

For Want of a Nail: Field Experimental Evidence on the Effectiveness of Direct Mail in a Political Campaign*

Peter John Loewen

Département de science politique, Université de Montréal
peter.john.loewen@umontreal.ca

Daniel Rubenson

Department of Politics, Ryerson University
rubenson@ryerson.ca

This version printed on May 24, 2007

*Note: This is a working dissertation article.
Please do not quote or cite without permission.*

Abstract

Direct mail is a pervasive feature of modern political campaigns. But does it work? This paper presents evidence from a field experiment into the effects of direct mail. Working with a front-running campaign during the race for the leadership for the Liberal Party of Canada, we randomly assigned a subset of convention delegates to receive a direct mail treatment. We then measured the effects of this treatment on delegates' ratings and preference ordering of leadership candidates using a survey instrument. Our results indicate that despite its ubiquity, the direct mail had minimal and probably negative persuasive effects.

1. INTRODUCTION

Does direct mail work? Political campaign managers certainly believe it does. In nearly every type of political campaign at every level of competition, some form of mail is used. Sometimes this mail serves the purpose of outlining a candidate's position, or casting an opponent's position in an unfavourable light. At other times it is used for fundraising. It also sometimes serves a mobilizing function, encouraging potential voters to participate in an election. Most often it takes up several of these tasks at once. Whatever its purpose, there seems little question that direct mail is a frequently used tool in politics generally. Its use in Canadian politics is no exception.

The ubiquity of direct mail is easily explained. It is a relatively cheap manner in which to reach a large number of voters. Moreover, when its design incorporates individual-level data on a voter's preferences or concerns (or even their consumer habits and financial status) it promises still greater potential effectiveness. In this sense, it epitomizes the "electronic politics" which close

*We thank the Michael Ignatieff Leadership Campaign for their cooperation. We also thank André Blais, Mehdi Mourali, Tom Flanagan, Clinton Free, Kevin Arceneux, David Nickerson, Richard Niemi, Amanda Bittner, Angelo Elias, Lorena Rucci, Stuart Soroka, Patrick Fournier, Brian Doody and Christophe Chowanietz for comments. Loewen thanks André Blais for the time and space to write the paper. He acknowledges support from SSHRC in the form of a CGS Doctoral Grant. Rubenson thanks Cheryl Johnston for research assistance and the Faculty of Arts, Ryerson University for financial assistance with the project.

observers of Canadian politics have identified as characteristic of our current party system [Carty et al., 2000]. Direct mail allows parties to personally connect with voters through potentially highly targeted messages. This combination of low cost and tailored messaging should increase the importance of direct mail into the future.

Political direct mail has not undergone systematic academic study in Canada and as such it is difficult to quantify its volume or effectiveness. However, some accounts have given us an idea of its importance to both federal elections [Carty et al., 2000, Carty and Eagles, 2005] and leadership races [Flanagan, 2003, Wearing, 1988]. Moreover, anecdotal evidence, personal experience, and conversations with campaign managers suggest to us that direct mail forms an important part of modern campaigns in Canada. But the question remains: Does it work?

This paper presents evidence from a recent field experiment into the effectiveness of direct mail in changing vote intentions and candidate evaluations in the 2006 federal Liberal leadership race. To our knowledge, this represents the first field experiment in modern Canadian political science and the first within the context of Canadian leadership elections. To anticipate our findings, we conclude that, for at least one front-running campaign, there was no positive effect from sending one or two pieces of direct mail to delegates. While this finding does not conclusively demonstrate the inutility of direct mail, it does represent one strong piece of evidence against its effectiveness. This is especially so for two reasons. First, because we worked in partnership with a campaign, our treatments are entirely realistic. As Gerber notes, “if field experiments can be conducted in cooperation with actual campaigns... questions of external validity will be minimized” [Gerber, 2004, 547]. Second, as we argue below, evidence from field experiments should be given substantially more weight than that drawn from purely observational accounts.

The paper is organized as follows. We begin by outlining current knowledge on the effects of political direct mail. We then briefly discuss the race in which the experiment occurred. In section four we outline our field experiment and justify its use in comparison to other inferential techniques. Section five presents our model and results. After discussing our findings, we conclude.

2. DIRECT MAIL

Direct mail is ubiquitous in Canadian politics. Older evidence suggesting the importance of printed materials, such as that presented by Paltiel [1974], has been confirmed by recent analyses of modern campaigns. Carty and Eagles [2005], in particular, have documented the importance of printed advertising for modern local campaigns. Using data from the 2000 Canadian federal election, they observe that print advertising was the largest expense of candidates in all parties. While this material encompasses much more than just direct mail, our own conversations with local campaign managers have suggested that direct mail makes up a large portion of this spending, and often the largest. Clearly, it is a tool frequently drawn from a campaign manager’s toolbox. This trend promises to continue as parties become increasingly adept at collecting individual-level data and mining it for insights which can then be leveraged through direct contact with individual voters [see Carty et al., 2000, and for a more popular account Wells, 2006].

The importance of direct mail in federal elections is probably surpassed by its importance in party leadership races. Whether conventions or direct elections, leadership races seem especially amenable to this campaign tool. These races are often paid little sustained attention by the media, especially for less competitive candidates. They tend to feature candidates who are often difficult to distinguish on ideological or policy grounds[Vavreck et al., 2002]. Moreover, party leadership campaigns are increasingly large-scaled affairs in which it is difficult for candidates to personally reach every member in the electorate through face-to-face meetings [Cross, 1996,

312]. At the same time, the number of eligible voters (i.e. party members) relative to the typical budget does make it possible to reach each voter by mail, often multiple times. Mail thus allows a candidate to speak directly to each party member or delegate. Finally, in races with many candidates, persuasion becomes a principal activity as campaigns seek to build coalitions which can deliver a majority of delegates or voters over a series of ballots. Direct mail plays an important role in this persuasion. Wearing's [1988] accounts of the 1976 and 1983 Progressive Conservative and 1984 Liberal leadership convention campaigns and Flanagan's [2003] account of the 2002 Harper campaign for the leadership of the Canadian Alliance Party both provide convincing evidence of the importance which campaign managers assigned to direct mail in these races—an importance which we think generalizes fairly easily to all leadership races. Whether direct mail actually works, however, remains unclear.

In contrast to the political science literature, marketing is one field in which direct mail has been extensively studied. As a result, a substantial and broad literature exists. Among its findings, the marketing literature includes theory and knowledge about the elements of direct mail which make for success [eg Nash, 1984, Elsner et al., 2004], how direct mail campaigns (especially coupons) affect purchasing [Bawa and Shoemaker, 1989, Bult and Wansbeek, 1995], how they affect incremental sales, and how direct mail campaigns can be optimized based on past purchasing information [Allenby et al., 1999, Neslin et al., 1985]. Moreover, much of this literature includes an experimental element. For example, Irons et al. [1999] present a meta-analysis of *sixty* field experiments on the effects of coupons on purchasing habits.

We can learn clear methodological lessons from this literature, particularly about the analytical power of field experiments. But despite this, it is unclear how much we can apply the lessons of consumer behaviour to electoral politics. The decision to consume more goods or change the mix of goods which an individual consumes does not accurately reflect the nature of political choice in which a decision is forced (you have to vote at a certain time), zero-sum (you have to vote for one candidate and not others), and essentially civic (in that one is likely, in making their choice, to think about more than self-interest or the meeting of a need). In short, individuals may bring a substantially different calculus to vote choice, one which is responsive in a different way—or not at all—to direct mail efforts.¹ Accordingly, we need to look principally to evidence within politics and political science.

Whether in general elections or leadership contests, there is a lack of systematic evidence in Canada on the effectiveness of political direct mail. As Carty and Eagles have observed, Canadian political scientists have lagged colleagues in other countries in studying local campaign methods, of which direct mail is just one [2005, 120]. What evidence does exist about local campaign effects in Canada generally supports the conclusion that local spending and campaign strength (i.e. volunteers and organization) matter for election outcomes. For two reasons however, an examination of direct mail using existing data on party or candidate spending cannot be undertaken. First, even if we can assume that spending is measured consistently and correctly [Ansolabehere and Gerber, 1994], it is still accounted for in a manner which prevents the identification of direct mail outlays specifically [Loewen, 2005]. Second, even if we could observe the different types of spending, we could not easily extract strong causal statements from these observations. While we discuss this in greater length in the next section, the basic problem is easily stated: Because spending decisions are not made randomly, we cannot determine if the effects of campaign practices are a function of the types and extent of the method or the unobserved factors which influence campaigns to choose some methods over others.² This problem is far from unique to Canada or Canadian political science—indeed, we can find many examples of observational re-

¹Not to mention that direct mail may vary systematically in its design from that in the commercial world.

²This is only in addition to the problems of the endogeneity of campaign fundraising.

search on campaign effectiveness in other countries which are confronted with the same empirical problem [for British examples, see Johnston and Pattie, 1998, Pattie et al., 1995, Whiteley and Seyd, 1994, for American examples, see Holbrook and McClurg, 2005, Vavreck et al., 2002]. What is required is some form of inquiry not subject to Leamer’s “inferential monsters lurking beyond our immediate field of vision” [1983, 83]. That is, some form of inquiry where we can reasonably limit the number of possible explanatory variables and focus on one in particular—ie direct mail.

A growing line of research in the United States has sought to confront this problem of unobserved heterogeneity in campaign effects by engaging in field experiments. This research program has been both wide and deep. It covers several different campaign methods including direct mail, door-to-door canvassing, various telephone techniques, and leafleting; and it reaches down into several types of elections, several different types of campaigns, and several different locales. The most important feature of these experiments is the randomized assignment of a treatment of interest to a well-defined population, followed by a statistical analysis of the effects of the treatment [Green and Gerber, 2004, 11–22]. While these studies have not been without criticism — especially in terms of execution and estimation [eg Imai, 2005]—they have allowed for strong conclusions to be drawn on the effects of direct mail, especially as it relates to mobilization. Following Green and Gerber’s [2004] summary, while non-partisan direct mail seems to increase turnout, that expressing opposition to a candidate does not seem to have an effect. Partisan mail is effective in mobilizing partisans but not in bringing “swing voters” to the polls. On balance, the mobilizing effects of direct mail appear highly conditional.

Less work has been undertaken on the persuasive effects of partisan direct mail. One early study examines the effects of a single candidate mailing in a weakly contested Democratic congressional primary [Miller and Robyn, 1975]. It found no effect, though it is subject to important questions of statistical power. Gerber [2004] single-handedly expanded the field, conducting field experiments with five different campaigns during the 1999–2000 election cycle. These experiments—conducted during a mayoral race, a New Jersey state assembly election, a state legislative race in Connecticut, a Congressional primary and a Congressional general election—examined the effect of campaign mailings on vote totals, which we take to be a test of the persuasive capacities of direct mail. In some cases, post-election surveys were used to estimate effects while in others they were measured by ward-level differences in vote totals. The results generally show that while incumbent mailings had little effect (except in primaries), challenger mailings were effective in some cases.

Taken together with the mobilization literature, it is difficult to arrive at a firm conclusion on the effectiveness of direct mail. Its utility is contingent both on the type of race and the type of candidate. As a consequence, these results do not directly inform our study of the persuasiveness of direct mail in a leadership race towards one expectation or another. However, they do demonstrate two things. First, we can effectively ascertain the causal properties of campaign methods through field experiments. Second, at least some of the claims of those who advocate direct mail appear to be false. The mobilizing capacity of direct mail has, at best, been overestimated by its advocates. Might it be the same for its persuasive properties? After describing the race for the Liberal leadership, we turn to this question.

3. THE RACE

After losing the January 2006 federal election, Prime Minister Paul Martin resigned as parliamentary leader of the Liberal Party of Canada. In the subsequent weeks the party outlined the conditions of its leadership selection process—much of which was predetermined by the party’s constitution. A leadership convention was held in Montreal on December 3, 2006. Delegates to

the convention were elected from among party members. In addition to ex-officio delegates who were guaranteed a place at the convention, these regular delegates from each federal electoral district were allotted to leadership candidates according to the total preferences of all members in that electoral district. On the first ballot non-ex-officio delegates were thus obliged to vote for the candidate to whom they were pledged. Indeed, they received marked ballots upon their arrival at the convention. This apportionment process occurred at a “Super Weekend” at the end of September. Only those party members who were of good standing as of July 1, 2006 were allowed to vote in the Super Weekend.

The race was nothing if not exciting. More than twenty names were identified as potential candidates and eleven officially entered. By the time of the delegate selection meetings the field had eventually narrowed to eight candidates. Michael Ignatieff was the clear front runner, obtaining the support of about 30 percent of pledged delegates, as well as many ex-officios. Ignatieff, recently returned from more than twenty years outside the country as an academic and journalist, was generally seen as being on the right of the party. He was a polarizing candidate. Bob Rae, a former Premier of Ontario (as the leader of the social democratic New Democratic Party), was the clear second place candidate. He can also be regarded as polarizing. Rounding out the top four were Gerard Kennedy, a former Ontario provincial cabinet minister and Stéphane Dion, a former federal cabinet minister (and political scientist) known far more for intellectual battles with sovereigntist/separatist leaders in Quebec than for his political *panache*. The bottom four was comprised of Ken Dryden, Joe Volpe, Scott Brison, and Martha Hall Findlay.

To the surprise of many, Dion would eventually win the leadership. Results from the pre-convention delegate selection and the four ballots at the convention are presented in Table 1. Dion finished in third place on the first ballot, just two delegates ahead of Kennedy. He would receive Kennedy’s endorsement after widening his lead on the second ballot, and on the strength of that hand-tipping he would vault past both Ignatieff and Rae on the third ballot, thus eliminating Rae. He defeated Ignatieff on the fourth and final ballot, receiving 54.7 percent of the votes to Ignatieff’s 45.3 percent. Rather than polarizing delegates, as the two front-runners had, Dion was successful in his portrayal of himself as a safe second-choice. Whether by luck or design, he showed himself a master of convention politics.

Table 1: 2006 Liberal Party leadership election results (%)

Candidate	Pre-convention delegates		1 st ballot		2 nd ballot		3 rd ballot		4 th ballot	
Ignatieff	1,377	(29.3)	1,412	(29.3)	1,481	(31.6)	1,660	(34.5)	2,084	(45.3)
Rae	943	(20.1)	977	(20.3)	1,132	(24.1)	1,375	(28.5)		
Kennedy	820	(17.5)	854	(17.7)	884	(18.8)				
Dion	754	(16.1)	856	(17.8)	974	(20.8)	1,782	(37.0)	2,521	(54.7)
Dryden	238	(5.1)	238	(4.9)	219	(4.7)				
Volpe	226	(4.8)	156	(3.2)						
Brison	181	(3.5)	192	(4.0)						
Hall Findlay	46	(1.0)	130	(2.7)						
Undeclared	112	(2.4)								
Total votes	4,697	(100.0)	4,815	(100.0)	4,690	(100.0)	4,817	(100.0)	4,605	(100.0)

Our own experiment was situated within the period between the election of delegates and the convention in Montreal; what Wearing calls the “second stage” of delegated conventions [1988]. This period provided a crucial test of the persuasive ability of campaigns. Rather than selling memberships and encouraging supporters to stand as delegates, campaigns in this period of the process were dedicated to ensuring delegates attended the convention and persuading delegates

to select their candidate as their next choice should their preferred leadership candidate fall-off the ballot or withdraw. Among many tactics, direct mail played an important role in this crucial period. For example, the Dion campaign sent a DVD featuring a series of short interviews with their candidate. The Ignatieff campaign sent a 40-page bilingual policy book entitled “Agenda for Nation Building: Liberal leadership for the 21st century”³. The book outlined in detail Ignatieff’s policy on the economy, the environment, the constitution, national unity and foreign affairs. The Ignatieff campaign also sent out a simple, colour brochure summarizing Ignatieff’s positions.

4. THE EXPERIMENTAL STUDY

Our experiment included two components. First, a randomized program of direct mail from a front-running campaign over a subset of elected delegates in the last week of October and first week of November. Second, an academic mail-back survey of the same delegates which measured, among other things, their evaluations of each candidate as well as their preferences between the various leadership candidates. We describe each in more detail below.

4.1. THE EXPERIMENT

Our experiment relied on a partnership with the Michael Ignatieff campaign. After the selection of delegates at the end of September, we randomly selected a subset of 800 delegates from those who had a current address on the official party list of delegates.⁴ In addition to restricting our sample to those delegates who had addresses, we also excluded Quebec, Manitoba and British Columbia.⁵ Among these 800 delegates, we identified those who had not pledged to support Ignatieff at delegate selection meetings, reducing our sample to 567. Among these remaining delegates, we randomly assigned 100 to receive two pieces of mail from the Ignatieff campaign and 200 to receive one piece of mail.⁶ All those who received mail received a copy of Ignatieff’s policy book. Those who were assigned to receive a second piece of mail also received a copy of a colour brochure. By randomly assigning mail we (theoretically) ensured that the reception of mail was not a function of a respondent’s personal characteristics or preferences. As with conventional random assignment in a laboratory, this affords us much analytical leverage.

4.2. THE SURVEY

Subsequent to the sending of the direct mail, we mailed each delegate within our subset an academic survey from the Department of Politics at Ryerson University. By sending the survey under the cover of the University, we concealed any connection between the survey and the experiment. Moreover, we excluded any mention of the field experiment on our respective academic websites. The survey included a postage-paid return envelope, as well as an ethics disclaimer and short introduction. The survey obviously made no mention of the experiment, though it did include recall questions on the reception of direct mail from campaigns since the selection of delegates. Most pertinent to our study, the survey included questions about preferences and the evaluations of candidates, which allowed us to test the effects of direct mail.

³The title on the French side of the book was “Bâtir notre nation: le leadership libéral pour le 21e siècle.”

⁴This represented approximately 16% of delegates.

⁵Excluding Quebec allowed us to avoid the use of a bilingual survey and ensured that all delegates were receiving the same mail, thus reducing the need to condition our results on the language of the mail. Manitoba and British Columbia were excluded from the party’s delegate list at the time because of incomplete delegate lists or disputes between several campaigns over the status of various delegates.

⁶We describe our treatment assignment procedure in more detail in Appendix A.

The advantages of combining a survey and an experiment become clear when we consider the most typical alternative approach to studying the impact of campaigns on individual voters, ie a survey which may or may not include contextual information about the campaign [see Perlin, 1988, Stewart, 1997, Vavreck et al., 2002, Bartels, 1987, for leadership campaign examples]. As Gerber and Green have argued, relying on a survey alone to gauge the effects of direct mail—and other campaign contacts more generally—suffers from two problems [2000b, 2004]. First, individual respondents are demonstrably poor at recalling whether or not they have received mail from a campaign. For example, our survey included a recall question which asked delegates to identify from which campaigns they had received mail since the conclusion of delegate selection meetings. Because we know which delegates received mail from the Ignatieff campaign we were able to measure the level of error in delegate recall. Of those who did not receive mail from the campaign, 85 percent correctly recalled that they received no mail. However, 15 percent did report receiving mail. The case is more grave with those who did receive mail, with less than two-thirds (64%) correctly recalling receiving mail. Moreover, based on a question-wording experiment embedded in our survey, we found that recall was not improved by giving some delegates a further prompt identifying the types of mail they may have received. Accordingly, even with a carefully designed survey we would risk serious measurement error in identifying who received direct mail from a campaign. Our study avoids this pitfall because we know to whom the campaign sent mail.⁷

Second, political campaigns are often strategic in their targeting of direct mail. Mailings are targeted and tailored to reflect a campaign’s beliefs about the recipient. For example, campaigns may be more likely to send mail to those who they believe are at least open to supporting the campaign. By merely observing the relationship between direct mail and behaviours, we cannot know whether any direct mail effect is the result of the mail itself or the individual in question being predisposed to support the party. Even with a bevy of control variables this problem cannot be easily solved statistically, if at all [Gerber et al., b]. However, in the case of our experiment we know that the assignment of mail was random and thus uncorrelated with individual characteristics. Any observed effect is likely the result of mail.

In comparison to the American experiments outlined above, one caveat is in order. Those experiments typically test the effects of a treatment—direct mail for example—on a directly observable behaviour such as voting as determined by an official record. Our experiment, by contrast, still relies on estimates of an effect drawn from a survey. As we could not peer inside the ballot boxes at the party convention, we are left to ascertain the effects of direct mail through our survey questions. We are thus left open to many of the problems associated with survey responses. However, we are not confronted with the more fundamental problems of respondent recall or the strategic allotment of a treatment. As a result of this, the effects we observe are “real” to the extent that surveys capture “real” aspects of delegates’ considerations and evaluations in the run-up to the convention.

Our final sample includes 161 respondents, representing a response rate of 28 percent. This sample is evenly balanced between those who did receive mail (81) and those who did not (80). Treatment is unrelated to the pledged support of delegates ($\chi^2 = 7.78$, $p = 0.35$), province of residence ($\chi^2 = 3.73$, $p = 0.81$), or delegate type⁸ ($\chi^2 = 8.95$, $p = 0.26$). Most importantly,

⁷We also know that the mail at least landed in their mailbox. We used the same addresses for the surveys as for the mail, so it is not possible that a delegate received and responded to our survey without receiving the mail. While delegates may very well be selective about what they choose to read [see Barlett et al., 1974], we can be certain that we are at least dealing with cases in which they had the opportunity to read the mail sent to their address.

⁸Delegates are classified by the party according to gender, age, and aboriginal status.

survey response is unrelated to our 3-category treatment assignment ($\chi^2 = 0.61, p = 0.74$).

5. RESULTS

We measure the effects of direct mail across three different variables. In each case, we compare those who did and did not receive mail using simple statistical tests.⁹ Rather than specifying complicated models, we rely on the power of random assignment.

Table 2: Effects of Ignatieff mail on average leadership candidate ratings^a

Candidate	Mean rating with no mail (s.d.)		Mean rating with mail (s.d.)		Pr(Effect on leader rating) ^b	Mail	N No mail
Michael Ignatieff	46.3	(32.6)	47.8	(33.7)	0.38	75	80
Bob Rae	61.4	(33.6)	63.6	(33.9)	0.66	75	81
Gerard Kennedy	73.6	(26.1)	73.9	(26.7)	0.53	74	80
Stéphane Dion	72.6	(23.7)	77.3	(21.4)	0.90	75	79
Ken Dryden	60.7	(26.0)	54.7	(26.9)	0.08	72	79
Joe Volpe	15.5	(22.6)	20.8	(25.5)	0.91	70	79
Scott Brison	42.4	(27.5)	54.3	(26.1)	0.99	67	76
Martha Hall Findlay	49.5	(27.0)	49.7	(27.9)	0.51	65	78

^a Note: Calculations of difference rely on an unpaired *t*-test with an assumption of unequal variance.

^b For Ignatieff, this is the probability that leader ratings are higher with mail; for the other candidates, it is the probability that leader ratings are higher without mail.

We first measure whether those who received direct mail evaluate the likeability of the eight candidates differently than those who did not. The expectation of those sending direct mail—at least for the campaign in question—was that mailers would increase evaluations of their own candidate and reduce evaluations of other candidates. Using a conventional 0-100 rating scale we find little evidence of a positive effect for direct mail (Table 2). Those who received mail did not give higher ratings to Ignatieff, on average. And only in the case of Ken Dryden did they move their evaluations down. Moreover, in the case of Dion, Brison, and Volpe, it appears that direct mail from the Ignatieff campaign *increased* the likeability of these candidates. On the whole, receiving mail did not move the opinions of those who were not already pledged to support Ignatieff in the expected direction.¹⁰

Given the multi-ballot nature of a competitive delegated convention, moving a candidate up in delegates' preference rankings is a principal objective for campaigns.¹¹ Receiving direct mail from the Ignatieff campaign appears to have done little to achieve this objective. Those who

⁹Our treatment regime specified that some individuals receive two pieces and others one piece. Because of our relatively small *n*, we have collapsed these two treatments into one in the analysis. Our substantive results do not change when we consider those who received two pieces of mail separately.

¹⁰We have also estimated these effects with separate OLS regressions for each candidate with leader rating on the left-hand side and mail and a small number of control variables on the right-hand side. Our results do not change.

¹¹Indeed, from the polling of delegates and party members that occurred between delegate selection and the convention, it is apparent that there was considerable movement among delegates' second choice preferences. In other words, delegates were not completely fixed in their preferences on second and subsequent ballots. As such, candidates did have an opportunity during this period to persuade potential supporters to move to them on later ballots [The Strategic Counsel, 2006b,a, Canadian Press, 2006, Laghi, 2006].

Table 3: Effects of Ignatieff mail on making Ignatieff a second or third preference

	No Mail (%)	Mail (%)	Total (N)
No	67.5	74.1	114
Yes, 2nd or 3rd preference	32.5	25.9	47
Total (N)	80	81	161

$\chi^2 = 0.84, p = 0.359$

Table 4: Effects of Ignatieff mail on respondent indicating they would *never* vote for Ignatieff

	No Mail (%)	Mail (%)	Total (N)
No	63.8	51.9	93
Yes, never	36.3	48.2	68
Total (N)	80	81	161

$\chi^2 = 2.33, p = 0.126$

received mail from the Ignatieff campaign were no more likely to identify him as a 2nd or 3rd choice than those who received no mail (Table 3). Similarly, receiving Ignatieff mail did nothing to prevent delegates from ruling him out altogether. Among those who received mail, nearly half (48 percent) indicate that they would never vote for Ignatieff. Only 36 percent of those who did not receive mail indicated that they would never vote for Ignatieff (Table 4).¹² Despite these distributions not reaching conventional levels of statistical significance, we do note that the direction of effects is counter to that expected. As with candidate ratings, those who received direct mail from the Ignatieff campaign appeared less likely to adopt Ignatieff as a second or third choice and more likely to indicate that they would never vote for him.

As each of these three sets of results are drawn from a relatively small sample ($n = 161$), there is a possibility that we have committed type II errors. Indeed, direct mail may exert a small but significant effect which our study does not have the power to detect. Consider the case of candidate ratings. Given our sample size and our prior knowledge of the average standard deviation associated with candidate ratings¹³ we can use standard statistical power calculations to determine the magnitude of detectable effects [R.V. Lenth, 2006]. Following these, we are able to detect candidate rating increases of 8 points at $p = 0.10$ and 9.5 points at $p = 0.05$. Given that some of our results are smaller than these, we may very well be making a type II error. We have two responses to this. First, if the insignificant effects we have found are in fact real, they are very small and substantively meaningless. If candidate ratings in leadership races resemble those of

¹²We have also estimated these effects with separate logistic regressions for each candidate with a second or third preference ranking of Ignatieff and for an indication of never voting for Ignatieff on the left-hand side and mail and a small number of control variables on the right-hand side. In no case do we find a positive effect for mail.

¹³We used a file of leader ratings from Canadian election studies from 1974 on, available at the Canadian Opinion Research Archive: <http://www.queensu.ca/cora/trends/tables/PartyLeaders.htm>.

leadership ratings in general elections, then we can say with some confidence that an increase of 1.5 points is *substantively* insignificant [see Blais and Perrella, 2006]. Second, with the exception of two estimates (the increase in Ignatieff’s rating and the decrease in Dryden’s rating as a result of mail) none of the results are even in the expected direction, including the results in Tables 2 and 3. Accordingly, we are confident that despite the potential for questions of inadequate statistical power, we are avoiding type II errors. And we are confident in our conclusion that direct mail did not have significant effects in the expected direction for the Ignatieff campaign. To the contrary, to the degree that our study was able to uncover significant effects, they were in the opposite direction to that expected, a result which puts into further question the utility of direct mail in this instance.

6. DISCUSSION

Taken together, these results lead us to a clear conclusion. Direct mail appears to have had little positive effect for the Ignatieff campaign. In the face of crystallized preferences, receiving one or two mailings from a campaign was not enough to alter delegates’ assessments or intentions. Rather, if it had any systematic effect it was in making delegates more negative towards Ignatieff’s candidacy. This finding should give campaign managers pause. Political direct mail is a communication of a message which can have three effects. It can increase the appeal of a candidate; it can have no effect; or it can decrease the appeal of a candidate. If direct mail makes clear positions or attributes which voters find objectionable, it may have such a negative effect. For a candidate as polarizing as Michael Ignatieff, this final outcome could have been very real.

The larger implications of these findings are clearly conditioned by the fact they derive from just one campaign. In other words, the external validity of the results is necessarily limited. Nonetheless, the study does represent a proper experiment with an entirely realistic treatment. As McDermott notes, experiments are often dismissed on account of criticisms about external validity, while what is even more important is internal validity [2002, 334–6]. We are confident about the internal validity of this study. Accordingly, these findings should be added to the larger (experimental) literature on direct mail effects, with which they largely agree [Gerber and Green, 2000b,a, Gerber et al., a, Gerber, 2004, Green and Gerber, 2004, Green, 2004].

Given that we have found no positive effect for direct mail, why do we observe campaigns devoting substantial resources to this tool? We have three explanations. First, campaign operatives are certain that these tools work. This message is rather consistently delivered in trade publications such as *Campaigns and Elections* and in operative training sessions, such as the “universities” which Canadian parties hold prior to elections. It only makes sense to use these tools, given the received wisdom. Second, it is not difficult to talk oneself into believing that a chosen campaign tool is working despite a lack of evidence of positive effects or evidence to the contrary. In the hubbub and stress of a campaign an operative will look for any affirmation that things are on the right track. A positive comment about direct mail can quickly become enough to convince one of larger effects. Similarly, it is easy to become convinced of the importance of, say, direct mail, when one knows it is being used by other campaigns. A third possibility exists—one which is less pessimistic about the analytical abilities of campaign managers. Even if direct mail were known to have very small effects, it may still be the most efficient use of resources. Volunteers cannot be bought, professional call centres and automated calls are demonstrably inefficient, a candidate can only work phones or shake hands a certain number of hours each day and time cannot be stretched. The implication of this is that a campaign which did not spend its remaining money on direct mail may not be able to spend it at all. Moreover, direct mail can be sent at a relatively low cost and can often be scaled up into repeated or more substantial mailings

fairly easily. Indeed, once a campaign has settled on a message and obtained a list of voters, the marginal cost of mailing consists only of the cost of producing materials and postage. Knowing this, why would a campaign not spend whatever extra resources it had on printed material? Perceiving that direct mail has some effect, knowing that it is widely used in other campaigns, and being able to send it rather economically, what campaign manager could be expected to take the risk of not sending the mail? Like the fall of a kingdom for the want of a nail, who would risk the loss of a campaign for the want of mail? Our results would suggest that even if direct mail is the most sensible expenditure given resource and time constraints, it may be a message better left unsent.

7. CONCLUSION

Direct mail is a widely and increasingly used tool within modern electoral politics, to which Canada is no exception. Nonetheless, its effectiveness is as much a matter of lore as fact. Going forward, operatives and political scientists both should be more sceptical of claims of the efficacy of direct mail and they should be more ready to engage analytical methods which will allow for a real testing of effects.

We have presented results from a field experiment into the effectiveness of direct mail in the 2006 leadership race of the Liberal Party of Canada. This study represents a first for field experimentation in Canadian politics. Our results suggest that—at least for the Michael Ignatieff campaign—direct mail had little positive effect. It did not improve delegates’ evaluations of the candidate or push downwards their evaluations of other candidates. Moreover, it did not systematically effect the candidate’s positioning within delegates’ preference rankings. Indeed, the only effects which it seemed to have were harmful to the campaign. This time, a kingdom was not lost for the want of a nail.

REFERENCES

- Greg M Allenby, Robert P Leone, and Lichung Jen. A dynamic model of purchase timing with application to direct marketing. *Journal of the American Statistical Association*, 94(446): 365–74, 1999.
- Stephen Ansolabehere and Alan Gerber. The mismeasure of campaign spending: Evidence from the 1990 u.s. house elections. *Journal of Politics*, 56:1106–1118, 1994.
- Dorothy L Barlett, Pamela B Drew, Eleanor G Fahle, and William A Watts. Selective exposure to a presidential campaign appeal. *Public Opinion Quarterly*, 38(2):264–70, 1974.
- Larry M Bartels. Candidate choice and the dynamics of the presidential nominating process. *American Journal of Political Science*, 31(1):1–30, 1987.
- Kapil Bawa and Robert W Shoemaker. Analyzing incremental sales from a direct mail coupon promotion. *Journal of Marketing*, 53:66–78, 1989.
- André Blais and Andrea Perrella. How much does a party/candidate need to be liked to get a vote? Working Paper, Université de Montréal, December 2006.
- Jan Roelf Bult and Tom Wansbeek. Optimal selection for direct mail. *Marketing Science*, 14(4): 378–94, 1995.

- Canadian Press. Ignatieff, Rae leadership favourites among Liberal party members. *The Globe and Mail*, Sep. 19 2006. On-line edition; accessed May 24, 2007.
- R. Kenneth Carty and Munroe Eagles. *Politics Is Local: National Politics at the Grassroots*. Oxford University Press, Don Mills, 2005.
- R. Kenneth Carty, William Cross, and Lisa Young. *Rebuilding Canadian Party Politics*. UBC Press, Vancouver, 2000.
- William Cross. Direct election of provincial party leaders in Canada, 1985-1995: The end of the leadership convention? *Canadian Journal of Political Science*, 29(2):295–315, 1996.
- Ralf Elsner, Manfred Krafft, and Arnd Huchzermeier. Optimizing rhenania’s direct marketing business through dynamic multilevel modeling (dmlm) in a multicatalog-brand environment. *Marketing Science*, 23(2):192–206, 2004.
- Tom Flanagan. Database party: The 2002 leadership campaign for the Canadian Alliance. *Canadian Parliamentary Review*, 26, 2003.
- Alan S Gerber. Does campaign spending work?: Field experiments provide evidence and suggest new theory. *American Behavioral Scientist*, 47(5):541–74, 2004.
- Alan S Gerber and Donald P Green. The effect of a nonpartisan get-out-the-vote drive: An experimental study of leafletting. *Journal of Politics*, 62(3):846–57, 2000a.
- Alan S. Gerber and Donald P. Green. The effects of canvassing, direct mail, and telephone contact on voter turnout: A field experiment. *American Political Science Review*, 94(3):653–63, 2000b.
- Alan S Gerber, Donald P Green, and Matthew N Green. The effects of partisan direct mail on voter turnout. *Electoral Studies*, 22:563–79, a.
- Alan S Gerber, Donald P Green, and Edward H Kaplan. The illusion of learning from observational research. In Shapiro et al., pages 251–73.
- Donald P Green. Mobilizing African-American voters using direct mail and commercial phone banks: A field experiment. *Political Research Quarterly*, 57(2):245–55, 2004.
- Donald P. Green and Alan S. Gerber. *Get Out the Vote*. Brookings Institution Press, Washington, 2004.
- Thomas M Holbrook and Scott D McClurg. The mobilization of core supporters: Campaigns, turnout, and electoral composition in United States presidential elections. *American Journal of Political Science*, 49(4):689–703, 2005.
- Kosuke Imai. Do get-out-the-vote calls reduce turnout? the importance of statistical methods for field experiments. *American Political Science Review*, 99(2):283–300, 2005.
- Karl W Irons, John DC Litle, and Robert L Klein. Determinants of coupon effectiveness. *Advances and Practices of Marketing Science*, Proceedings of the 1983 ORSA/TIMS Marketing Science Conference (March), 1999.
- Ronald J Johnston and Charles J Pattie. Campaigning and advertising: An evaluation of the components of constituency activism at recent British general elections. *British Journal of Political Science*, 28(4):677–85, 1998.

- Brian Laghi. Dion picks up steam in home stretch. *The Globe and Mail*, page A1, Nov. 28 2006.
- Edward E. Leamer. Let's take the con out of econometrics. *American Economic Review*, 73(1): 31–43, 1983.
- Peter John Loewen. How do local candidates spend their money? And does it matter? Paper presented at the Canadian Political Science Association Annual Meeting, May 2005.
- Rose McDermott. Experimental methodology in political science. *Political Analysis*, 10(4):325–42, 2002.
- R Miller and D Robyn. A field experimental study of direct mail in a congressional primary campaign: What effects last until election day? *Experimental Study of Politics*, 5:1–37, 1975.
- Edward L Nash. *The Direct Marketing Handbook*. McGraw-Hill, New York, 1984.
- Scott A Neslin, Carline Henderson, and John Quelch. Consumer promotions and the acceleration of product purchases. *Marketing Science*, 4(2):147–65, 1985.
- Khayyam Z Paltiel. Party and candidate expenditures in the Canadian general election of 1972. *Canadian Journal of Political Science*, 7(2):341–52, 1974.
- Charles J Pattie, Ronald J Johnston, and Edward A Fieldhouse. Winning the local vote: The effectiveness of constituency campaign spending in Great Britain, 1983-1992. *American Political Science Review*, 89(4):969–83, 1995.
- George Perlin, editor. *Party Democracy in Canada: The Politics of National Party Conventions*. Prentice-Hall Canada Inc., Scarborough, 1988.
- R.V. Lenth. Java applets for power and sample size (computer software). Retrieved February 13, 2007 from <http://www.stat.uiowa.edu/~rlenth/Power>, 2006.
- Ian Shapiro, Rogers Smith, and Tarek Massoud, editors. *Problems and Methods in the Study of Politics*. New York.
- David K Stewart. The changing leadership electorate: An examination of participants in the 1992 Alberta Conservative leadership election. *Canadian Journal of Political Science*, 30(1): 107–28, 1997.
- The Strategic Counsel. Liberal leadership survey, 2006a.
- The Strategic Counsel. Liberal party members survey, 2006b.
- Lynn Vavreck, Constantine J Spiliotes, and Linda L Fowler. The effects of retail politics in the New Hampshire primary. *American Journal of Political Science*, 46(3):595–610, 2002.
- Joseph Wearing. The high cost of high tech: Financing the modern leadership campaign. In Perlin [1988], pages 72–83.
- Paul Wells. *Right Side Up: The Fall of Paul Martin and the Rise of Stephen Harper's New Conservatism*. Douglas Gibson Books, Toronto, 2006.
- Paul F Whiteley and Patrick Seyd. Local party campaigning and electoral mobilization in Britain. *Journal of Politics*, 56(1):242–52, 1994.

APPENDIX A: TREATMENT ASSIGNMENT PROCEDURE

Our treatment assignment procedure occurred in three steps:

- An official list of delegates was provided to campaigns by the Liberal Party of Canada following delegate selection meetings. We first excluded all those who did not have a proper address and then excluded those from three provinces Quebec, Manitoba and British Columbia. Delegates from Quebec were excluded as they were subject to a different ad campaign by the Ignatieff campaign. Those in Manitoba and British Columbia were excluded because delegate lists were not finalized at the time of treatment assignment due to disputes over the eligibility of several delegates. Using the random number generator function in Excel, we assigned each delegate a random number and then ranked delegates from largest to smallest number. The first 800 delegates were selected for the study. We originally included delegates pledged to support Ignatieff in our first sample because we expected a second campaign to participate in the experiment. We included Ignatieff delegates to allow us to test the effectiveness of the second campaign's direct mail on delegates committed to other candidates. Ultimately, the second campaign did not participate, but not before we had sent a treatment schedule to the Ignatieff campaign.
- Among the 800 selected delegates, we identified and excluded all those who were not pledged to support Michael Ignatieff (the leadership selection process of the Liberal party requires those who stands as delegate candidates to formally declare their allegiance prior to delegate selection meetings. This information is retained in the official party list). This left 567 delegates.
- Among the remaining delegates, we assigned them a second random number and ranked them from largest to smallest number. The first 100 delegates were assigned to receive two pieces of mail from the Ignatieff campaign. The next 200 delegates were assigned to receive one piece of mail. The remaining delegates (267) were assigned to receive no mail for the period of the study.

In the course of receiving completed surveys we identified as many as four individuals in our control condition who may have been treated by the campaign. Because the campaign eventually mailed every delegate, those from whom we received completed surveys after November 27th may have received mail from the campaign. However, when we rerun our analysis with these individuals excluded our statistical and substantive results do not change.