

Do we need polls? Why Twitter will not replace opinion surveys, but can complement them.

Javier Sajuria
Department of Political Science
University College London

Jorge Fábrega
School of Government
Universidad Adolfo Ibáñez

Can we observe public attitudes using social media data? Or more concretely, can we disregard polls as the preferred method for observing public opinion? Our comparison between public opinion survey data and Twitter data from the 2013 Chilean presidential election shows a nuanced picture. We use social network analysis to estimate political positions of Twitter users and estimate their sentiment towards public issues based on their public tweets. We work with the public opinion data from the Chilean CEP survey to compare the Twitter data to the polls. Our focus is on the relationship between political position — proxied by their support for a candidate — and their views about given relevant political topics. The results show that, in most cases, support for certain policies has a correlate in the online world, with a positive tone of tweets. However, there are some interesting differences among supporters of various candidates. Those who support the leading candidate tend to tweet with a more positive tone, regardless of the issue. On the other hand, supporters of other candidates are less likely to tweet with a positive tone, even when it is about a topic they support. These findings show that Twitter data can provide interesting insight about how people frame political discussions, depending on the electoral viability of the candidate they support.

Keywords— political position estimation, social media, network homophily, electoral viability

Introduction

Monitoring and using social media to understand — or influence — public opinion is not a new thing. Companies, political parties and organisations alike are keen to observe what their followers say, what people are commenting on their Facebook pages, and what is said in the comments sections of Youtube and Instagram. Moreover, a great deal of work has been done in building social media teams in charge of both engaging and analysing what people exchange through these platforms. To some extent, these phenomena have questioned whether traditional, more expensive, ways to observe public opinion are still required. The regular route for understanding

public opinion, both at the consumer and the political levels, relies heavily on surveys. These instruments present their own advantages depending on the scope of the research. Moreover, they enjoy a fair amount of validity among the scientific community as proper instruments to analyse public attitudes.

Twitter, on the other hand, has been widely contested by the academic community as a valid way to analyse public attitudes and behaviour. Different attempts to predict election outcomes from Twitter have failed, and scholars (Gayo-Avello 2012; DiGrazia et al. 2013) have argued about the usefulness of social media data to understand large-scale political events. The same has been argued in relation to other events, such as the Eurovision contest or popular TV Shows (e.g. The X Factor show). The underlying consensus is that Twitter does not present the conditions required by traditional research approaches to produce accurate forecasts. Hence, some recent attempts have pursued a different route: comparing Twitter data to opinion polls. Some recent efforts (e.g. Beauchamp 2013) aim to forecast candidates' approval ratings by matching them with Twitter data from the previous period. In that way, the goal is no longer in predicting elections (that is left to opinion polls), but to analyse how close are the discussions on Twitter to more valid representations of public opinion.

This chapter aims to expand this line of research by using two different strategies. On the one hand, we use retweet networks to estimate the political position of Twitter users. Second, we take the content of the tweets from those users to compare their views on different topics with data from public opinion surveys. Scholars (Ansolabehere, Rodden, and Snyder 2008; Bartels 2010; Iyengar, Sood, and Lelkes 2012) have already established the presence of a relationship between political positions and attitudes towards public issues, such as equal marriage and the electoral system. We use the estimation from Twitter data to compare the results with opinion polls and provide a more informed picture of when, if possible, social media can substitute or complement them.

We use data from Chile for 2013, focusing in the period before the presidential election. The Twitter data was gathered from 17 September to 17th December of 2013 using the Twitter public streaming API. Survey data, on the other hand, comes from a mainstream source in Chile: The seasonal survey from the Centro de Estudios Públicos (CEP). We discuss the validity issues of each source and the strategy to assess their accuracy.

Our results show that, according to the expectations, Twitter data is still not appropriate for substituting opinion polls. However, there is an interesting story to be told in relation to candidate support and tone. In the Chilean case, the supporters of the leading candidate, and later President, Michelle Bachelet, are more likely to express their views on Twitter with a positive tone. Moreover, when the support for Bachelet predicts significantly the support for certain policies (according to survey data), there is a correspondence in the positive tone of the tweets from her Twitter supporters. This is something that does not happen in the case of those who support other candidates. In other cases, we see supporters that tend to be less likely to use a

positive tone on their tweets, even when survey data says that, on average, they support the policies that are talking about. This is consistent on our hypothesis that electoral viability is related to higher likelihood of a positive tone in the tweets.

This chapter will go as follows. First, we discuss the literature on opinion formation and the role of political position. Then, we move into the discussion of using social media data to forecast political events and understand public attitudes. This is followed by a discussion of the Chilean case and the elements of electoral viability. We then explain the methods used to estimate political positions and filter the relevant topics. The results are presented to demonstrate how we derive the conclusions stated above. In our discussion, we extrapolate from these results to make a compelling case of how much researchers should rely on these sources and what is the actual potential of new media for valid academic research.

Literature Review

Ideology and Public Opinion

Political ideology is relevant for public opinion. This is a bold statement, but not unjustified. Zaller argued in 1992 in his well-known RAS model that people who are more politically aware tend to have more stable and defined attitudes. For Zaller, ideology was a product of this awareness. The more aware a person is, the more stable are their ideological positions. Then, people with more defined ideology or systems of belief will look for information from partisan voices. That is, liberals will search for opinions from liberal elites, and will reinforce their own liberal views. This will reflect on preferences for public policies (such as redistribution or welfare in the case of liberals), and approval ratings. Conversely, Zaller claims that less aware people have also less stable ideological positions. In turn, this will reduce their “attitude constraints” creating inconsistency. In short, more ideology leads to more consistent attitudes.

The empirical evidence supports this view. For example, Bartels (2005) studies how views on tax reform in the US are explained by ideology and levels of education. Converse (1975) makes a similar case in relation with voting behaviour, while Dalton (2000) analyses the role of party identification (usually used as a measure for ideological position) in today’s politics. Outside the US, there is growing body of literature on the topic. López-Sáez & Martínez-Rubio (2005) explain how ideological positions change the level of credibility in governmental information. Based on the case of the 11-M terrorist attacks in Madrid, they found that right-wing people believed that their voting behaviour had been affected more by official information, while left-wing respondents were more influenced by unofficial information.

Estimating people’s political position is, then, extremely relevant to understand political attitudes. Traditional measures have relied on survey questions where people can position themselves in 1-10 (or 1-7) scales, indicate their preferences for the existing political parties, or self-identify as liberals or conservatives. The validity

of each of these measures has been widely discussed (see discussion on Ansolabehere and Hersh 2012) and is usually dependent on the political system of each country. In multiparty systems, measures of left-right scale might be an over-reduction of the complexity in which people can position politically. Accordingly, countries with two-party systems are more suitable for such scales. In the case of Chile — the case under study — the presence of two big coalitions for the last 20 years allows us to use methods that are similar to two-party systems.

Using Social media for forecasting elections and understanding public attitudes

Since Nate Silver's fairly accurate predictions of the last US elections, forecasting events has become an attractive topic. Nowadays, we can find statistical models to predict the outcome of the Fifa World Cup, the winner of the Eurovision contest, or the next armed conflict in the world. With unequal results, the advancements of forecasting models rely heavily on the quality of the information they use to base their predictions. For examples, attempts to produce similar election forecasts in Chile (Bunker n.d.) have failed, mainly due to the inability of pollsters to estimate turnout. Similar situations can be observed in other Latin American countries, where the low quality of survey data produces bad forecasts.

A similar discussion has been taking place in academic spheres studying social media. Daniel Gayo-Avello's (2012) now-seminal piece on the topic explained, in a very crude way, why most attempts to provide electoral predictions using Twitter have failed. In Gayo-Avello's opinion, there are no real studies predicting elections, but they fit their models against past events (he calls them *post-hoc* analysis). This practice creates a large number of successful "predictions", but there are big questions about how these results would actually fit with future events. Another criticism is that there is no clarity on how to equate tweets to votes. In that regards, a recent study (DiGrazia et al. 2013) attempted to correlate the raw count of tweets with the raw count of votes, although this approach has been heavily questioned for its validity. In general, criticisms about prediction attempts are wide and well-grounded.

The field has now started to move away from electoral outcomes to compare Twitter data with some indicators from public opinion surveys. The substantive grounds for the move are clear. The same as responses in a survey, tweets can be understood as public expressions of private attitudes. On the other hand, voting is a different kind of action, more complex, which is constituted by a large array of preferences, attitudes and motivations. The relationship between tweets and survey responses, hence, is more justifiable and sensible. Both respond to a process of opinion formation and thus, can be compared as units of observation of a similar nature.

On another note, information from Twitter is relevant as much as any other elite source is. As discussed above, Zaller's RAS model relies on the notion that people who don't hold strong views about certain topics will rely on whatever is more salient to them at the moment they need to elicit a response. In that way, priming plays a significant effect. People can be primed by the media, through pundits, news and other

means. But they can also use trusted sources of information, usually, elites. Elite discourse is key in the process of opinion formation, as it influences people's opinions by reinforcing their preconceptions and affecting the level of salience of certain information.

Nick Beauchamp (2013) has inaugurated the field by using Twitter data to observe how it can predict the popularity of different candidates in state-level surveys. His results show how can we use Twitter to interpolate the text from the tweets to indicators of candidates' popularity. His preliminary findings have clear implications for cases when accurate polling data is not easily available.

Our approach follows from this logic. If we approach tweets as expressions of personal attitudes, we can use them to understand key elements of public opinion. Moreover, there is more in Twitter than tweets. For example, Barberá (2015) has been successful in estimating users' political positions based on who they follow. We take a slightly different approach. As we explain below, instead of using following networks, we prefer retweets. We take advantage of the current research on the level of homophily of these networks (Lietz et al. 2014) to estimate the political position of those who participate in them. Our preference for retweet networks stems from the substantive difference between following someone and retweeting that person. Although both actions can be motivated by ironic or sarcastic motives, one is a one-off, relatively costless and private action (following), while the other exposes the user's timeline to someone else's tweets. The action behind it reflects a more active statement, whatever that is.

Another way to obtain public attitudes from social media is through text analysis. The use of text as data is a growing area of work in political research. From scaling text to obtain political positions (Lowe 2008; Lowe 2013; Laver, Benoit, and Garry 2003) to the use of topic modeling for understanding representatives' responsiveness on social media, the field has grown significantly in the last decade. We are interested in a particular area, sentiment analysis. The purpose of this approach is to use semi-automated methods to understand the polarity of certain sections of texts. By polarity, we mean the position of the text in terms of positive or negative dimensions. The use of sentiment analysis is also well documented (Godbole, Srinivasaiah, and Skiena 2007; Pang and Lee 2008) and it has been applied to Twitter.

This paper takes advantage of this literature and aims to advance it by looking beyond candidates. The field has been mildly obsessed with forecasting candidate support, while much has yet to be said about policy positions.

The case of Chile in the context of electoral viability

The 2013 presidential election in Chile was particularly different from previous opportunities, at least since the end of Pinochet's dictatorship in 1990. First, it was the end of the first democratic right-wing government in over 50 years. President Piñera's administration had won the election in 2009, and was facing a difficult situation in terms of public approval. Moreover, his coalition suffered from the resignation of their

candidate due to mental health concerns. A new candidate, Evelyn Matthei, former minister of Labour, had to take the role and agreed to compete.

Second, the candidate from the left-wing coalition, Michelle Bachelet, was the front-runner by a large margin. The CEP survey showed a support of 47%, while Ms Matthei, her closest contender, only reached 14%. Ms Bachelet left the Presidency in 2010 with arguably the highest approval ratings in Chile's history. While she remained absent from Chilean politics, she was always in the lead of the opinion polls.

Third, each sector had other interesting, yet unlikely-to-win candidates. On the one hand, former socialist Marco Enríquez-Ominami was pursuing the presidency for the second time, after obtaining a respectable 20% in 2009. In 2013, however, his share was half of that. From the right wing, an independent, Franco Parisi, was also gaining strong support, risking Matthei's options. The CEP survey showed a support of 10%, which is close to his final share in the election.

Fourth, there was a group of 5 other candidates, which raised the number to 9, the highest in Chile's recent history. None of them got more than 5% of the vote share, nevertheless.

We understand electoral viability as the perceived capability of a candidate of winning an election, and this is an attribute that only Michelle Bachelet had. According to the survey data from CEP, 78.2% of the respondents believed that Ms Bachelet would become the next president, while only 5.2% believed that Ms Matthei had a chance. None of the remaining 7 candidates got more than 1% of the mentions. In a scenario like this, the low level of competition should be reflected in the discussions among supporters of each candidate. That is, supporters of Michelle Bachelet should be more confident, and eventually more positive, when engaging in political discussions in the time prior to the election. There are no motives for them to engage in negativity.

Hypotheses

Based on the theory discussed above, we aim to test the following three hypotheses,

Hypothesis 1: Consistent with the expectations of the literature, the share of supporters of each candidate on Twitter will not be useful to predict the actual share on the survey.

The main argument in support of this hypothesis lays in the lack of probabilistic sampling of Twitter data. Although the information obtained through Twitter's public streaming API is a sample of 1% of the total, research (Morstatter et al. 2013) has shown that it does not consist in a probabilistic random sample. Any result we get from our estimation is likely to differ from the results of the opinion survey.

Hypothesis 2: Michelle Bachelet's Twitter supporters are more likely to use a positive tone when tweeting about political issues.

This hypothesis stems from the electoral viability argument. Supporters of Bachelet are in no need to engage in violent, negative discussions on Twitter. They are the most

likely winners of the forthcoming election, and even the supporters of the other candidates recognise that.

Hypothesis 3: Where there is a statistically significant relationship in the survey data between supporting a candidate and supporting a given policy, that relationship is reflected in a higher probability of a positive tone on Twitter.

This hypothesis consists on the idea that support for a policy can influence the way in which users who support the same candidate talk about the policy issue. We do not expect a significant difference among candidates, as the same argument should run for all of them.

Data and Methods

Data Sources

Our data comes from Twitter's public Streaming API according with two criteria. On one hand, during the second half of October 2013 (two weeks before the election), we collected twitter accounts who explicitly declared their intention to vote for a particular presidential candidate using hashtags with the following structure #yovotoZZZ (meaning "I will vote for ZZZ") where ZZZ represents the name or last name of one of the four main candidates: Michelle Bachelet, Evelyn Matthei, Marco Enriquez and Franco Parisi. Each tweet contained all the relevant meta-data, such as location, URLs, date of creation and the users who tweeted them. This filter allowed us to identify the political preferences of 4,111 accounts.

On the other hand, we collected tweets that were posted between the months of September to December 2013 in Chile using any of the following concepts that were relevant during the presidential campaign. Each of them were topics in which candidates showed polarized positions: equal marriage ("matrimonio igualitario", "#avp"), changes to the electoral system ("binominal"), abortion ("aborto"), copper ("codelco", "nacionalizaci.n del cobre"), and constitutional changes ("asambleaconstituyente*", "#ac"). To validate the method, we also included a search on a topic for which we do not expect polarized opinions based on political preferences: tweets about the Chilean national football team manager ("sampaoli").

To guarantee that tweets were written in Chile we filtered the collected data by its location (either self reported location or geographical coordinates when they were available). By means of these filters we obtained 152,240 messages including tweets and their retweets. Finally, from the above collection we selected all the users' accounts who retweeted messages from any of the politically identified users. This allowed us to expand the original dataset from 4,111 to 5,603 accounts. The steps followed to assign political preference to those retweeters are explained in the next subsection.

The survey data comes from the quarterly survey from the Centro de Estudios Públicos (CEP), a recognised Chilean think tank. The survey was conducted during

13th September and 14th October 2013. This is a nationally representative survey with a sample of 1,437 respondents. The margin of error is 3.0%. All calculations were conducted using the weighting instructions from CEP.

We have chosen 5 different topics that were asked in the CEP survey as dependent variables: constitutional reform, equal marriage, electoral reform, abortion, and the ownership structure of the copper mines. Constitutional reform relates to a campaign started in 2013 to establish a constitutional assembly in the country. In the case of electoral reform, this refers mostly to the different attempts to change the way in which the members of Congress are elected. Equal marriage was a hot topic during the campaign, with most candidates supporting some sort of legal protection for same-sex couples. The case of abortion is particular, as Chile is one of the few countries in the world that does not allow abortion under any circumstances. As such, some of the candidates, including Bachelet, showed their support for allowing abortion under certain circumstances. Finally, Chile's copper mine create the greatest income of the country. Prior to the dictatorship, all the operation of the mines was under the state-owned company Codelco. There are some political actors (and some of the candidates) promoting that the state should re-gain control of the mines.

Estimating political positions from Retweet Networks

Let S represents a set of individuals who explicitly declared their intention to vote for some candidate in the set C . Such that S_c represents the subset of individuals in S supporting a given candidate $c \in C$

Let M be the set of tweets written by individuals in S and M_c be the subset of messages written by each individual $s \in S_c$ supporting the candidate $c \in C$.

Let R represent another set of individuals who have retweeted at least one message written by individuals in S . Then, for each individual $r \in R$,

$$r_c = \text{count}(M_c)$$

and

$$r_{\text{total}} = \sum c r_c$$

Then,

$$r_{wc} = \frac{r_c}{r_{\text{total}}}$$

is the weight of individuals in S_c within the set of individual r 's retweets.

Finally, we defined r_p as individual r 's political position as,

$$r_p = \max(r_{wc}) | r_p > 0.5$$

Figure shows the results for $c =$ Bachelet, Matthei, Parisi, Enriquez and the data specified in the previous subsection. As shown, for all candidates and controversial

issues retweeters predominantly retweet messages from individuals with political affinity giving support for the claim that retweeting on political issues is mostly a homophily-like behavior.

Calculating the sentiment behind the tweets

For the purpose of this exercise, we rely on an unsupervised method to calculate sentiment analysis. Unsupervised methods use several ways to compare the text with some baseline lexicons or dictionaries, and assigns scores to the different units of observations.

In this case, we have used our own Spanish translation from the lexicon created by Wilson et al. (2005), and the sentiment R packages (Jurka 2012; Sajuria 2014). The original English lexicon uses a trained naive Bayes algorithm to classify the polarity of the words between negative and positive. Since this is a preliminary exercise, instead of re-training the algorithm, we use the voter algorithm. Each tweet is deconstructed into single words, and they are compared to the lexicon. The algorithm then counts the number of positives and negative words in each tweet. Whenever the majority of words are positive, the whole tweet is classified as such. The same procedure operates for words classified as negative. When the words that do not have a pre-set polarity are a majority, the algorithm classifies tweet as neutral. Subsequently, we created a dummy variable called “positive” for each tweet, with value of 1 if the classification of that tweet is such, and 0 otherwise.

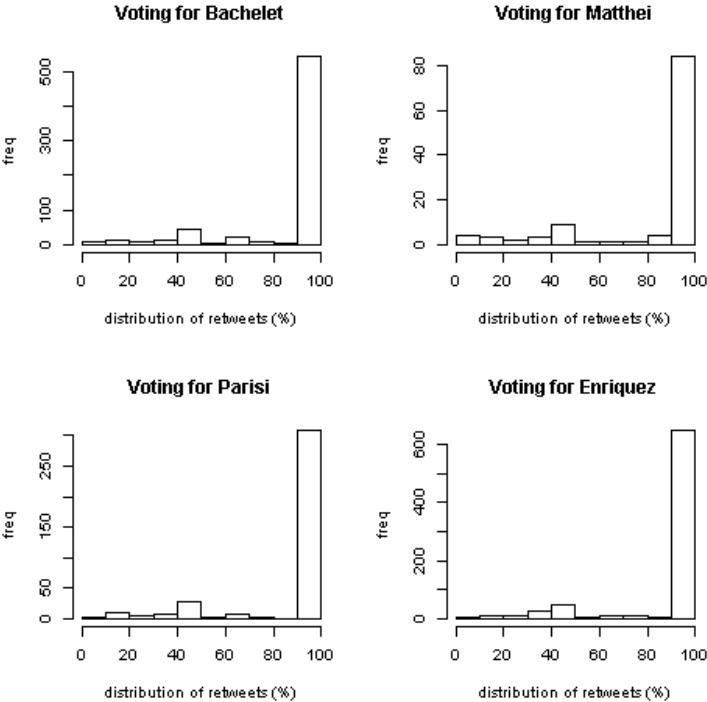


Figure 1: Retweeters’ distribution of their retweets by political affinity. The x-axis shows that there is a high homophily level on each of the retweet networks.

Modelling support for candidates on tone and support for policies

As Barberá (2015) explains, electoral support should be orthogonal to support on Twitter for a given candidate. However, we are departing from that discussion in two ways. First, we are looking at support for policies, which are not subject to a vote. Moreover, we aim to compare if the support for a candidate can be used as predictor for both support for a given policy, and the sentiment of those talking about that policy. Therefore, we need to estimate different models for each dataset.

For the Twitter data, we model the positive tone of tweets about a given policy (positive), expressed as a dummy variable, as a function of the support for a given candidate (candidate), the sex of the respondent (sex), and whether the respondent lives in Santiago or somewhere else in Chile (santiago). The last two variables are traditional controls that can be used when understanding political support. We developed our own algorithm to detect the sex of the Twitter user and its location (when not geotagged). The link function is a logistic regression, and can be expressed as follows:

$$p(\text{positive} = 1) = \frac{1}{1 + \exp(-\beta_0 + \beta_1 \text{candidate}_1 + \beta_2 \text{sex}_2 + \beta_3 \text{santiago}_3 + \epsilon)}$$

In the case of the survey data, the dependent variable changes for support (support) for a policy. The independent variables remain the same and the model has a similar expression,

$$p(\text{support} = 1) = \frac{1}{1 + \exp(-\beta_0 + \beta_1 \text{candidate}_1 + \beta_2 \text{sex}_2 + \beta_3 \text{santiago}_3 + \epsilon)}$$

In the case of each datasets, we estimated five models, one for each of the policies. The 'candidates' variable only considers 4 contenders, although the last Chilean election had 9. We are focusing on those who got the highest share on the first round: Michelle Bachelet (who went onto the run-off election and became finally elected), Evelyn Matthei (the other candidate that went onto the run-off election), Marco Enríquez-Ominami, and Franco Parisi. Bachelet and Enríquez-Ominami are considered centre-left/left wing candidates, while the other two are centre-right/right. Together, these four candidates obtained 92.82% of the votes during the first round of the presidential election.

Results

In relation to the share of supporters for each candidate, our estimation on Twitter should be independent to the support shown on the survey. Table 1 shows the identification or support for all the candidates included in our models per dataset. The results provide support for Hypothesis 1, as we can see no evident connection between the share of those who support any given candidate on the survey and those who we estimate on Twitter.

Table 1: Support for candidates

	Twitter	CEP Survey
Bachelet	39%	47%
Matthei	4%	14%
Enríquez-Ominami	32%	7%
Parisi	17%	10%

Table 2 shows the estimation results using the CEP survey dataset. In essence, we can observe that the support for certain candidates can act as a predictor of the support for certain reforms. In particular, the support for candidates (any) seems to be positively related with support for some of the policies. Women, on the other hand, tend to be more supportive of equal marriage than men, and less supportive of electoral reform. People living in Santiago are also more supportive of equal marriage, but have a negative likelihood to support abortion law, and transferring the ownership of the copper mines back to the State.

In the case of the Twitter data, the picture looks a bit different. As explained above, our data does not come from a probabilistic sample. As such, the relevance of statistical significance is lower. Without a random sample, we cannot assume that our estimations will reflect the larger population of users on Twitter. Hence, we focus more on the direction of the coefficients, and less on the standard errors. The results shown in Table 3 express the coefficients from the logistic regression on the probability of a tweet having a positive tone.

Table 2: Logistic regression using CEP survey data

	<i>Dependent variable:</i>				
	Constitutional change	Equal Marriage	Electoral Reform	Abortion	Copper ownership
	(1)	(2)	(3)	(4)	(5)
Other Candidates	2.056*** (0.441)	2.130*** (0.424)	1.866*** (0.485)	1.459** (0.623)	3.392*** (1.028)
Evelyn Matthei	0.310 (0.269)	0.206 (0.296)	0.884*** (0.277)	0.582** (0.287)	0.403 (0.331)
Franco Parisi	0.596* (0.355)	0.516 (0.326)	0.734* (0.386)	1.045*** (0.361)	0.979** (0.402)
Marco Enríquez-Ominami	0.776** (0.323)	1.293*** (0.342)	0.849*** (0.328)	1.404*** (0.366)	1.564*** (0.498)
Michelle Bachelet	0.311 (0.215)	0.259 (0.223)	0.553** (0.224)	0.366* (0.218)	0.748*** (0.263)
Sex (women)	-0.297* (0.162)	0.531*** (0.170)	-0.302* (0.167)	0.239 (0.175)	0.102 (0.208)
Santiago (location)	0.212 (0.162)	0.353** (0.167)	0.190 (0.164)	-0.242 (0.179)	-0.412** (0.208)
Constant	-0.516** (0.216)	-1.381*** (0.230)	-0.474** (0.233)	0.490** (0.226)	1.053*** (0.283)
Observations	1,437	1,437	1,437	1,437	1,437
Log Likelihood	-918.345	-856.686	-919.870	-765.994	-594.736
Akaike Inf. Crit.	1,852.689	1,729.371	1,855.741	1,547.987	1,205.473

Note:

*p<0.1; **p<0.05; ***p<0.01

As predicted, supporters of Michelle Bachelet are more likely to tweet with a positive tone on almost all the models. The only exception consists on the model about equal marriage. In the case of supporters for the other candidates, the tones differ across the models. Women are less likely to use a positive tone on every one of the topics than men, while people tweeting from Santiago are more likely to use a positive one.

Table 4 shows no clear support for hypothesis 3. Only in the case of supporters of Michelle Bachelet there is a consistency between the support for an issue in the survey and the occurrence of a positive tone on Twitter. That is the case for the topics of electoral reform, abortion law, and the ownership of the copper mines. Interestingly, supporters of Mr Enríquez-Ominami show some decoupling on this regard. Their tone when discussing equal marriage, electoral reform or abortion is not positive, even when the survey data shows that, in average, they are more likely to support these issues. A similar case takes place among supporters of Matthei and electoral reform, and supporters of Parisi and constitutional reform.

Table 3: Logistic regressions using Twitter data

	<i>Dependent variable:</i>				
	pol.num				
	Constitutional change	Equal Marriage	Electoral Reform	Abortion	Copper ownership
	(1)	(2)	(3)	(4)	(5)
Michelle Bachelet	0.395 (0.307)	-0.381 (0.554)	0.447 (1.066)	0.069 (0.325)	16.128 (1,348.160)
Evelyn Matthei	0.970** (0.446)	-0.118 (0.794)	-14.537 (1,135.237)	0.386 (0.467)	18.038 (1,348.160)
Marco Enríquez-Ominami	0.018 (0.316)	-0.252 (0.561)	-0.215 (1.112)	-0.053 (0.363)	14.826 (1,348.160)
Franco Parisi	-0.559 (0.436)	-0.147 (0.975)	0.230 (1.213)	0.323 (0.460)	15.602 (1,348.160)
Sex (women)	-0.232 (0.157)	-0.232 (0.290)	-0.466 (0.572)	-0.330 (0.247)	-1.015 (0.944)
Location (Santiago)	0.204 (0.150)	0.328 (0.266)	0.373 (0.435)	-0.198 (0.248)	0.858 (0.627)
Constant	-2.409*** (0.305)	-1.106** (0.554)	-3.162*** (1.044)	-1.647*** (0.319)	-18.017 (1,348.160)
Observations	2,195	434	478	696	137
Log Likelihood	-716.758	-230.063	-98.071	-274.456	-38.025
Akaike Inf. Crit.	1,447.515	474.126	210.143	562.912	90.051

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 4: Comparison between support and positive tone

Light grey: Increase in probability; Dark grey: Decrease in probability

	Constitution		Equal Marriage		Electoral reform		Abortion Law		Copper owner	
	Survey	Twitter	Survey	Twitter	Survey	Twitter	Survey	Twitter	Survey	Twitter
Michelle Bachelet					*		*		*	
Evelyn Matthei					*		*			
Marco Enríquez-Ominami	*		*		*		*		*	
Franco Parisi	*				*		*			
Sex (women)	*		*		*					
Location (Santiago)			*						*	

* p<0.05 (only for survey data).

Discussion

This exercise has been an attempt to expand the field of understanding public opinion through social media data. Much of the discussion has been exploring the notion that Twitter, and other social media, can be a useful — and cheaper — alternative to public opinion surveys. As shown above, there have been some attempts both to analyse Twitter data and compare it to polls. Our approach is an extension of that literature, by incorporating two new elements. On the one hand, we use an innovative way to estimate the political position of Twitter users, by looking at their support for candidates. On the other hand, we depart from the traditional question of relating candidate support on Twitter with similar indicators from surveys. We believe that, given the special nature of this social media platform — such as its non-representative population of users — support for a candidate on Twitter should be orthogonal to similar measures from probabilistic samples. Our results are an initial confirmation of that.

Another contribution of this chapter comes from the notion of electoral viability. We have used that framework to propose that supporters of the leading candidate are more likely to use a positive tone on Twitter. Our results show a preliminary support for this, but they should be taken carefully. There are still some validation tests to perform, in order to provide error measurement to this relationship. However, this is a promising initial attempt.

With regards to the comparison between survey and Twitter data, our results are inconclusive. The supporters of one particular candidate, Enríquez-Ominami, are less likely to use a positive tone on Twitter when discussing topics they support. This is surprising, and might reflect deeper dynamics of public discussion. An option for future research could be to focus on this particular case and understand what are the drivers of this phenomenon. One option is electoral viability — Enríquez-Ominami suffered throughout the campaign from a lack of momentum, especially compared to his performance in 2009 — as a driver of tone. However, given the impressive advantage that Ms Bachelet had on this regard, it becomes difficult to model this assertion with the data we have available.

A word of caution should be given with regards to our estimation of political positions. As we have discussed before, there are other ways to estimate political positions of Twitter users that use followers-following relationships. We believe that there is something intrinsic in retweet networks that make them more useful in assessing if a given user holds a clear political position. While following another user is a rather costless and private act, retweeting is none of the above. However, we also know that revealed preferences are not necessarily the same as the real. As such, we still need to develop a way to test our estimations against well-respected procedures.

Finally, sentiment analysis is not free from criticism. Machine learning processes, or even counting algorithms such as the one we use, are incapable of understanding sarcasm and irony. Furthermore, in the case of our paper, there are not many tools available to produce sentiment analysis in Spanish. Hence, the translation we

produced for this chapter relies on the accuracy of the English lexicon. Further developments on this point would consist of training some sort of naive Bayesian classifier using only tweets in a language different than English.

In summary, our goal is to propose new roads for future investigation, based on the notion that social media interactions can produce meaningful, complementary information to opinion polls.

References

- Ansolabehere, Stephen and Eitan Hersh (2012). "Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate." In: *Political Analysis* 20.4.
- Ansolabehere, Stephen, Jonathan Rodden, and James M Snyder (2008). "The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting". In: *American Political Science Review* 102.02, pp. 215–232.
- Barberá, Pablo (2015). "Birds of the Same Feather Tweet Together. Bayesian Ideal Point Estimation Using Twitter Data". In: *Political Analysis*.
- Barberá, Pablo et al. (2013). *Is There Anybody Out There? The Effects of Legislators' Communication with their Constituents*. Tech. rep.
- Bartels, Larry M (2005). "Homer gets a tax cut: Inequality and public policy in the American mind". In: *Perspectives on Politics* 3.01, pp. 15–31.
- (2010). "The study of electoral behavior". In: *The Oxford handbook of American elections and political behavior*, pp. 239–261.
- Beauchamp, Nick (2013). "Predicting and interpolating state-level polling using twitter textual data". In: *Meeting on automated text analysis*, London School of Economics, London.
- Bunker, Kenneth. *Tresquintos: Análisis Políticos y Predicciones Electorales*. www.tresquintos.com. Accessed: 2014-08-01.
- Converse, Philip E (1975). "Public opinion and voting behavior". In: *Handbook of political science* 4, pp. 75–169.
- Dalton, Russell J (2000). *The decline of party identification*. Oxford University Press.
- DiGrazia, Joseph et al. (2013). "More tweets, more votes: Social media as a quantitative indicator of political behavior". In: *PloS one* 8.11, e79449.
- Gayo-Avello, Daniel (2012). "' I Wanted to Predict Elections with Twitter and all I got was this Lousy Paper'—A Balanced Survey on Election Prediction using Twitter Data". In: *arXiv preprint arXiv:1204.6441*.

- Godbole, Namrata, Manja Srinivasaiah, and Steven Skiena (2007). "Large-Scale Sentiment Analysis for News and Blogs." In: ICWSM 7.
- Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes (2012). "Affect, not ideology a social identity perspective on polarization". In: Public Opinion Quarterly 76.3, pp. 405–431.
- Jurka, Timothy P. (2012). sentiment: Tools for Sentiment Analysis. R package version 0.2. url: <http://CRAN.R-project.org/package=sentiment>.
- Laver, Michael, Kenneth Benoit, and John Garry (2003). "Extracting policy positions from political texts using words as data". In: American Political Science Review 97.02, pp. 311–331.
- Lietz, Haiko et al. (2014). "When Politicians Talk: Assessing Online Conversational Practices of Political Parties on Twitter". In: CoRR abs/1405.6824.
- López-Sáez, Mercedes and José-Luis Martínez-Rubio (2005). "¿Influyeron los procesos de comunicación sobre los sucesos del 11-M en las votaciones del 14-M? La percepción de los jóvenes en función de su ideología política". In: Revista de Psicología Social 20.3, pp. 351–367.
- Lowe, Will (2008). "Understanding wordscores". In: Political Analysis 16.4, pp. 356–371.
- (2013). "There's (Basically) Only One Way to Do it". In: Available at SSRN. Morstatter, Fred et al. (2013). "Is the sample good enough? comparing data from twitter's streaming api with twitter's firehose". In: Proceedings of ICWSM.
- Pak, Alexander and Patrick Paroubek (2010). "Twitter as a Corpus for Sentiment Analysis and Opinion Mining." In: LREC.
- Pang, Bo and Lillian Lee (2008). "Opinion mining and sentiment analysis". In: Foundations and trends in information retrieval 2.1-2, pp. 1–135.
- Sajuria, Javier (2014). sentimiento: Package for sentiment analysis in Spanish [beta]. R package version 0.1.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffmann (2005). "Recognizing contextual polarity in phrase-level sentiment analysis". In: Proceedings of the conference on human language technology and empirical methods in natural language processing. Association for Computational Linguistics, pp. 347–354.
- Zaller, John (1992). The nature and origins of mass opinion. Cambridge university press.