations in Pollsters' and Media Reports

	Pollsters' Reports			Media Reports		
	No statistical difference	Statistical difference	Total	No statistical difference	Statistical difference	Total
Accurate	5	17	22	16	50	66
Interpretations	(33.3)	(100)	(68.7)	(24.2)	(94.3)	(55.5)
Inaccurate	10	0	10	50	3	53
Interpretations	(66.7)		(31.3)	(75.8)	(5.7)	(44.5)
Total	15	17	32	66	53	119
	(100)	(100)	(100)	(100)	(100)	(100)

Notes: Column percentages in parentheses
$\chi^2 = 16.49$ (pollsters) $p \leq .01$; $\chi^2 = 58.47$ (media) $p \leq .01$
Ambiguous interpretations and media interpretations with no identifiable pollsters' source were dropped from the analysis.

in the analysis would drown instances of inaccurate statements in a sea of accurate statements, and this would introduce a bias in the results. For this reason, we will only test the accuracy of horse-race interpretations in the cases in which party scores are not statistically different, when pollsters and journalists are faced with the real possibility of making a mistake.

Table 3 reports the variation in the level of accuracy of pollsters' interpretations of 48 horse races in which the difference between party scores is not statistically significant (33 synchronous + 15 diachronic interpretations). The overall level of accuracy is 56 per cent on average and ranges from a high of 100 per cent (Segma) to a low of 25 per cent (Ipsos Reid). What accounts for the variation in accuracy levels? Quebec seems to be a factor. Interpretations by Quebec-based polling firms (CROP, Léger Marketing, Segma, GPS) are accurate 78 per cent of the time on average, whereas interpretations by other firms are accurate only 43 per cent of the time.[9] Tracking polls are also a factor. One-half of all the media interpretations under study are based on tracking polls conducted by two polling organizations: Strategic Counsel and Harris-Decima (see the appendix for details). The Strategic Counsel conducted daily tracking polls in key ridings for *The Globe and Mail* and CTV. Harris-Decima did daily tracking polls for Canadian Press and the results were reported in *Le Devoir*.[10] The average accuracy level for these two firms is a paltry 36 per cent, much lower than the average among the remaining firms. Two factors contribute to increased likelihood of inaccuracy in tracking poll interpretations. First, they have larger margins of error than regular polls (due to their smaller sample size); second, because they focus on predicting the race and nothing else, they can only be framed in terms of