

Deception in Speeches of Candidates for Public Office

D.B. Skillicorn, School of Computing, Queen's University, Kingston, Canada

skill@cs.queensu.ca, +1 613 533 6065

C. Leuprecht, Department of Political Science, Royal Military College of Canada, Kingston, Canada

christian.leuprecht@rmc.ca, +1 613 541-6000 ext. 6428

ABSTRACT

The contribution of this article is twofold: the adaptation and application of models of deception from psychology, combined with data-mining techniques, to the text of speeches given by candidates in the 2008 U.S. presidential election; and the observation of both short-term and medium-term differences in the levels of deception, primarily persona deception, the presentation of being a different, better person than one might actually be. Rather than considering the effect of deception on voters, deception is used as a lens through which to observe the self-perceptions of candidates and campaigns. The method of analysis is fully automated and requires no human coding, and so can be applied to many other domains in a straightforward way. The authors posit explanations for the observed variation in terms of a dynamic tension between the goals of campaigns at each moment in time, for example gaps between their view of the candidate's persona and the persona expected for the position; and the difficulties of crafting and sustaining a persona, for example, the cognitive cost and the need for apparent continuity with past actions and perceptions. The changes in the resulting balance provide a new channel by which to understand the drivers of political campaigning, a channel that is hard to manipulate because its markers are created subconsciously.

Deception in Speeches of Candidates for Public Office

ABSTRACT

The contribution of this article is twofold: the adaptation and application of models of deception from psychology, combined with data-mining techniques, to the text of speeches given by candidates in the 2008 U.S. presidential election; and the observation of both short-term and medium-term differences in the levels of deception, primarily persona deception, the presentation of being a different, better person than one might actually be. Rather than considering the effect of deception on voters, deception is used as a lens through which to observe the self-perceptions of candidates and campaigns. The method of analysis is fully automated and requires no human coding, and so can be applied to many other domains in a straightforward way. The authors posit explanations for the observed variation in terms of a dynamic tension between the goals of campaigns at each moment in time, for example gaps between their view of the candidate's persona and the persona expected for the position; and the difficulties of crafting and sustaining a persona, for example, the cognitive cost and the need for apparent continuity with past actions and perceptions. The changes in the resulting balance provide a new channel by which to understand the drivers of political campaigning, a channel that is hard to manipulate because its markers are created subconsciously.

Keywords: corpus analytics, political discourse, U.S. presidential elections, deception

Introduction

Political speeches are a campaign's most direct channel for influencing voters. They provide an important opportunity both to convey attractive ideas and policies, and to present an attractive and competent candidate. In other words, to appear more like the ideal electoral choice for the office concerned and better than the other candidates, campaigns appeal to voters' desires, perhaps more than presenting policy contrasts. There is considerable motivation to use deception, across a full range of communication including direct factual deception, evasion, spin (presenting agreed-upon facts in the best possible light), non-responsiveness, and negative advertising. However, in this article we consider a different, and thus far largely uninvestigated, form of deception, which we call persona deception. Persona deception describes the way in which candidates and their campaign staff present themselves as 'better', across multiple dimensions, than they actually are – more experienced, more able, more knowledgeable, and so on. Although it is possible to present oneself as better than one is intentionally, psychological research suggests that there is a parallel, predominantly subconscious, process driven by an individual's self-perception – and as such not subject to deliberate manipulation.

This approach is the first compound application to politics of two developments in other disciplines. The first is the discovery by psychologists of an internal state channel, carried by the small, functional words in utterances, that reveals much of the (subconscious) state of the

speaker/writer. This includes deception (our focus here), status in interactions, personality, and even health (Chung & Pennebaker, 2007). The second is the application of data-analysis techniques using matrix decompositions, particularly singular value decomposition, that allows more sophisticated factor analysis (Skillicorn, 2007). Dimensional analysis of word frequencies has a tradition in political science (Slapin & Proksch, 2008; Lowe, 2008; Laver, Benoit, & Garry, 2003). These new data-analysis techniques open up ways to analyze, more deeply, what lies behind the attitudes and actions of individuals in many political settings. If deception is detectable, how might it relate to the candidate running for public office and the staff running the candidate's campaign? Is it just a function of personality? Or is there a pattern whereby subconscious deception is greater when speaking to certain kinds of audiences or dealing with certain types of policy issues? The novelty in the approach taken by this article, then, is deception as revelatory of the self-perceived state of each campaign and their collective view of the candidate pool, rather than the effect of deception on a target audience of potential voters.

Using computational data-mining techniques to analyze the large number of observations generated by candidates' speeches throughout a campaign, the article develops a model to theorize aspects of persona deception. The article assesses the model using data from the 2008 U.S. presidential campaign's three eventual contenders: Hillary Clinton, John McCain, and Barack Obama. There is precedent for the comparative analysis of speeches by presidential contenders (e.g. Corcoran, 1991). This particular setting was chosen because of its inherent high profile; the rarity of a presidential election without an incumbent; and the ready availability of textual data. The model has two key advantages over much of the literature dealing with slant and bias: text-based analysis obviates the difficult, cumbersome, and inherently subjective process of coding speech; and the large number of data points makes it possible to identify patterns across multiple times scales. Levels of persona deception among the three campaigns differ considerably, and all show variations over medium and even quite short time scales. The variations *per se* produce interesting results insofar as they show distinct patterns in areas where we did not anticipate them and no patterns in areas where one might reasonably have anticipated a relationship. Several intuitively plausible drivers, such as low levels of persona deception for speeches given to friendly audiences, are not supported by the evidence. More subtle associations, however, such as the tactical purpose of each speech, do appear to be associated with levels of persona deception in predictable ways.

Since speech is a largely subconscious process, variation allows for fairly unmediated observations of each candidate and campaign at any given moment in time. Evidence is growing that this remains true for speeches that are constructed with great care and by many participants, since overt attempts to manipulate tone founder on the subconscious nature of even written language; and those involved in a campaign surely live within a tightly shared subculture.

The article posits levels of persona deception as the function of a complex and dynamic balance between a desire for higher levels of deception (because it is effective), the limitations of the emotional and cognitive cost of deception, and the difficulty of maintaining, internally and

externally, a consistent persona for different kinds of supporters. Whereas the evidence supports a consistent correlation between levels of deception across competing campaigns, the relationship between deception and actual success at the polls is indeterminate. Notwithstanding the limitations inherent to our case study, we believe that the prospect of replicating, testing, and applying our method to an almost unlimited array of electoral campaigns, with the sole empirical limitation being the availability of machine-readable speeches, holds considerable promise for social scientists interested in the study of elections, and indeed of other forums such as legislatures. In principle, the model is applicable to a wide variety of settings in the social sciences, and the ability to replicate the method in other settings is limited only by the number of available data points.

Deception

Political science normally treats campaigns as a means to disseminate policy positions in a Downsian type of political competition. In a Single Member Plurality electoral system only one candidate can carry an electoral district. To compete successfully (measured as a function of getting re/elected) in this type of political marketplace, politicians have to attract votes across a wide spectrum of the electorate. To this end, they have to appeal to highly diverse audiences on a broad array of policy issues. In doing so, candidates for public office have to overcome sizeable hurdles: Politicians naturally have a greater affinity with some audiences than with others, they are more familiar with some policy areas than others, and they have to sell audiences on policy solutions to problems that are largely beyond their control. The more populous and diverse the constituency, the higher the transaction costs.

To get elected, then, candidates for public office have to be strategic about reducing transaction costs. Deception is one such strategy: present a different persona to different audiences, convey a better grasp of a given policy area than one really has, or convince audiences that one can make an impact in an area that transcends a candidate's sovereign jurisdiction. Just how difficult it is to reconcile these competing demands is exemplified by the issue-image distinction that the functional theory of campaign discourse (Benoit, 2007) draws between candidates' position on policy issues and the persona they try to impress on their audience (Hacker, Zakahi, Giles, & McQuitty, 2000; Stuckey & Antczwk, 1995).

The differences between what is said or written by an individual and the actual views of that individual about a situation can be placed on a spectrum. At one end are factual discrepancies: something is said that is objectively untrue, and the individual involved knows it. At the other end are discrepancies that are related to character: what is said presents someone, often the individual making the statement, as different from who they actually are. Deception also varies in severity, from outright lying, which can attract legal penalties, to negotiation, where there is an expectation by both parties that neither is presenting a complete and accurate picture of their goals and constraints.

The term spin is often used to describe the attempt to present the best possible interpretation of facts on which there is broad consensus (e.g. Alia, 2004). This kind of spin is not really a form of deception (e.g. Napper, 1972) – the different interpretations are often due to assumptions about the way facts are perceived, and are often believed, with some level of sincerity, by those who present them.

The range of deception that we consider is illustrated in Figure 1. At one extreme, perjury (top left) represents factual deception serious enough to be considered a felony. Defense counsel, by contrast, typically tries to present a defendant's character in the best possible light (top right), even if this is a serious misrepresentation. In dating, factual misrepresentation may have negative consequences (bottom left), but misrepresentation of character (bottom right) is commonplace. In all aspects of society, factual misrepresentation is regarded more negatively than persona deception.

2.1. Psycholinguistic views of deception

From the psycholinguistic perspective, speech and writing are produced by largely subconscious processes and so, whenever people speak or write, they reveal information about themselves, including the fact that they are being deceptive.

There are three main theories to explain why deception creates detectable signals (Meservy *et al.*, 2007):

1. Deception causes *emotional arousal* associated with emotions such as fear, and perhaps delight. These in turn may cause physiological changes that are visible in body language and voice qualities.
2. Deception requires *cognitive effort* which, in turn, causes verbal disfluencies and body language associated with mental effort. Content may be simpler than expected, and delivery less smooth.
3. Deception is a *violation* of social and interpersonal norms of behavior, and so posture and voice may be more guarded and content more complex. There may, for example, be more intense scrutiny of the other parties to assess whether the deception is being effective.

These different theories make overlapping, but different, predictions about how deception should be detectable in communication and, of course, deception may involve all three processes.

2.2. Modelling deception

The literature has tended to focus on aural patterns of speech, such as intonation, as well as the microanalysis of facial expressions (Bull, 2002, 2003; Ekman 2002). Yet, these techniques are replete with problems of coding (Benoit, Laver, & Mikhaylov 2009). As coding requires interpretation, the findings may not be robust. Moreover, studies of this sort are difficult and costly to carry out consistently on a large scale and over a long period of time, such as a

presidential campaign since they may require, for example, high-speed video capture. In an attempt to transcend the limitations of the standard application of Correspondence Analysis to analyze political speeches (eg. Schonhard-Bailey, 2008), we develop a method that obviates the problem of coding while transforming a large quantity of unstructured data into structured data that lend themselves to statistical analysis. Drawing on recent advances in natural language processing (e.g. Beigman Klebanov, Diermeier, & Beigman, 2008; Laver, Benoit and Garry, 2003) our approach is based on text, either written or as a record of speech.

Using text-based analysis to investigate deception has precedent. For example, Gentzkow and Shapiro (2010) and Ansolabehere, Lessem, and Snyder (2006) observe that newspaper editorials heavily favor incumbents. Similarly, Larcinese, Puglisi, and Snyder (2006) observe partisan bias favoring incumbents in economic reporting. Most of the work in this field which focuses on the effect of political discourse on the *audience* in areas such as persuasion (Merton, 1946) and the manipulation of public opinion (Esser & Pfetsch, 2004; Lilleker, 2006; Stanyer, 2007; McNair, 2007). Campaign researchers tend to be more interested in the audience than the candidate (Benoit, 2007; Tuman, 2008). In its study of deception as an explicit political strategy, however, this literature is quite different from our approach, which focuses on political speeches as a way to understand *candidates* (and campaigns) rather than voters.

To this end, it is helpful to consider text as carrying two channels of information (Pennebaker & King, 1999). The first is the content channel, which is carried by ‘big’ words: nouns and verbs. The content of this channel is created (mostly) consciously by the speaker/writer and is available for rational analysis by listeners/readers.

The second channel is not as obvious. It is mediated by the small, functional words: conjunctions, auxiliary verbs, and adverbs. These are only 0.04% of English words, but make up 50% of words used. They carry little content but play a structuring role for the content-filled words that surround them. Changes in the frequencies of use of these words are strong signals for various kinds of internal mental states and activities, but humans are almost entirely blind to them, both as speakers/writers and as listeners/readers. Humans lack the ‘hardware’ to keep track of frequencies and whether or not these frequencies are abnormal or changing. As a result, the existence of this internal-state channel has been under-appreciated.

The internal-state channel is a reliable way to detect mental state in a speaker or author because its content is almost entirely subconsciously generated, and so cannot be directly manipulated to achieve a particular effect. A model of deception that uses text as input and looks for variations in word use in the internal-state channel has been derived by Pennebaker’s group (Newman, Pennebaker *et al.*, 2003, Tausczik & Pennebaker, 2010). The model is empirical, and was derived initially by asking students to support positions that they did and did not actually hold. The model has since been validated by collection of data from much larger and more varied populations (Chung & Pennebaker, 2007) and has proven to be remarkably robust.

This deception model uses the frequencies of words in four classes as markers. These classes are:

- First-person singular pronouns, for example, “I”, “me”.
- Exclusive words, words that introduce a subsidiary phrase or clause that refines or modifies the meaning of what has gone before, for example, “but”, “or”, “whereas”.
- Negative-emotion words, for example, “lie”, “fear”.
- Verbs of action, for example, “go”, “take”.

The complete set of 86 words used in the model is shown in Table 1. In this model, deception is signaled by:

- A decrease in the frequency of first-person singular pronouns, possibly to distance the author from the deceptive content;
- A decrease in the frequency of exclusive words, possibly to simplify the story and so reduce the cognitive overhead of constructing actions and emotions that have not really happened or been felt;
- An increase in the frequency of negative-emotion words, possibly a reaction to awareness of socially disapproved behavior; and
- An increase in the frequency of action verbs, possibly to keep the narrative flowing and reduce opportunities for hearers/readers to detect flaws.

The deception model will not detect those who genuinely believe something that is not, in actual fact, true. Rather it detects utterances that are, in some way, at variance with the beliefs of those who make them.

Pennebaker’s model of deception appears to detect deception across the entire spectrum described in Figure 1, presumably because they all share an underlying psychological process, a reflexive, though subconscious, awareness by the person being deceptive of what they are doing. The Enron email corpus, a large collection of email to and from employees of Enron, was made public by the U.S. Department of Justice as a result of its investigation of the company (Shetty & Adibi, 2004). Because nobody expected their emails to be made public, this corpus is a useful example of real-world texts of all sizes and purposes. When the deception model was applied to the emails in this corpus (Keila & Skillicorn, 2005; Gupta, 2007), emails that were detected as deceptive ranged from what appeared to be (sometimes with hindsight) egregious falsehoods, to less obviously false statements, and to forms of socially sanctioned deception such as negotiations.

2.3. Deception in political campaigns

Half a decade of research on vote models posits a complex interplay of three key variables: party, personality, and policies (Markus & Converse, 1979; Miller, 1991). The relative importance of each of these concepts to voters remains controversial, and varies depending on time and circumstances (e.g. Vavreck, 2001). The study of deception, however, is significant insofar as it has a bearing on all three concepts. Moreover, theories of voting suggest voters use cues, shortcuts, and symbols to minimize the costs associated with becoming informed about candidates' policy positions (Popkin, 1991; Fogarty & Wolak, 2009). Electoral politics are complicated, there are too many elections, and most voters are simply too busy to spend much time researching the positions of all the candidates for whom they could cast ballots (Converse, 1964; Delli Carpini, & Keeter, 1996; Lupia & McCubbins, 1998). In this light, candidates are highly motivated to endear themselves to as many potential voters as possible to maximize the payoff in the electoral marketplace. Candidates may thus downplay or conceal some of their beliefs, feelings, and character, while creating a persona or facade that presents beliefs, feelings, and character that they believe to be attractive to the largest audience. Deception should thus be observable across the entire spectrum in politicians campaigning for office. One manifestation consists of factual misstatements, for example Hillary Clinton's misremembering coming under sniper fire at the Sarajevo airport.

Although voters are prone to say that they prize honesty in politicians, the authors speculate that persona deception is a successful strategy for politicians. From social-identity theory, for instance, it follows that reaching out to a set of voters who are not yet committed to a candidate should work well if it appeals to human needs for affinity and relationship; therefore, for many voters an attractive and appealing personality may be more important, at the margin, than a robust and plausible set of policies. This claim is readily substantiated by research on political behavior which consistently shows large swaths of voters expressing preferences that are not necessarily premised on rational policy-maximizing choice. As a matter of fact, the decisive part of political behavior, the swing/independent vote, by and large, is not policy-utility maximizing. The model that is being posited in this article works towards bridging this gap.

Method

The article investigates persona deception empirically using the speeches of the major contenders for the presidency of the U.S. in 2008. It scrutinizes variations in the levels of persona deception of each candidate and campaign across time, and looks for correlation between levels of deception and other factors: time, location, audience, content, polls, and the activities of the other candidates. The deception model is a particularly useful tool because it is relatively immune to manipulation by a speaker, even one who is familiar with its markers and measures. This is because humans have little direct control over the way in which they use the internal-state channel, and the words of the model are predominantly those used in this channel. Hence deception and, particularly, changing levels of deception, are windows into the subconscious minds of speakers and writers.

Unstructured data are gathered by means of the candidates' speeches, retaining the entire text of each speech as it appears online. Attention is restricted to settings in which candidates were able to speak as they wished. Other settings such as debates, press conferences, and the like are complicated by the fact that questions impose some limits on the form of answers. For example, many questions are of the form "Did you ...?" which almost forces a response beginning "Yes/No, I did ... " in which a first-person singular pronoun appears. In fact, interrogative settings appear to influence the global structure of responses, changing the pattern of word usage corresponding to deception (Little & Skillicorn, 2008, Lamb & Skillicorn, 2013).

Political speeches are now largely written and edited by a team. The word usage in a particular speech is therefore a blending of language that originates with the speechwriting team, and with the candidate. The relative contribution of each component remains obscure; but it seems plausible to assume that the language patterns reflect the perceptions of the campaign team as a whole, its view of the environment in which it is operating day to day, and its view of the current degree of success and challenges faced. Since our sample texts were downloaded from web sites, it was not always possible to determine whether, in each case, they represented the text as written or the text as actually delivered¹.

Although the deception model is conceptually simple, it cannot be trivially applied to an individual speech to compute a deception score for two reasons. The first is obvious – measuring a decrease in frequency requires some estimate of the baseline or typical frequency. In computing a deception score for an individual text (un)certainty is determined within the context of a collection of similar texts. This collection could be all documents in English, but this will clearly give poor results because norms of usage differ by setting. For example, the use of first-person pronouns is rare in business writing.

The second reason that the model cannot be applied on a per-speech basis is that it implicitly assumes that each word of the deception model is a signal of equal importance. If, in a given

¹ The authors have discussed this issue with speechwriters and believe that the candidate is in fact the principal driver of speech language (PBS 2008; Lepore, 2009). First, a good speechwriter necessarily becomes an alter ego for the person for whom the speeches are written, and must be able to produce language that "sounds like" the candidate to be good at the job. Second, candidates often take a speech and, in the course of preparing to deliver it, and perhaps also *while* delivering it, make small changes. These alterations are often in the kinds of 'little' words that are important in the deception model. In other words, the candidate imposes his or her internal mental state on each speech by adjusting the language patterns to reflect their integrated perception of the current situation. Some evidence for this can be seen in the closing stages of the campaign where both candidates delivered speeches that were clearly based on the same templates – but there is considerable variation in the 'little' word frequencies and consequently in the resulting deception scores.

setting, most texts have a greatly increased frequency of a particular word, then it is probably not a marker for deception, but a stylistic marker typical of the setting. Hence, that word is less important as a marker of deception than the increased frequency would indicate. Conversely, even a very small increase in frequency of a word that appears in only a few speeches may be more important as a marker of deception than it seems.

To take such contextual information into account, the authors extend the basic deception model. The approach taken applies only to a set of texts, implicitly texts that are comparable in the sense of coming from similar settings. In this case, we make the assumption that candidates are driven by the same goal, and are competing in the same arena, so that, while each has a different history and set of abilities, their speeches can plausibly be considered to be comparable.

A matrix is constructed from the texts of speeches, one row corresponding to each speech, and one column to each word of the deception model. The ij th entry of the matrix is the frequency of word j in speech i . Each row of the matrix is normalized by dividing by the length of the corresponding speech, in words, to compensate for the naturally greater frequencies of all words in longer speeches.

Each column of the matrix is then zero-centered around the mean of its non-zero elements. This is done because the matrix is, in general, sparse, and the large number of zero-valued entries makes the denominator large and so the mean small, which blurs available information. The non-zero entries of columns are then scaled to standard deviation 1 (that is, converted to z-scores). The effect of this transformation is to remove the effect of the base frequency of each word in English.

For two classes of marker words (first-person singular pronouns and exclusive words), a *decrease* in frequency signals deception, while for the other two classes an increase in frequency signals deception. Since the normalized values in each column are centered at zero, the signs of the columns of words in the first two classes are reversed, so that larger positive magnitudes in the matrix are always associated with stronger signals of deception.

A singular value decomposition (SVD) (Golub & van Loan, 1996) is performed on the matrix of n speeches and m words. If this matrix is referred to as A , then a singular value decomposition expresses it as the product of three other matrices, thus:

$$A = USV'$$

where U is an n by m matrix, S is an m by m diagonal matrix, and V is an m by m matrix. U and V are both orthogonal matrices (their columns, as vectors, are mutually at right angles to one another), and the dash indicates matrix transposition.

Each of the n rows of A can be interpreted as defining a point in an m -dimensional space. In this space, rows that are similar to one another will be represented by points that are close to

one another, so geometric analysis techniques can be used for clustering. However, m is quite large so this geometric space is hard to work in and, in particular, hard to visualize.

One way to interpret a singular value decomposition is as a change of basis, from the usual orthogonal axes to a new set of orthogonal axes described by the rows of V . In this interpretation, U expresses the coordinates of each text relative to these new axes. This new basis has the property that the first axis is aligned along the direction of greatest variation among the texts; the second axis is aligned along the direction of greatest remaining orthogonal (that is, uncorrelated) variation, and so on. The entries of the diagonal of S capture how much variation exists along each axis. Singular value decomposition is akin to Principal Component Analysis (Jolliffe, 2002), but differs in capturing, simultaneously, variations in the word use of the speeches and the speech use of the words.

Just as in Principal Component Analysis, the new axes can be considered as latent factors capturing variation in the way words are used. There are typically far fewer latent factors than there are words, since languages are highly redundant, so the variation along many of the later axes will be negligible, and can be ignored. Also, since the majority of the variation is captured along the first few axes, the first two (respectively, three) columns of U can be interpreted as coordinates in two (resp., three) dimensions, and the similarities among the speeches visualized accordingly.

Ranking versus scoring

Rather than treating deceptiveness as an inherent property of documents, which is problematic for the reasons discussed above, we produce a ranking of a set of documents from most- to least-deceptive. Such a ranking is more robust than a per-document measure since it enables contextual information to be taken fully into account.

One would expect that, because all of the explicit factors (word frequencies) are markers of deception, the greatest variation in deception among speeches will be along the first axis of the transformed space. This implicitly assumes that deception is a single-factor property. However, even a cursory analysis of speeches suggests that deception is multifactorial, so some of the variation along other axes is also related to deception. In other words, a deceptive pattern of speech might include reduced use of first-person singular pronouns, or reduced use of exclusive words, and these variations need not be completely correlated. In what follows, it is assumed that the first three dimensions are those significantly related to deception. The significance of each dimension is estimated by the magnitude of the corresponding element of the diagonal of S and, in particular, one can tell how much structure in the speeches is being discounted by ignoring later dimensions from the magnitude of the first element of the diagonal of S that is ignored. In all of the results presented here, approximately 85% of the variation is captured in the first three dimensions. To compute a deception score for each speech, the points that correspond to each are projected onto a line from the origin passing

through the point (s_1, s_2, s_3) , the first three entries of the diagonal of S . This weights each factor exactly by its importance.

The singular value decomposition is entirely symmetric with respect to the rows and columns of the matrix A . Hence, exactly the same analysis can be carried out for the words. Points corresponding to each word can be visualized in three dimensions. This is useful because it enables us to determine how important each word is for the model; and how words are related to one another so, in particular, when two or more words signal approximately the same information.

Previous work has shown that the frequencies of the 86 words of the deception model are neither independent nor equally important (Gupta, 2007; Little and Skillicorn, 2009). Although it is, of course, domain dependent, in general first-person singular pronouns are the most significant markers; the words “but” and “or” are most significant among the exclusive words; and “go” and “going” are most significant among the action verbs. As classes, the first-person pronouns are more significant than the exclusive words, which are more significant than the action verbs, which are much more significant than the negative-emotion words.

Figure 2 shows the relationship of the different word classes for the complete set of election speeches, with the dashed lines indicating the direction of increased frequency and the distance from the origin indicating the significance of each word class as a marker of deception. First-person singular pronouns are the most important attribute; the words “or” and “but” and the action verbs are also important, but almost uncorrelated, both with the pronouns and with each other; and the other attributes have very little impact. Thus it is simpler, and just as effective, to start from a model that counts aggregated word frequencies for only the following six categories: the first-person singular pronouns; “but”; “or”; the remaining exclusive words; the action verbs; and the negative-emotion words. Obviously, deception is correlated with decreased frequencies of words in the first four categories and increased frequencies of words in the last two categories.

Model Uncertainties

The deception scores presented in what follows are the results of projecting a point corresponding to each speech onto a deception axis that has been induced from the total set of documents using the singular value decomposition. Since each deception score is a measurement within a discrete system, it does not have an associated error. However, it is necessary to consider uncertainties in the scores; that is, how much the score would change in response to a small change in the use of a single word in a document. The relationship between small changes in the input matrix and the resulting changes in the decomposition has been studied using perturbation analysis, which has shown that SVD is both insensitive to perturbations of the matrix and robust with respect to round-off errors caused by the use of limited-precision real arithmetic (Stewart, 1998, p. 69). Hence one can conclude that

small changes by a speaker have small effects on the resulting deception score; and also that errors in transcription or the like also have small effects.

A second issue is whether deception scores are distorted by the length of each speech. Varying lengths have been partially accounted for by dividing the word frequencies of each word class by the total length of the speech in which it appears. However, longer speeches contain more opportunities for variation in word frequencies. For example, a 1-word speech can produce a scaled word count of only 0 or 1 for each word; a 10-word speech can produce scaled word counts of 0, 0.1, 0.2, ... or 1.0, and a longer speech can produce even more finely divided scaled word counts. Does this matter to the computation of deception scores? Quantization, that is, using only a small discrete set of values for each matrix entry, can be modeled as adding a random matrix to the data matrix, and it is extremely unlikely that such a random matrix has significant variational structure (Achlioptas and MacSherry 2001). Hence, with high probability, the higher level of variability in the normalized scores from longer speeches has no effect on the computed deception scores.

To illustrate the method, consider the following four example texts. These are artificial examples, but have been derived from actual sentences used by the candidates. Words relevant to the deception model are underlined.

1. However, it is whether or not we have people who feel that they are going toward the American dream or whether they're disappointed because it looks like it's getting further away.
2. I am happy to tell you my story but I can't start without thanking my friends and colleagues.
3. But don't be afraid we are going to move into the future without fear, and I am going to lead you.
4. We know it takes more than one night, or even one election, to take action and overcome decades of money.

Which utterance is the most deceptive? It is relatively difficult to notice the words from the deception model, even on this modest amount of text (one has not been underlined as an exercise for the reader). It is also difficult to judge, intuitively, how high the score of some sentences should be. For example, sentence 1 contains several exclusive words (good) but also a negative-emotion word, and an action verb (bad). Should it be considered deceptive overall?

Table 2 shows the matrix calculated from these four sentences, where each entry is the frequency of words in the class divided by the length of the sentence. Table 3 shows the frequency values after the normalization procedure. The results are affected by the non-inclusion of zero-valued entries in the calculation of means and standard deviations, which creates some surprising results. For example, the raw frequencies of first-person singular pronouns in Sentences 2 and 3

are quite different, but the normalization has the effect of mapping them to “low” and “high”, centered at the origin.

Figure 3 shows a plot of the points in the transformed space. The projection of each point onto the line indicates the deception score associated with each speech (deception increasing in the direction of the solid line). Within the context of this very small set of texts, sentences 1 and 2 are not at all deceptive, although they are quite different from one another. Sentence 4 is neutral with respect to deception, and sentence 3 is the most deceptive in the set.

Results

Table 4 shows the characteristic rates of word use by each of the candidates, for the six marker words or classes used in the deception model. Rates are in bold font when they disagree with the expectations of the deception model.

The average persona deception score from the beginning of 2008 until the election for the three candidates is: McCain: 0.04 (SD 0.84); Obama: 0.11 (SD 0.52); and Clinton: -0.49 (SD 0.50). These averages conceal very large variations both from speech to speech, and over time. No candidate can be labeled, overall, as high or low in persona deception.

Comparing candidates

Figure 4 plots points corresponding to the entire set of speeches in three dimensions after the singular value decomposition. A deception axis is drawn passing from the origin to the points (s_1, s_2, s_3) and $(-s_1, -s_2, -s_3)$, where $s_1, s_2,$ and s_3 are the first three entries of the diagonal matrix S . The point corresponding to each speech can be projected onto this line to obtain a deception score, which is larger in magnitude the higher the level of persona deception relative to all of the other speeches in the set under consideration. The absolute magnitude of the deception scores has no intrinsic meaning, but relative magnitudes are meaningful. It makes sense to say that one speech is more deceptive than another, but not that one speech is, say, twice as deceptive as another.

The figure is cluttered, but the low-deception speeches are predominantly those given by McCain and Clinton. There is considerable variation in level of deception for each candidate, especially McCain. The speeches by Obama and Clinton are qualitatively quite different, with most of Obama’s towards the back of the figure, and Clinton’s towards the front. The nominal axes in this figure represent linear combinations of the original attributes (word frequencies) and so do not have a directly accessible interpretation, so we omit them. However, geometric proximity (closeness) does correspond to underlying similarity.

Zooming in to the region close to the origin (Figure 5) shows that all three candidates have some speeches with both high and low levels of deception, although Clinton’s are almost entirely low.

The patterns of persona deception for each candidate are more obvious by plotting levels of deception by candidate, in the order in which the speeches were given. Figure 6 shows the persona deception scores of McCain's speeches over the entire election, divided into three time periods: the initial period when there were two Democratic contenders, the period from Clinton's exit from the race (early June) to his party's convention (late August), and the period from the convention to the election. Positive numbers represent high levels of deception. The level of deception begins fairly low, then rises during the middle period of the campaign, dropping again towards the end.

Figure 7 shows the level of persona deception in Obama's speeches, divided into the same three periods (although the timing of the convention is slightly different—early September). The level of deception is generally higher than that of McCain throughout. As the election draws to a close, his levels of deception decrease. In the last few weeks, levels of deception increase once again.

Figure 8 shows the level of deception in Clinton's speeches. Her levels of deception are much lower across the board than for the other two candidates, but for two different reasons. In the early stages of the campaign, her speeches are characterized by complex policy statements that result in high levels of exclusive word use. In contrast, in the second half of the primary campaign, she reaches out in a very personal way, characterized by high rates of first-person singular pronoun use.

Figure 9 shows the persona deception levels of all three candidates aligned by time, and Figure 10 shows the same, but only for the Democratic candidates. It is striking that there are periods where the levels of persona deception in different campaigns change in a correlated way, at least to some extent. For example, in the first period, Obama and Clinton's levels of deception are similar, while McCain's are quite different. During this period McCain spent his time becoming known to Republicans, and effectively ignored what was happening in the Democratic campaign. This suggests that attention to the other campaign, and perhaps visiting similar locations and addressing groups with similar issues are relevant factors of this apparent coupling.

Similarly, in the second period, there are similarities between McCain and Obama, notably their decreased levels of deception in July as the economy suddenly became the major election issue, to which it took both of their campaigns some time to adapt; and the following period of high deception as they both strove to become the better expert on the economy, a role neither was well-equipped to play.

In the third period, there is again a substantial similarity between their levels of persona deception. This perhaps represents a consolidation of their appeal to those who are already committed to vote for them, and a certain amount of confidence about the outcome.

Medium-term variation

Clearly there are differences between candidates in the level of persona deception in their speeches. Some conclusions about possible explanations for the observed variation can be drawn.

First, it is plausible that the base level of persona deception differs from person to person, for four reasons:

1. Individuals have different abilities in delivering convincing speeches at any given level of persona deception. In a sense, an actor is indulging in a kind of deception (with the audience's permission) and a poor actor is one who fails to do so convincingly. Having briefly examined the levels of persona deception in the speeches of Bill Clinton and Ronald Reagan, both seem to be able to use high levels of persona deception in settings where most other politicians cannot. Speechwriters necessarily have to adapt the language they use to a level where it can be delivered convincingly by their candidate. Hence, one might expect that there will be significant differences between politicians independent of any political factors.
2. Campaigns must bridge the gap between the current perception of the candidates' personas, including the candidates' current self-perception, and the personas associated with the role for which they are running. Hence we might expect that politicians running for lower offices, or those for whom the step between their current role and level of the experience and the role for which they are running is small will require lower levels of persona deception.
3. Someone starting out in a political career may create, partly consciously and partly unconsciously, a persona to appeal to voters. With time, they may grow into this persona so that it becomes less of a facade. Hence, one might expect that mature politicians should exhibit lower levels of persona deception across the board.
4. Once a politician has used a particular persona, it becomes hard to create and use another persona – doing so creates the opportunity for charges of hypocrisy, and for the new persona to seem unconvincing. A politician who is well-known within a particular scope and over a period of time cannot perhaps create a fresh persona unless they campaign within a larger scope and for a qualitatively different role.

One cannot, of course, assess the base level of persona deception of any of the candidates (although it would be interesting to estimate this for novice politicians from their maiden speeches when elected). However, the other three factors are supported by the data: both McCain and Clinton exhibit lower average levels of persona deception than Obama. Thus name recognition may actually be a handicap for a previously successful candidate who now aspires to a qualitatively different office.

There are also changes in levels of persona deception over time periods of weeks and months. Because the internal-state channel is subconsciously mediated, these cannot be the direct result of decisions to change presentation style. However, they do reflect and integrate the entire election scenario from the candidate's, and campaign team's, point of view. As such they provide a new point of view and, moreover, one that cannot be manipulated by campaigns.

The level of persona deception over a medium-length period of time (weeks) is the result of a complex interaction between a perception of the primary audience being appealed to and the campaign's perception of the predisposition of the audience. This finding echoes work by Gordon and Miller (2004) on the contingent effectiveness in the way issues are framed. In an election, audiences can be divided into three categories: a base of people already willing to support a candidate more or less regardless, a group of people who can potentially be convinced to support the candidate (moderates, independents, center), and a group who are unlikely to support the candidate under any circumstances. At different stages of a campaign, one or other of these groups is the focus of attention. During primaries, campaigns appeal mostly to voters of their own party affiliation who perhaps already agree with the candidate on many issues, and are selecting the candidate based on less-objective criteria such as personality. In the early phases of an election, campaigns reach out to moderates whom they hope can be convinced to become supporters. In the closing stages, in a tight race, campaigns may even reach out to groups from opposing parties in the hope of gaining a small extra amount of support. At any given moment, a campaign will have a perception about how well their candidate is doing, and can potentially do, with each of these three groups; this opinion may be deeply visceral and so unconnected with polling results. Hence, one would expect that levels of persona deception would, overall, increase over the course of a campaign, as campaigns increasingly reach out to groups that are less their natural constituencies. Once, and if, these groups have been won over, levels of persona deception might once again decrease.

The use of a persona is also not a trivial artifact to construct, even though most of the construction is subconscious. Hence, one would expect that persona deception levels might be lower when unexpected issues suddenly surface that require new aspects of a persona to be developed. In other words, when a new issue arises, a campaign takes time, not only to decide how to handle it as a policy issue, but how to integrate their response to it into the persona they are using. The former is a conscious process; the latter is primarily subconscious. Hence, one would expect that levels of persona deception should fall, temporarily, when candidates are blindsided by unforeseen issues or by attacks.

Both of these phenomena are visible in the data through the course of the election. McCain's situation was unusual in that his natural constituency was moderate voters, and he was much less popular in his own party. His deception scores were low in the early part of the election as he found himself, surprisingly, the Republican Party nominee. During the middle third of the election, his persona deception became high as he reached out (to Democratic-leaning moderates and Republicans). After the conventions, his problem with the Republican base was solved, for a time, by the selection of Palin, and his deception scores decreased, rising in the closing stages as he reached out to the remaining swing vote.

Obama's level of persona deception was consistently high, reflecting the fact that he began with no guarantee of the nomination as he was in a race with Clinton for the first half of the

year, and he had the least experience and standing as a potential presidential candidate. He also reached out to moderates almost from the beginning of the campaign. As his success became increasingly assured, his levels of deception dropped. As with McCain, levels rose in the closing stages as there was a brief period when his win did not seem certain. There were also several moments early in the campaign when Obama's deception dropped suddenly. These correspond to the moment when it became mathematically clear that he will win the party nomination (late in February); and when the Pastor Wright controversy erupted, causing questions about his attitudes to race and issues surrounding it. In the first case, the increased security seems to have emboldened his campaign, briefly, to start using a more open style; in the second, it took his campaign team some time to decide how to react. His speech on race represents one of the lowest levels of deception in the first half of the campaign.

Clinton's level of persona deception was much lower than that of the other two candidates throughout. She was only involved in the primary portion of the campaign, so her audience is primarily Democrats. She clearly felt no need to use a persona, since she was, in part, relying on her name and character recognition – indeed she was so well known that it would have been difficult for her to recreate herself.

Short-term variation

What about the surprisingly large amount of day-to-day variability in levels of persona deception, which is not accounted for by the model that has been proposed so far? We computed rolling averages over various periods of time to examine the possibility that the observed scores fit a slowly varying underlying change in deception, overlaid with some form of short-term noise. However, no smoothing was observed over any of the time periods considered, which leads us to conclude that the short-term variability captures an underlying rapid real change in levels of deception scores.

The deception model implicitly suggests that deception requires a speech act that is perceived by the speaker to be, at some level, socially inappropriate, and that the increase in negative-emotion words is a reflection of the negative emotions that such an activity causes in the speaker and writer. Thus any kind of sustained deception requires balancing a (subconscious) desire to present the candidate as better than he or she actually is, with the resulting negative emotional impact. Given this perspective, it is less surprising that the level of deception changes from day to day. There is also a certain amount of cognitive effort involved in presenting a persona, and this may be difficult to maintain at an even level over time.

Some other plausible hypotheses about the short-term variability are:

- Persona deception depends on the target audience – deception levels will be lower when speaking to a 'friendly' audience, for example, veterans for McCain and union members for Obama.

- Persona deception depends on topic – some topics are more natural for a candidate and so speeches in these topics will have lower levels of deception.
- Persona deception depends on recent success – a candidate who has recently done well, for example by winning an important state primary, will have lower levels of deception immediately afterwards.

The authors investigated whether these reasons had any predictive power by building decision-tree predictors for level of deception, coding each of these potential reasons as a set of attribute values. All of these predictors performed at the level of chance, or slightly worse, so one can conclude that these reasons do not explain the short-term variation in levels of deception.

Instead we posit a functional theory of tactical purpose to explain this variation. Examination of the kinds of speeches given in campaigns suggests that it is useful to classify them into three categories by purpose:

1. Blue-skies policy speeches. Such speeches enunciate some policy proposal that is unconnected to anything that the candidate has done before. Such a speech has little connection to the candidate as an individual; and could be given interchangeably by different people as the content is not predicated on the speaker. The writers of such a speech do not need to consider much about the particular candidate when crafting such a speech; it could be given with any level of persona deception, from very low to very high (and so would normally be as high as the candidate is able to deliver convincingly).

For example, here is a fragment of a typical blue-skies policy speech:

Washington is still on the wrong track and we still need change. The status quo is not on the ballot. We are going to see change in Washington. The question is: in what direction will we go? Will our country be a better place under the leadership of the next president – a more secure, prosperous, and just society? Will you be better off, in the jobs you hold now and in the opportunities you hope for? Will your sons and daughters grow up in the kind of country you wish for them, rising in the world and finding in their own lives the best of America?

This was actually said by McCain (October 6th, Albuquerque) but, from the sentiments expressed and even the style, it could easily have been said by Obama or Clinton.

2. Track-record policy speeches. Such speeches also enunciate a policy proposal, but in a way that is tied to the personality, history, or track record of the candidate. Such a speech only works if it is given by the candidate for whom it was written – it can no longer be given interchangeably by different people. Because of the speech's purpose, its level of deception is now more limited; it requires much greater skill to be written and delivered in a high-deception way because of the need somehow to associate it with the speaker. Potential discrepancy between the speaker's actions and past persona limit the extent to which a new persona can be presented. Our examination of speeches by Bill Clinton and Ronald Reagan suggest that they were especially skilled at giving high-persona-deception track-record speeches.

Here is a fragment of typical track-record policy speech:

But I also know where I want the fuel-efficient cars of tomorrow to be built — not in Japan, not in China, but right here in the United States of America. Right here in the state of Michigan.

We can do this. When I arrived in Washington, I reached across the aisle to come up with a plan to raise the mileage standards in our cars for the first time in thirty years — a plan that won support from Democrats and Republicans who had never supported raising fuel standards before. I also led the bipartisan effort to invest in the technology necessary to build plug-in hybrid cars.

This was said by Obama (August 4th, Lansing) and uses something he has done to add credibility to something he plans to do. The same words would not work for McCain or Clinton because they have not taken the actions described.

3. Manifesto speeches. These speeches contain a smattering of policy material, often well-worn and diffuse, and are primarily intended both to convey the personality and good qualities of the speaker, and to reassure supporters of the bond between them and the candidate, and between each other as supporters. Again, the purpose tends to limit the level of persona deception achievable in such a speech, because it almost forces consistency between the persona being presented, and its previous incarnations.

Here is a fragment of a manifesto speech:

I am grateful for this show of overwhelming support. I came to Puerto Rico to listen to your voices because your voices deserve to be heard. And I hear you, and I see you, and I will always stand up for you.

This was said by Clinton (June 1st, Puerto Rico). There is no policy content; it is entirely about Clinton as a person.

In one sense, these different kinds of speeches can be seen as intended for different audiences. Manifesto speeches are intended for the candidate's base; track-record policy speeches are intended for those who are primed to align with the candidate but need to be told what specific goals and plans are intended; while blue-skies policy speeches are intended to reach out to independents and moderates and those who would not naturally consider the candidate. Following Gentzkow and Shapiro (2006, 2010), this is to be anticipated: candidates respond to consumer preferences by slanting speeches toward the prior beliefs of their audiences.

The differences among the levels of persona deception for each campaign are almost completely accounted for by the choice of purpose for each speech. For all three candidates, blue-skies policy speeches have high levels of deception; and track-record policy speeches have moderate levels of deception. For McCain and Clinton, manifesto speeches have low levels of deception, but Obama's manifesto speeches have moderate levels.

One can, therefore, examine the way in which each campaign *chooses*, tactically, the purpose of each speech as another way to understand their strategies. It seems unlikely that these choices

are entirely conscious. For example, there is a strong saw-tooth pattern in the level of persona deception of all of the candidates, so that a high-deception speech is followed immediately by a low-deception speech, and vice versa. In fact, a blue-skies policy speech is often followed immediately by a track-record or manifesto speech. This suggests that, at some level, campaigns are uncomfortable with both extremes: high persona deception, in case it comes to seem hypocritical and low persona deception, in case it is too revealing.

Table 5 shows the distribution of speeches into these three categories, and the resulting classification of each speech as high- or low-deception (deception score positive or zero; or negative). The average level of persona deception within each category is also provided, but these averages conceal large variations, some of them with a strong temporal component. Notice the consistent average deception scores for the three different kinds of speeches across the three candidates despite the large differences in the proportions of each kind used by candidates (with the exception of Obama's manifesto speeches).

McCain chose to give many policy speeches, suggesting what needs to be done in areas such as foreign and economic policy. At the beginning of the primary campaign, and once he became the presumptive nominee, he gave more frequent manifesto speeches, but his speeches leading up to the convention were almost entirely policy speeches, and these account for the higher average levels of persona deception over this period. Manifesto speeches were relatively rare, perhaps reflecting the fact that he did not play well to the Republican base. His track-record policy speeches in the first half of the campaign tended to be high-deception. This is because he developed a technique of wrapping typical blue-skies policy content inside opening and closing paragraphs that are much more personal. This enables him to fulfill the purpose of a track-record policy speech, making it clear that he has performed in certain areas, while still allowing a great deal of blue-skies content. In the second half of the campaign, the focus turned to the economy, and he gave many track-record speeches. These tended to be low-deception, at least initially, because they were cast in terms of what he had done, and what he planned to do.

In contrast, both Democratic contenders had a greater proportion of manifesto speeches, particularly in the earlier part of the campaign. This is unsurprising, given that they were competing with each other for the Democratic base during that time. In the period between the time when Obama became the presumptive Democratic nominee mathematically, and the time when Clinton actually withdrew, Obama gave track-record policy speeches, while Clinton gave almost entirely manifesto speeches. As a primary strategy, this clearly worked – between March and June, Clinton did much better, compared to Obama, than she had before.

Obama and Clinton's manifesto speeches also show markedly different levels of persona deception. It seems clear that Clinton addresses her supporters without any facade (or, conceivably, as an experienced politician, her facade has become her personality). Obama, in contrast, shows signs of presenting quite a strong facade, even to his supporters. Arguably, part of his success was because voters projected their desires for a presidential candidate onto this facade.

Obama and McCain both developed track-record speeches in the last third of the campaign that were much lower in persona deception than either had managed up to that point. This is largely because each was putting forward an economic plan that depended on their personal credibility, rather than on content as such.

There are trends over time in what kind of speeches all of the campaigns used. In particular, manifesto speeches disappeared almost completely in the last third of the campaign period, being replaced by track-record speeches. This is plausible given that, by that stage, both candidates had achieved their party's nomination and were concentrating on voters in the middle ground.

Many of the speeches given by both McCain and Obama over the last months of the campaign were essentially identical from day to day – they were driven by the same template.

Nevertheless, it is revealing how much the persona deception scores differ from day to day, indicating how the text is being altered in small ways from the underlying template, reflecting the changing feelings and attitudes of each candidate and campaign on a daily timescale.

Our results suggest that the level of persona deception achieved by a campaign at any given moment comes from a dynamic tension between factors that increase persona deception, and factors that act to limit it. The factors that increase persona deception are:

- A global desire to increase a candidate's appeal to the greatest number of voters in the electoral marketplace;
- A desire to reach out to more distant groups as campaigns progress – first consolidating the 'base', and then trying to convince a wider audience;

and the factors that limit persona deception are:

- The ability of a campaign to create a differentiated persona for their candidate that is not obvious, that is which does not appear to be artificial or phony;
- Psychological pressures associated with negativity and cognitive load;
- The tactical purpose of each speech in the bigger picture.

Levels of persona deception are therefore the result of a complex balance between perceptions and abilities. Because much of this balancing takes place subconsciously, however, it provides a new window into how candidates and campaign teams are viewing a campaign at any particular moment. In particular, changes in levels of persona deception indicate changes in these views; therefore, they signal that the available information is being integrated in a new way.

Persona deception and polls

Figures 11 (McCain) and 12 (Obama) show the relationship between levels of persona deception and favorable ratings in the daily Gallup tracking poll, with the level of deception

shown as a solid line, and the poll data as a dashed line. Favorable poll ratings have been scaled to comparable magnitudes to allow for easier comparison.

There is no clear dependency between polls and levels of deception in either direction, although there are indications of some weak coupling. This is interesting, and perhaps surprising, because it might be expected that poor polling numbers would increase efforts to reach out, and so would increase levels of deception. As Obama's lead increased in the last few weeks of the campaign, his levels of persona deception decreased – but so did McCain's. Further investigation is needed here.

First-person singular versus first-person plural pronouns

The deception model does not involve the first-person plural pronouns, “we”, “us” and so on. It is commonly believed, based on intuition, that these pronouns play an important role in openness and relative power, and so might be relevant to deception. In fact, research suggests that this is not the case (Chung & Pennebaker, 2007). When used by women, pronouns such as “we” may have an element of inclusiveness, but when used by men, they tend to signal an artificial inclusiveness, making what is actually a command sound a little softer. The fact that Obama used high rates of “we” while Clinton uses high rates of “I” has been taken to indicate that Obama is inclusive and Clinton is egotistical. In fact, the opposite is true. In a dialogue, the lower-status person tends to use higher rates of “I”. In contrast, using “we” is a weak form of manipulation. A good example is an Obama statement, “we need to think about who we elect in the Fall”. Of course, Obama already knew who he was going to vote for, so this is really code for “you need to think about who you elect in the Fall”, expressed in a way that hides the implicit command.

Conclusion

Politicians who run for mainstream parties with large electorates need to maximize their appeal to get re/elected. Although they can be factually deceptive, this may be ineffective, since voters tend to believe that all politicians are factually deceptive, or risky, since their deception may be demonstrated objectively. They can also be deceptive in a subtler and less visible way, by presenting themselves as ‘better’, in a broad sense, than they are, using what we have called persona deception.

It seems plausible that persona deception is an effective strategy, and levels of persona deception did agree with outcomes in this case. However, it is also clear that persona deception is highly variable. To synthesize the empirical findings, we consider possible sources of this variation and assess their plausibility. These sources of, and constraints on, variation fall naturally into four categories: (1) aspects of the personality of each candidate; (2) aspects of the audience, both physical and virtual, to which each speech is made; (3) aspects of the content chosen for each speech; and (4) aspects of the current success or failure of the campaigning candidate at the time each speech was made.

Potential personal factors that might affect levels of persona deception are:

- The raw ability of a candidate to present a credible persona that differs from that candidate's 'natural' persona. The difference between good and bad actors suggests that humans vary substantially in this ability and it would be surprising if this were not also true of politicians. This is difficult to measure directly because it is confounded by other personal factors (but would be interesting to assess for novice campaigners).
- The limits imposed by the need for consistency with an existing political persona. Presenting a persona different from that used in the current role seems both psychologically difficult and an invitation to be considered hypocritical – a particular form of the challenges of incumbents who claim, if elected, to be able to solve problems they have been unable to solve thus far. The data does suggest that persona deception is higher for candidates with less experience and/or track record.
- The gap between each campaigner's current persona and the persona perceived to be required to fill the role in question. It seems plausible that persona deception is greater for inexperienced or junior candidates than for those who are more senior or more experienced.
- The effort required to produce and maintain a persona. This suggests that persona deception will tend to wane over time, so that increases should be considered as more significant than decreases.

Hence personal factors largely determine the baseline persona deception available to each campaign as a function of the candidate's personality and abilities, and the candidate's personal political situation at the outset of the campaign.

Potential audience factors that might affect levels of persona deception are:

- The natural attraction between candidate and the audience present at a speech's delivery. Although this seems superficially plausible, the data does not support the obvious conclusion that levels of persona deception will be lower in speeches given to 'friendly' audiences.
- The virtual audiences that candidates conceive themselves as addressing. The data indicates a tactical consideration in each speech, reaching out to a particular set of potential voters that both the content and the presentation are crafted to reach. Levels of persona deception will be lower when the tactical purpose of a speech is to reach a natural constituency (rather than a natural audience).

Hence, audience factors affect the medium variation in the level of persona deception as campaigns change tactics, typically moving from attracting their base, to attracting independent voters, and perhaps even reaching out to natural opponents.

Potential content factors that might affect levels of persona deception are:

- Familiarity with the particular content of a speech. The data does not support the obvious conclusion that speeches in areas with which the candidate is especially familiar or expert have lower levels of persona deception.
- The extent to which the content of the speech represents an area where the candidate, if elected, has the ability to exert control or force change. The data suggests that the less the candidate has control of an area, the greater will be the level of persona deception in speeches about that area.

Hence, content factors affect medium and short-term variation in the level of persona deception as a result of each campaign's choices about what topics to address in each speech. However, the relationship is indirect and underappreciated.

Potential success-related factors that might affect levels of persona deception are:

- Recent high or increasing approval ratings in tracking polls. The data does not support the conclusion that better polling numbers imply lower persona deception.
- Perception by a campaign team that their candidate is doing well. Although it is hard to elicit such perceptions, there does appear to be an association between lower levels of persona deception and moments in the campaign when a campaign might reasonably have felt that they were doing well.
- Recent successful events, such as winning a primary. The data does not support an association between such events and lower levels of persona deception.
- The relative success of other campaigns. The data suggests that levels of persona deception rise and fall in a correlated way with competing campaigns, although we do not have an explanation for why this should happen.

Hence success-related factors affect medium and short-term variation in the level of persona deception and, as such, provide a window into each campaign's view of how well it is doing that cannot easily be seen in any other way.

Overall, variations in the level of persona deception in the medium term appear to be driven by a complex, dynamic, and subconscious integration of the campaign's intended audience at a particular stage, the campaign's belief about the receptiveness of that audience, and (in a way that is not yet fully understood) the actions of competing campaigns. Over the short term, changes appear to be driven by a balance in individual psychological and cognitive factors, and the tactical purpose of each speech, a finer-grained version of the (subconscious) calculation that drives medium-term variation. As these changes are largely subconsciously mediated, they cannot be directly manipulated and so provide a new way to gain insight into how each candidate and campaign regard the election arena at any given moment.

References

- Chung, C., Pennebaker, J.: The Psychological Function of Function Words. In: Fiedler, K. (ed.) *Social communication: Frontiers of Social Psychology*, pp. 343-359. Psychology Press, New York (2007)
- Skillicorn, D.: *Understanding Complex Datasets: Matrix Decompositions for Data Mining. Volume I*. Chapman & Hall/CRC Data Mining and Knowledge Discovery Series, Boca Raton (2008)
- Slapin, J., Proksch, S.-O.: A Scaling Model for Estimating Time-Series Party Positions from Texts. *American Journal of Political Science* 52(3), 705-722 (2008)
- Lowe, W.: Understanding Wordscores. *Political Analysis* 16(4), 356-371 (2008)
- Laver, M., Benoit, K., Garry, J.: Extracting Policy Positions from Political Texts Using Words as Data. *American Political Science Review* 97(2), 311-331 (2003)
- Corcoran, P.: Presidential Concession Speeches: The Rhetoric of Defeat. *Political Communication* 11(2), 109-131 (1994)
- Benoit, W.: *Communication in Political Campaigns*. Peter Lang, New York (2007)
- Hacker, K., Zakahi, W., Giles, M., McQuitty, S.: Components of Candidate Images: Statistical Analysis of the Issue-Persona Dichotomy in the Presidential Campaign of 1996. *Communication Monographs* 67(30), 227-238 (2000)
- Stuckey, M., Antczak, F.: The Battle of Issues and Images: Establishing Interpretive Dominance. In: K. Kendall (ed.) *Presidential Campaign Discourse: Strategic Communication Problems*, pp. 117-134. State University of New York Press, Albany (1995)
- Alia, V.: *Media Ethics and Social Change*. Routledge, New York (2004)
- Napper, C.: *The Art of Political Deception*. Johnson, London (1972)
- Meservy, T., Jensen, M., Kruse, W., Burgoon, J., Nunamaker, J.: Automatic Extraction of Deceptive Behavioral Cues From Video. In: Chen, H., Reid, E., Sinai, J., Silke, A., Ganor, B., Joslin, K. (eds.) *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security Series*, pp. 495-516. Springer, Boston (2007)
- Bull, P.: *Communication Under the Microscope: the Theory and Practice of Microanalysis*. Routledge, New York (2002)
- Bull, P.: *The Microanalysis of Political Communication: Claptrap and Ambiguity*. Routledge, New York (2003)
- Ekman, P.: *Telling Lies: Clues to Deceit in the Marketplace, Marriage, and Politics*. W.W. Norton, New York (2002)
- Benoit, K., Laver, M., Mikhaylov, S.: Treating Words as Data with Error: Uncertainty in Text Statements of Policy Positions. *American Journal of Political Science* 53(2), 495-513 (2009)

Schonhardt-Bailey, C.: The congressional debate on partial-abortion: Constitutional gravitas and moral passion. *British Journal of Political Science* 38(3), 383-410 (2008)

Beigman Klebanov, B., Diermeier, D., Beigman, E.: Lexical Cohesion Analysis of Political Speech. *Political Analysis* 16(4), 447-463 (2008)

Gentzkow, M., Shapiro, J.: What Drives Media Slant? Evidence from U.S. Daily Newspapers, *Econometrica* 78(1), 35-71 (2010)

Ansolabehere, S., Lessem, R., Snyder, J. : The Political Orientation of Newspaper Endorsements in U.S. Elections, 1940-2002. *Quarterly Journal of Political Science* 1(4), 393-404 (2006)

Ansolabehere, S., Snowberg, E., Snyder, J.: Television and the Incumbency Advantage in U.S. Elections. *Legislative Studies Quarterly* 31(4), 469-490 (2006)

Larcinese, V., Puglisi, R., Snyder, J.: Partisan Bias in Economic News: Evidence on the Agenda-Setting Behavior of U.S. Newspapers. Working paper (2006)

Merton, R.: *Mass Persuasion: The social psychology of war bond drive*. Harper, New York (1946)

Esser, F., Pfetsch, B.: *Comparing Political Communication: Theories, Cases, and Challenges*. Cambridge University Press, Cambridge (2004)

Lilleker, D.: *Key Concepts in Political Communication*. Sage Publications, London (2006)

Stanyer, J.: *Modern Political Communication: Mediated Politics in Uncertain Time*. Cambridge, MA: Polity Press (2007)

McNair, B.: *An Introduction to Political Communication*. Routledge, London (2007)

Tuman, J.: *Political Communication in American Campaigns*. Sage Publications, Los Angeles (2008)

Pennebaker, J., King, L.: Linguistic styles: language use as an individual difference. *Journal of personality and social psychology* 77(6), 1296-1312 (1999)

Newman, M., Pennebaker, J., Berry, D., Richards, J.: Lying Words: Predicting Deception from Linguistic Style. *Personality and Social Psychology Bulletin* 29(5), 665-675 (2003)

Tausczik, Y., Pennebaker, J.: The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology* 29(1), 24-54 (2010)

Shetty, J., Adibi, J.: *The Enron Email Dataset Database Schema and Brief Statistical Report*. Information Sciences Institute, University of Southern California, Los Angeles (2004)

Keila, P., Skillicorn, D.: Structure in the Enron Email Dataset. *Computational and Mathematical Organization Theory* 11(3), 183-199 (2005)

Gupta, S.: *Modeling Deception Detection in Text*. Master's Thesis, Queen's University School of Computing, Kingston (2007)

- Markus, G., Converse, P.: A dynamic simultaneous equation model of electoral choice. *American Political Science Review* 73(4), 1055-1070 (1979)
- Miller, W.: Party identification, realignment, and party voting: Back to the basics. *American Political Science Review* 85(2), 557-568 (1991)
- Vavreck, L.: The Reasoning Voter Meets The Strategic Candidate: Signals and Specificity in Campaign Advertising, 1998. *American Politics Research* 29(5), 507-529 (2001)
- Popkin, S.: The reasoning voter: Communication and persuasion in presidential campaigns. University of Chicago Press, Chicago (1991)
- Fogarty, B., Wolak, J.: The Effects of Media Interpretation for Citizen Evaluations of Politicians' Messages. *American Politics Research* 37(1), 129-154 (2009)
- Converse, P.: The nature of belief systems in mass publics. In D. Apter (Ed.), *Ideology and discontent*, pp. 206-261. Free Press, New York (1964)
- Delli Carpini, X., Keeter, S. : What Americans know about politics and why it matters. Yale University Press, New Haven (1996)
- Lupia, A., McCubbins, M.: The democratic dilemma: Can citizens learn what they need to know? Cambridge University Press, New York (1998)
- Lamb, C., Skillicorn, D.: Extending Textual Models of Deception to Interrogation Settings, *Linguistic Evidence In Security, Law And Intelligence (LESLI)* 1(1), 13-40 (2013)
- Little, A., Skillicorn, D.: Patterns of Word Use for Deception in Testimony. In: Chen, H. Reid, E. Sinai, J. Silke, A., Ganor, B. *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security Series*, pp. 425-450. Springer, Boston. (2008)
- PBS: Inside the Art of Political Speech Writing. Retrieved from <http://www.npr.org/templates/story/story.php?storyId=94074060>
- Lepore, J.: The Speech. *New Yorker*, 48-53 (12 January 2009)
- Golub, G., van Loan, C.: Matrix computations. Johns Hopkins Studies in Mathematical Sciences (3rd ed) The Johns Hopkins University Press, Baltimore, MD (1996)
- Jolliffe, I.: Principal Component Analysis (2nd ed) Springer-Verlag Inc., New York (2002)
- Stewart, G.: Matrix Algorithms, Volume 1: Basic Decompositions. Society for Industrial and Applied Mathematics (1998)
- Achlioptas, D., McSherry, F.: Fast Computation of Low Rank Matrix Approximations. *Journal of the ACM* 54(2), 1-19 (2007)
- Gordon, A., Miller, J. Values and Persuasion During the First Bush-Gore Presidential Debate. *Political Communication* 21(1), 71-92 (2004)

Gentzkow, M., Shapiro, J. Media Bias and Reputation. *Journal of Political Economy* 114(2), 280-316 (2006)

Chung, C., Pennebaker, J.: Revealing Dimensions of Thinking in Open-Ended Self-descriptions: An Automated Meaning Extraction Method for Natural Language. *Journal of Research in Personality* 42(1), 96–132 (2008).

Fridkin, K., Kenney, P., Gershon, S., Woodall, G.: Spinning Debates. *International Journal of Press/Politics* 13(1), 29-51 (2008)

Gerber, A., Karlan, D., Bergan, D.: Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics* 1(2), 35-52 (2009)

Kelley, C.: *Post-9/11 American Presidential Rhetoric: a Study of Protofascist Discourse*. Lexington Books, Lanham, MD (2007)

Table 1: The 86 words used by the Pennebaker model of deception

Categories	Keywords
First-person singular	I, me, my, mine, myself, I'd, I'll, I'm, I've
Exclusive words	but, except, without, although, besides, however, nor, or, rather, unless, whereas
Negative-emotion words	hate, anger, enemy, despise, dislike, abandon, afraid, agony, anguish, bastard, bitch, boring, crazy, dumb, disappointed, disappointing, f-word, suspicious, stressed, sorry, jerk, tragedy, weak, worthless, ignorant, inadequate, inferior, jerked, lie, lied, lies, lonely, loss, terrible, hated, hates, greed, fear, devil, lame, vain, wicked
Motion verbs	walk, move, go, carry, run, lead, going, taking, action, arrive, arrives, arrived, bringing, driven, carrying, fled, flew, follow, followed, look, take, moved, goes, drive

Table 2: Word frequencies for the example texts

	First-person singular pronouns	“but”	“or”	Exclusive words	Negative emotion words	Action verbs
Sentence 1	0	0	0.065	0	0.032	0.032
Sentence 2	0.222	0.056	0	0.056	0	0
Sentence 3	0.048	0.048	0	0.048	0.048	0.190
Sentence 4	0	0	0.05	0	0	0.1

Table 3: Normalized word frequencies for the example texts

	First-person singular pronouns	“but”	“or”	Exclusive words	Negative emotion words	Action verbs
Sentence 1	0	0	-0.007	0	-0.008	-0.075
Sentence 2	-0.087	-0.004	0	-0.004	0	0
Sentence 3	0.087	0.004	0	0.004	0.008	0.083
Sentence 4	0	0	0.007	0	0	-0.008

Table 4: Characteristic word use by each of the candidates compared to the pattern of extreme deception (bold indicates deviation from the model)

	First-person singular pronouns	“but”	“or”	Exclusive words	Negative emotion words	Action verbs
Extreme deception	Low	Low	Low	Low	High	High
McCain	High	Low	Low	High	High	Low
Obama	Low	High	High	Low	High	Low
Clinton	V. High	Low	Low	High	Low	High

Table 5: Distribution of speech types and their related deception scores

	Blue-skies policy	Track-record policy	Manifesto
McCain	28/76 = 37% 27 high, 1 low Av. DS 0.77	35/76 = 46% 13 high, 22 low Av. DS -0.22	13/76 = 17% 2 high, 11 low Av. DS -0.81
Obama	25/114 = 21% 25 high, 0 low Av. DS 0.61	55/114 = 48% 24 high, 31 low Av. DS -0.07	34/114 = 30% 20 high, 14 low Av. DS 0.07
Clinton	1/32 = 3% 1 high, 0 low Av. DS 0.41	5/32 = 16% 3 high, 2 low Av. DS -0.19	26/32 = 81% 1 high, 25 low Av. DS -0.59

Figure 1: The spectrum of deception

Figure 2: Relationships among word frequencies

Figure 3: Deceptiveness for the example texts. The axes are linear combinations of the frequencies of the six word classes, and so have no directly accessible interpretation; however, proximity is equivalent to similarity. The line indicates the deception axis, with a solid line in the direction of positive score and the dashed line in the direction of negative score. Deception scores are projections onto this line.

Figure 4: Relationships among speeches for the entire campaign. Speeches by McCain are shown by circles, Obama by stars, and Clinton by hexagons. The low-deception end of the deception score axis is dotted. Showing only three dimensions captures 83% of the variation in this data.

Figure 5: Same figure zoomed in on the region around the origin

Figure 6: McCain's persona deception score over time for the whole election with the vertical lines indicating key moments: Clinton dropping out; Republican convention; and election. (Scores are normalized across all three candidates, so deception score magnitudes are comparable between candidates)

Figure 7: Obama's persona deception score over time for the whole election

Figure 8: Clinton's persona deception score over time until she drops out of the race

Figure 9: Deception score by day of the year from the beginning of 2008 until the election. McCain, red solid line; Obama blue solid line; Clinton, blue dashed line.

Figure 10: Deception score by day of the year for the Democratic candidates (Obama, blue solid line; Clinton, blue dashed line). Note the parallels in levels of deception between the two campaigns.

Figure 11: Deception score (solid line) and Gallup daily tracking poll approval (dashed) for McCain (scaled by matching means and standard deviations to allow comparison)

Figure 12: Deception score (solid line) and Gallup daily tracking poll approval (dashed) for Obama (scaled by matching means and standard deviations to allow comparisons)