

# Predicting Elections from Biographical Information about Candidates

J. Scott Armstrong

The Wharton School

University of Pennsylvania, Philadelphia, PA

[armstrong@wharton.upenn.edu](mailto:armstrong@wharton.upenn.edu)

Andreas Graefe

Institute for Technology Assessment and Systems Analysis

Karlsruhe Institute of Technology, Germany

[graefe@kit.edu](mailto:graefe@kit.edu)

January 5, 2010

**Abstract.** Traditional election forecasting models are estimated from time-series data on relevant variables and that limits the type and number of variables that can be used. Index models do not suffer from the same restrictions. We used as many as 60 biographical variables to create an index model for forecasting U.S. Presidential Elections. For each candidate, we simply counted the number of variables for which the candidate was rated favorably. The index model forecast was that candidate A would win the popular vote if he had a higher index score than candidate B. We used simple linear regression to estimate a relationship between the index score of the candidate of the incumbent party and his share of the popular vote. We tested the model for the 29 U.S. presidential elections from 1896 to 2008. The model's forecasts, calculated by cross-validation, correctly predicted the popular vote winner for 27 of the 29 elections and were more accurate than those from polls (15 out of 19), prediction markets (22 out of 26), and three regression models (12 to 13 out of 15 to 16). Out-of-sample forecasts of the two-party popular vote shares were more accurate for the last four elections from 1996 to 2008 than those from seven prominent regression models. By relying on different information and including more variables than traditional models, the biographical index model can improve the accuracy of long-term election forecasting. In addition, it can help parties to select the candidates running for office.

---

Presented at the *Symposium on Leadership and Individual Differences*, Lausanne, Switzerland, November 30 - December 1, 2009. An earlier version was presented at the *29<sup>th</sup> International Symposium on Forecasting*, Hong Kong, June 21-24, 2009.

For three decades now, economists and political scientists have used regression models to estimate the impact of variables such as economic growth and the incumbent president's popularity on the outcomes of U.S. Presidential Elections. The strong correlation between the popular vote for a candidate and these variables has fostered the view by some researchers that campaigns have little impact on the election outcome.

This is surprising as candidates play a major role in U.S. Presidential Elections. No matter how irrelevant some individual differences between candidates may appear in respect to their performance once in office (such as skin color), they are extensively discussed in the media. Furthermore, many researchers have studied the impact of biographical traits of politicians on their chances of being elected. For example, factors such as candidate's height or facial appearance have been found to have an impact on the outcome of elections.

Multiple regression—the dominant method in election forecasting—cannot be used to estimate models with many variables and relatively few observations, such as is the case for forecasting U.S. Presidential Elections. In such situations, the *index method* is an attractive alternative. We used the index method to develop a biographical (in the following referred to as *bio*) index for predicting the outcomes of U.S. Presidential Elections.

### **Index method**

Subjective indexes, also known as "experience tables", "unit weighting" (Einhorn & Hogarth 1975), or "Dawes' rule" (Czerlinski et al. 1999), have long been used for forecasting. Analysts prepare a list of key variables and specify from prior evidence whether they are favorable (+1), unfavorable (-1), or indeterminate (0) in their influence on a certain outcome. Alternatively, the scoring could be 1 for a positive position and zero otherwise. Then, the analysts simply add the scores and use the total to calculate the forecast.

The index method has been used for various types of forecasting problems. For example, Burgess (1939) described its use in predicting the success of paroling individuals from prison. Based on a list of 25 factors, which were rated either "favorable" (+1) or "unfavorable" (0), an index score was calculated for each individual to determine the chance of successful parole. This approach was questioned since Burgess (1939) did not assess the relative importance of different variables and no consideration was given to their magnitude (i.e. how favorable the ratings were). However, in addressing these issues, Gough (1962) did not find evidence that supported the use of regression models over index scores.

### ***Index method versus multiple regression***

Einhorn and Hogarth (1975) compared the predictive performance of multiple regression and unit weighting for a varying number of observations and predictor variables. They showed that unit weighting outperforms regression if the sample size is small and the number of—and inter-correlation among—predictor variables is high. Empirical studies have been consistent with this

theoretical result. In analyzing published data in the domain of applied psychology, Schmidt (1971) found regression to be less accurate than unit weighting. In his review of the literature, Armstrong (1985, p.230) found regression to be slightly more accurate in three studies (for academic performance, personnel selection, and medicine) but less accurate in five (three on academic performance, and one each on personnel selection and psychology).

Multiple regression can be useful for estimating the relative influence of causal variables on the outcome variable. Yet the method's ability to incorporate prior domain knowledge is limited. Although regression can use some prior knowledge for selecting variables, the variable weights are typically estimated from the dataset. While this makes multiple regression well-suited for explaining data (i.e., data fitting), it can harm the predictive accuracy of a model. The reason is that, in order to get a better fit, multiple regression often extracts too much information (i.e., noise) from existing datasets, which does not generalize to other datasets. Czerlinski et al. (1999) compared multiple regression and unit weighting for 20 prediction problems (including psychological, economic, environmental, biological, and health problems), for which the number of variables varied between 3 and 19. Most of these examples were taken from statistical textbooks where they were being used to demonstrate the application of multiple regression. The authors reported that, not surprisingly, multiple regression had the best fit. However, unit weighting showed higher out-of-sample predictive accuracy.

Regression modelers face a trade-off between data fitting and prediction. Einhorn and Hogarth (1975) showed that increasing the number of variables decreases a regression model's out-of-sample predictive accuracy given a constant sample size. In order to use more variables, one needs to have a large number of observations. Numerous rules of thumb exist for the necessary ratio of observations to predictors. Based on their analysis of the relative performance of multiple regression and unit weighting for five real social science datasets and a large number of synthetic datasets, Dana and Dawes (2004), found that regression should not be used unless sample size is larger than 100 observations per predictor. Because it is rare to have such large samples per variable in the social sciences, Dana and Dawes (2004, p. 328) concluded that "regression coefficients should almost never be used for social science predictions." Furthermore, we believe that for non-experimental data where the relationships are conditional on a number of factors, it is unlikely that regression can untangle the effects even with massive sample sizes.

### ***When does the index method work?***

Unlike regression, the index method does not estimate weights from the data, so the issue of sample size is not relevant. In using unit or equal weights, the forecaster assesses the directional influence of a variable on the outcome by examining prior research and by using experts' domain knowledge. If little knowledge exists, one might question the relevance of including a variable in the model. Thus, the index method is particularly valuable in situations with good prior domain knowledge.

The index method is not limited in the number of variables that one can incorporate in the model. Furthermore, different variables can be used when forecasting new events. These are

important advantages of the index method as it allows for using all cumulative knowledge in a domain.

In cases involving uncertainty about the relative importance of variables, a good starting point is to use equal weights. If many factors are expected to have an influence on the outcome, having all relevant variables in the model is likely to be more important than their weighting. As knowledge is gained, weights might be used.

In sum, the index method is useful in situations involving many causal variables, a limited number of observations, and good prior knowledge about the influence of the variables on the outcome. In addition, the index method is easier to understand than regression.

Despite growing evidence on the advantages of the index method, few researchers appear to have received the message. In the run-up to a talk at the *2009 International Symposium on Forecasting*, we conducted a small survey among forecasters and asked them for their expectations about the relative performance of the index method, multiple regression, and step-wise regression in situations with a large number of variables and few observations. On average, the 13 experts, who rated themselves as high on 'expertise with forecasting methods', expected regression to yield the most accurate results, followed by the index method, and step-wise regression.

### **Use of the index method in election forecasting**

For forecasting U.S. presidential elections, data for the majority of regression models is limited to about 25 elections. In fact, most models use no more than 15 observations and include from two to sometimes as many as seven explanatory variables (Jones & Cuzán 2008). Given that the number of potential variables is large and the number of observations small, forecasting of U.S. Presidential elections lends itself to the use of index models.

Lichtman (2006) was the first to use the index model to forecast U.S. presidential election winners. His model provided correct forecasts retrospectively for all of 31 elections and prospectively for all of the last 7 elections. No regression model has matched this level of accuracy in picking the winner. This model used the same variables for all elections and was based only on the judgments of a single rater, Lichtman.

Armstrong and Cuzán (2006) transformed Lichtman's model into a quantitative model and compared the derived forecasts against forecasts from three traditional regression models for six U.S. presidential elections from 1984 to 2004. Lichtman's "Keys" performed well, leading to forecast errors almost as low as those of the best regression models. For the 2008 election, the 'Keys' forecast – provided in August 2007, more than a year before Election Day – was again more accurate than the out-of-sample forecasts derived from the same three models and missed the actual outcome by only 0.3 percentage points.

Cuzán and Bundrick (2009) applied an equal-weighting approach to three regression models: Fair's equation (Fair 1978) and two variations of the fiscal model (Cuzán & Heggen 1984). Over 23 elections from 1916 to 2004, the equal weighting scheme outperformed two of the three regression models – and did equally well as the third – when making out-of-sample predictions.

When the authors used data from the 32 elections from 1880 to 2004, they found that equal weighting yielded a lower mean absolute error than all three regression models.

Although these studies demonstrate the value of the index method for forecasting U.S. Presidential Elections, they do not use much prior research for selecting and coding the variables. Furthermore, none of the existing models incorporates information about individual traits that might help candidates to get elected.

### **Predictors of leadership**

A vast amount of literature has analyzed the impact of individual differences on leadership, which can be harnessed to develop index models for forecasting election winners. While a large number of traits have been found to help individuals to emerge as leaders, only few seem to be objectively related to leader performance (Antonakis, in press). Many studies found that voters select their political leaders based on criteria that are irrelevant to performance. Similarly, for the task of predicting election winners, it is the candidates' electability – and not their ability to do the job – that matters most.

Ideally, voters should evaluate candidates along traits that actually matter for leader effectiveness by, for example, selecting the most intelligent candidate. In fact, intelligence has been found to be a major predictor of leadership. Meta-analyses revealed intelligence to be positively correlated with leader emergence ( $r = .5$ , Lord et al. 1986) and leader effectiveness ( $r = .33$ , Judge et al. 2004). For the sample of the 42 U.S. presidents before Barack Obama, Simonton (2006) found intelligence to be positively correlated with presidential performance. However, one might question to what extent voters are actually attracted to highly intelligent candidates, as they might be perceived as being 'out of touch' with the people. Results from the meta-analysis by Judge et al. (2004) support this hypothesis: correlations for intelligence and perceived effectiveness ( $r = .17$ ) were lower than for intelligence and objective effectiveness ( $r = .33$ ).

Thus, it might well be that voters evaluate candidates on traits that have no bearing on leadership performance. For example, candidates' facial competence has been found to be a highly accurate predictor of electoral success. Todorov et al. (2005) presented 31 subjects with pictures of candidates running in U.S. House and Senate elections. Based on one-second exposures, the subjects rated each candidate's competence (subjects who recognized a candidate were excluded). For the three Senate elections from 2000 to 2004, the most competent-looking candidates won 71% of the 95 races. For the two House elections in 2002 and 2004, the most competent-looking candidate won 67% of the 600 races in their sample. In a study by Antonakis and Dalgas (2009), subjects in Switzerland were asked to rate 57 pairs of black and white photos of faces of candidates in the 2002 French parliamentary election (none of the subjects recognized the candidates). In their first experiment, each of 684 university students rated 12 of the pairs of candidates for competency; the candidates with the highest average competency ratings won in 72% of the elections. In their second experiment, they tested Plato's observation by presenting 2,814 children with a pair of photos for a computer-simulated trip from Troy to Ithaca; 72% of the children selected the most

competent looking candidates as their captain. Similarly, Armstrong et al. (2009) found facial competence to be highly predictive for the outcome of the 2008 U.S. Presidential Primaries. In turn, perceptions of leadership might be affected by factors that influence facial appearance such as eyeglasses. In analyzing results from a lab experiment, Thornton (1944) found people wearing eyeglasses were perceived as more industrious, dependable, and honest. Another lab experiment found that eyeglasses enhanced an individual's perceived authority (Bartolini et al. 1988).

Some variables can also influence leader performance as well as voter acceptance. An example is height. In their meta-analysis, Judge and Cable (2004) found height to be positively correlated with social esteem ( $r = .41$ ), leader emergence ( $r = .24$ ), performance ( $r = .18$ ) and income ( $r = .26$ ).

Finally, there might be traits that people do not evaluate when selecting their leaders but that nonetheless have an impact on leader emergence. An example is birth order. Newman and Taylor (1994) analyzed samples of 45 male U.S. Governors and 24 Australian prime ministers. Compared to the population at large, the politicians in both samples were more likely to be first-born and less likely to be middle-born. Similarly, Andeweg and Van Den Berg (2003) analyzed birth-order data for almost 1,200 Dutch politicians. Compared to the general population, they found single children to be overrepresented, whereas middle-children were underrepresented. Another example is the experience of traumatic or adverse events like the early loss of a parent, which is often assumed to contribute to the development of leadership personalities. Simonton (1999) reported on various studies that found the incidence of orphanhood for geniuses from various fields to be higher than in the population at large. For example, one of these studies analyzed a sample of 24 British prime ministers, of which 15 were orphans.

In sum, there is a large body of theory and empirical research that explains and demonstrates the relevance of numerous biographical traits for the emergence of leaders (cf. Appendix 1). Given the vast amount of variables, the index method seems to be the appropriate choice for predicting election winners based on biographical traits. It allows for using extensive prior knowledge for selecting and coding the variables.

### **Biographical index**

Our forecasting environment consists of a set of variables (or cues) that are used to predict the popular vote in elections.

#### ***Variables***

We created a list of 60 variables (or cues) from biographical information about candidates that were expected to have an influence on the election outcome (see Appendix 1). Then, based on prior literature and common sense, we specified whether a cue has a positive or negative influence on the election outcome.

We distinguished two types of cues: (1) Yes / no cues indicate whether a candidate has a certain characteristic or not. (2) More / less cues are more complex as they also incorporate information about the relative value of the cue for the candidates that run against each other in a particular election. Here, the candidate who achieved a more favorable value on a cue was assigned a score of 1 and 0 otherwise. We used two independent coders. If these coders disagreed, a third coder made the final decision. (The final coding is available online at [tinyurl.com/pollybio-coding](http://tinyurl.com/pollybio-coding).) Finally, the sum of cue values for each candidate in a particular election determined his *bio-index score* (B).

### **Data**

We collected biographical data on the candidates of the two major parties that ran for office in the 29 elections from 1896 to 2008. All data referred to the candidate's biography at the time of the respective election campaign. We searched candidate's biographies, fact books, encyclopaedias and used data from earlier studies. For more information see Appendix 1.

### **Performance of the bio-index model**

The bio-index incorporates two ways for predicting the outcome of elections: (1) a heuristic to predict the election winner and (2) a model to predict the popular two-party vote shares of the candidates running for office.

#### ***Predicting the winner – a heuristic based approach***

A simple heuristic was used to forecast the election outcome: the candidate with the higher bio-index score (B) was predicted as the winner of the popular vote. Note that this approach does not require sample size (i.e., information about historical elections). To apply the heuristic, one only has to assess the direction for how a cue will influence the election outcome, assign cue values to the candidates, and then sum them up to calculate the index scores.

Table 1 shows the candidates' index scores in each election year. For the 29 elections, the heuristic correctly predicted the winner 27 times and was wrong twice. Thus, the proportion of correct forecasts (i.e., hit rate) was 0.93. In 1992, it did not predict Bill Clinton to succeed George Bush and, in 1976, it wrongly predicted Gerald Ford to win against Jimmy Carter.

**Table 1: Bio-index scores of presidential candidates (1896-2008)**

(grey= incorrect forecasts)

Election year	Election winner (W)*	Election loser (L)	Index score	
			W	L
1896	McKinley	Bryan	20	14
1900	McKinley	Bryan	21	14
1904	Roosevelt	Parker	24	14
1908	Taft	Bryan	22	16
1912	Wilson	Taft	28	23
1916	Wilson	Hughes	26	20
1920	Harding	Cox	19	14
1924	Coolidge	Davis	23	22
1928	Hoover	Smith	18	14
1932	Roosevelt	Hoover	26	19
1936	Roosevelt	Landon	24	19
1940	Roosevelt	Willkie	23	14
1944	Roosevelt	Dewey	23	16
1948	Truman	Dewey	20	17
1952	Eisenhower	Stevenson	20	15
1956	Eisenhower	Stevenson	21	15
1960	Kennedy	Nixon	28	19
1964	Johnson	Goldwater	24	17
1968	Nixon	Humphrey	22	17
1972	Nixon	McGovern	24	20
1976	Carter	Ford	21	27
1980	Reagan	Carter	22	20
1984	Reagan	Mondale	23	17
1988	Bush H	Dukakis	29	20
1992	Clinton	Bush	23	26
1996	Clinton	Dole	28	17
2000	Gore	Bush	23	20
2004	Bush	Kerry	23	21
2008	Obama	McCain	25	20

\* based on the popular vote

### *Bio-index heuristic versus polls*

Campaign – or trial heat – polls reveal voter support for candidates in an election. Although polls are only assessments of current opinion or ‘snapshots’, their results are routinely interpreted as forecasts and projected to Election Day. For example, the trial-heat forecasting model by Campbell (1996) uses the economic growth rate and Gallup trial-heat polls as its predictor variables. However, polls conducted early in the campaign are commonly seen as unreliable, which is why Campbell adjusts their results according to the historical relationship between the vote and the polls.

We compared the performance of the bio-index to the predicted two-party vote shares from the final pre-election Gallup poll. The Gallup polling data for the 18 elections from 1936 to 2004 were obtained from the Appendix to Snowberg et al. (2007). For the 2008 election, the final pre-election poll was obtained from [gallup.com](http://gallup.com). The hit rate, shown in Table 2, is the proportion of forecasts that correctly determined the election winner. Four times out of the last 19 elections, the

final pre-election Gallup poll predicted the wrong candidate to win the election, which yielded a hit rate of 0.79. By comparison, in failing twice for the same sample of 19 elections, the hit rate of the bio-index heuristic was 0.89.

**Table 2: Hit rate of the bio-index heuristic forecasts (made in January) and benchmark approaches**

Benchmark method	Approx. date of forecast	Sample of Elections	Benchmark method		Bio-index hit rate (same sample)
			Correct forecasts	Hit rate	
Gallup poll	Final pre-election poll	19	15	.79	.89
Prediction markets	Final pre-election market price	26	22	.85	.92
<i>Regression Models</i>					
Abramowitz (1996)	Late July / early August	16	12	.75	.88
Wlezien & Erikson (Wlezien 2001)	Late August	15	12	.80	.87
Campbell (1996)	Early September	16	13	.81	.88

### *Bio-index heuristic versus prediction markets*

Prediction markets to forecast election outcomes were popular in the late 19th century. Rhode and Strumpf (2004) studied historical betting markets that existed for the 15 presidential elections from 1884 through 1940 and found that these markets “did a remarkable job forecasting elections in an era before scientific polling”. In 1988, the Iowa Electronic Market (IEM) was launched as an internet-based futures market in which contracts were traded on the outcome of the presidential election that year. In an initial assessment, the IEM provided more accurate election forecasts than traditional opinion polls. In analyzing 964 polls for the five presidential elections from 1988 to 2004, Berg et al. (2008) found that IEM market forecasts were closer to the actual election results 74% of the time. However, this advantage seems to disappear when comparing the market forecasts to a more sophisticated reading of polls. In analyzing data from the same elections, Erikson and Wlezien (2008) found that polls, which had been combined and damped, were more accurate than both the IEM winner-take-all and the vote-share markets.

We compared the bio-index to prediction market prices from the last day prior to Election Day. Prediction market data were collected from various sources and were available for 26 of the last 29 elections. For the period from 1896 to 1960, we used prices from the historical Wall Street Curb markets as described in Rhode and Strumpf (2004). For the four elections from 1976 to 1988, we used betting odds from British bookmakers. Both datasets were taken from the Appendix to Snowberg et al. (2007). For the last five elections from 1992 to 2008, we used prices from the IEM. (For the three elections from 1964 to 1972, we were unable to obtain prediction market data.) The three datasets were slightly different. While the Wall Street Curb markets and the bookmakers predicted the Electoral College winner, the IEM provided a forecast of the popular vote winner. Nonetheless, each market provided winner-take-all prices. This price can be interpreted as the probability with which the market expects a candidate to win. For example, a market price of \$80

indicates an 80% chance of winning. Thus, if the price of a candidate exceeded 50%, this candidate was predicted to be the election winner. The results are shown in Table 2. The prediction markets got 22 out of the last 26 elections correct which led to a hit rate of 0.85, compared to 0.92 for the bio-index heuristic.

### *Bio-index heuristic versus regression models*

We also compared the hit rate of the bio-index heuristic to three well-established regression models for which we could obtain out-of-sample forecasts, calculated by N-1 cross-validation, for early elections. This means that the forecasters used N-1 observations from the dataset to build the model and then made a forecast for the one remaining election. Abramowitz (1996) and Campbell (1996) published such cross-validated forecasts from 1948; Wlezien and Erikson's forecasts were available from 1952 (Wlezien 2001). For the three most recent elections, ex ante forecasts, published before the actual Election Day, were derived from the authors' respective publications in the elections symposia in *PS: Political Science and Politics*, 34(1), 37(4), and 41(4). The results are shown in Table . In predicting 16 elections, Abramowitz's model failed four times, yielding a hit rate of 0.75. Both Campbell (n=16 elections) and Wlezien and Erikson (n=15 elections) missed the correct winner three times, which led to hit rates of 0.81 and 0.80, respectively. Compared to each of the three models, the bio-index heuristic yielded a higher hit rate.

In sum, none of the benchmark approaches achieved a hit rate as high as the bio-index. This performance was achieved even though our simple heuristic used only information from the respective election year and did not draw on historical information from other elections. Also, note that bio-index forecasts can be made as soon as the candidates are known; they can even be issued before, conditional on who is expected to run for office. By comparison, the forecasts of the three regression models are issued much later, usually around August and September in the election year. Furthermore, even the final Gallup pre-election polls as well as election-eve prediction market prices were less accurate than the bio-index heuristic.

### ***Predicting the vote share***

Although predicting the correct winner of an election might be the most important criteria to assess the performance of a forecasting method, quantitative models usually provide predictions of the actual vote shares. We tested how well the bio-index forecasts the incumbent candidates' percentage of the two-party vote for the past 29 elections.

To do this, we used information from other election years. We used the relative bio-index score (I) of the candidate of the incumbent party as our predictor variable. I is the percentage of cues that favored the candidate of the incumbent party. It is defined as:

$$I = [B_{\text{Incumbent}} / (B_{\text{Incumbent}} + B_{\text{Challenger}})] * 100.$$

We related I to the dependent variable, which was the actual two-party vote share received by the candidate of the incumbent party (V). That is, we used only a single predictor variable to

represent all issues. We performed a simple linear regression by relating V to I for the period from 1896 to 2008 and obtained the following vote equation:  $V = 15.1 + 0.70 * I$ .

Thus, the model predicts that an incumbent would start with 15.1% of the vote, plus a share depending on I. If the percentage of biographical cues favoring the incumbent went up by 10 percentage points, the incumbent's vote share would go up by 7%.

**Table 3: Out-of-sample forecasts of the bio-index model and actual election outcomes  
(grey: incorrect forecasts)**

Election year	Open-seat election	Incumbent's share of two-party popular vote		
		Actual	Predicted	AE
1896	1	47.3	43.5	3.8
1900	0	53.2	57.5	4.3
1904	0	60.0	59.3	0.7
1908	1	54.5	55.7	1.2
1912	0	35.6	47.6	11.9
1916	0	51.6	54.9	3.2
1920	1	36.1	44.7	8.6
1924	0	65.2	50.4	14.8
1928	1	58.8	54.3	4.5
1932	0	40.9	45.2	4.3
1936	0	62.5	53.9	8.6
1940	0	55.0	59.1	4.1
1944	0	53.8	56.6	2.8
1948	0	52.4	53.0	0.6
1952	1	44.5	45.2	0.7
1956	0	57.8	55.9	1.9
1960	1	49.9	42.4	7.5
1964	0	61.3	55.8	5.5
1968	1	49.6	44.0	5.6
1972	0	61.8	53.0	8.8
1976	0	48.9	54.8	5.8
1980	0	44.7	48.7	4.0
1984	0	59.2	55.2	4.0
1988	1	53.9	56.1	2.2
1992	0	46.5	51.8	5.2
1996	0	54.7	59.1	4.4
2000	1	50.3	52.7	2.4
2004	0	51.2	51.8	0.5
2008	1	46.3	46.2	0.1
<b>Sum / MAE</b>	<b>10</b>	-	-	<b>4.6</b>

### *Accuracy of the bio-index model*

We used jackknifing to calculate out-of-sample forecasts for the bio-index model. The results are shown in Table 3. As with the heuristic-based approach, the model-based approach correctly

predicted 27 elections and failed for the elections in 1976 and 1992. Over all 29 elections, the mean absolute error (MAE) of the bio-index model was 4.6 percentage points.

The bio-index model appears to be valuable for predicting who will win. It correctly predicted 27 out of 29 election winners and, as with the heuristic approach, yielded a hit rate that outperformed alternative methods like polls, prediction markets, and other regression models. In addition, it identified the winner in close elections such as 1916, 1948, 1960, 1968, 2000 and 2004. Only in 1976, was the bio-index model unable to correctly predict the close race between Carter and Ford. Forecasting errors were high for the two elections in 1912 and 1924, which were hard to predict as they had a strong third candidate in the race.

### *Bio-index model versus regression models*

To further examine the relative performance of the bio-index model, we generated 'ex ante' forecasts for the last four elections from 1996 to 2008 by successive updating. That is, we only used data from historical elections prior to the respective election year for building the model. For example, to predict the 2008 Election, we used data on the 28 elections from 1896 to 2004. To predict the 2004 Election, we used data on the 27 elections from 1896 to 2000, and so on. The results are shown in Table 4, along with ex ante forecasts from seven well-established regression models. Most of these forecasts were published in *American Politics Quarterly* 24(4) and *PS: Political Science and Politics*, 34(1), 37(4), and 41(4). The forecasts for Fair's model were obtained from his website (fairmodel.econ.yale.edu). For an overview of the predictor variables used in most of the models see Jones and Cuzán (2008).

The bio-index model was accurate when compared to the seven regression models. It yielded a MAE almost as low as the most accurate regression model. Recall that this was achieved with a forecast that was issued long before most of the benchmark forecasts were available.

**Table 4: Bio-index model vs. quantitative models: Errors of out-of-sample forecasts (1996-2008, calculated through successive updating)**

Model	Approximate date of forecast	Forecast error				MAE
		1996	2000	2004	2008	
PollyBio	January, or as (potential) candidates are known	4.5	2.4	0.5	0.1	<b>1.9</b>
<i>Regression models</i>						
Norpoth	January	2.4	4.7	3.5	3.6	<b>3.5</b>
Fair	Late July	3.5	0.5	6.3	2.2	<b>3.1</b>
Abramowitz	Late July / early August	2.1	2.9	2.5	0.6	<b>2.0</b>
Lewis-Beck and Tien	Late August	0.1	5.1	1.3*	3.6	<b>2.5</b>
Wlezien and Erikson	Late August	0.2	4.9	0.5	1.5	<b>1.8</b>
Holbrook	Late August / early September	2.5	10.0	3.3	2.0	<b>4.4</b>
Campbell	Early September	3.4	2.5	2.6	6.4*	<b>3.7</b>

\* incorrect prediction

## Discussion

Candidates' biographies incorporate much information about their chances to win elections. Based on a simple heuristic, the bio-index was able to correctly predict the winner of the popular vote for 27 out of the last 29 U.S. presidential elections; a performance that compared favorably to polls, prediction markets, and three established regression models. This was achieved by assigning unit weights (and directional information) to biographical information from 60 cues. Furthermore, in using it in combination with simple linear regression, the bio-index model was accurate, compared to seven benchmark models, for predicting the four elections from 1996 to 2008.

Bio-indexes do not require sample size for choosing and weighting cues. It is fast and easy to understand. Furthermore, we believe that this approach comes close to how most informed voters process information. They briefly consider the background of the candidates and draw an overall impression. A further advantage is the model's flexibility. Since the index method is (1) not limited in the number of variables and (2) does not weight variables according to their importance, different variables can be used when forecasting new events. For example, for predicting different-gender races, one might want to exclude cues that are only relevant for same-gender races (e.g. height and weight). Furthermore, the index method allows for adding variables once new information becomes available, for example, if a new cue is discovered that is not yet incorporated in the model (e.g., if a candidate was awarded with the Nobel Peace Prize). This is an important advantage as it allows for using all cumulative knowledge in a domain.

### ***When is a bio-index most effective?***

Most traditional regression models consider a presidential election to be a referendum on the president's popularity or, more narrowly, his ability to handle the economy: if the economy is doing well, voters will support the president. Otherwise, they will support the candidate of the other party. For example, Abramowitz (1996) as well as Wlezien and Erikson (Wlezien 2001) use economic growth and presidential approval as predictor variables; Campbell (1996) uses economic growth and measures public support for candidates by incorporating trial-heat polls.

However, in looking at presidential elections as a referendum on the incumbent's performance, forecasting becomes difficult for open-seat elections (i.e., without an incumbent in the race). In general, open-seat elections are considered as harder to forecast. Campbell (2008) compared the outcomes of the 13 open-seat elections to the 22 elections with an incumbent in the race that were held between 1868 and 2004. He found that open-seat elections were more often near dead heats than elections with an incumbent running. Also, out of the 11 elections in his sample that were decided by a landslide, only two were open-seat.

A closer look at the performance of the three models for which we could obtain forecasts for early elections supports the speculation that traditional election forecasting models have difficulties in predicting open-seat elections. All three regression models failed in correctly predicting the winner of the elections in 1960 and 1968; Campbell's model also failed in 2008. Each of these elections was an open-seat election. By comparison, as shown in Table 3, the bio-index correctly

predicted the winner for each of the ten open-seat elections in our sample. The reason might be that bio-indexes do not incorporate a measure that relates to the incumbent president's performance. Although drawing on a small sample, the results suggest that our model is helpful in predicting the outcome of open-seat elections.

### ***Bio-indexes as nomination helper***

Bio-index forecasts can be made as soon as the candidates are known; they can be issued even before, conditional on who is expected to be in the race. Thus, bio-indexes can advise candidates whether they should run for office or not. In addition, it can help parties in nominating their candidates. Parties should select the candidate who achieves a high index score – possibly conditional to a specific opponent. For example, assuming Barack Obama will run again in 2012, the Republican Party should search for a candidate whose index score is similar to or higher than Obama's.

### ***Beauty and resistance of simple models***

Bio-indexes are simple to use and easy to understand. If one only wants to predict the winner, a simple heuristic can be used that does not require information from previous elections. In addition, they can be used in combination with regression models to allow for quantitative predictions. However, a disadvantage is the cost of summarizing knowledge to develop the model.

Unfortunately, the simplicity of the index model may be the method's biggest drawback. Summarizing evidence from the literature, Hogarth (2006) showed that people exhibit a resistance to simple solutions. Although there is evidence that simple models can outperform more complicated ones, there is a belief that complex methods are necessary to solve complex problems.

## **Conclusion**

We applied the index method to the 29 U.S. Presidential Elections from 1896 to 2008 and provided a forecast based on biographic information about candidates. For 27 of the 29 elections, the bio-index model correctly predicted the popular vote winner, a performance that compared well to polls, prediction markets, and three regression models. In addition, its out-of-sample forecasts of the popular vote for the four elections from 1996 to 2008 were almost as accurate as the best model out of seven regression models.

In using a different method and drawing on different information than traditional election forecasting models, we believe our approach will make a useful contribution to forecasting accuracy. It is simple to use and easy to understand. Moreover, it can help political parties in nominating candidates running for office.

## Acknowledgments

John Antonakis, Roy Batchelor, Alfred Cuzán, Ray Fair, Dan Goldstein, Robin Hogarth, Randall Jones, Frank L. Schmidt, Dean K. Simonton and Christopher Wlezien provided helpful comments. We also received suggestions when we presented the paper at the *2009 International Symposium on Forecasting* in Hong Kong and the *Symposium on Leadership and Individual Differences* in Lausanne. We asked authors of key papers whether we cited their work correctly. Thanks to Rudy Andeweg, Alice Eagly, Timothy Judge, John Krantz, Andrew Leigh, James Nelson, Panu Poutvaara and Burt Pryor for responding. Andrew Dalzell, Ishika Das, Rui Du, Max Feldman, Rong Fu, Greg Lafata and Martin Yu helped with collecting data.

## References

- Abramowitz, A. I. (1996). Bill and Al's excellent adventure: Forecasting the 1996 presidential election, *American Politics Research*, 24, 434-442.
- Andeweg, R. B. & Van Den Berg, S. B. (2003), Linking birth order to political leadership: The impact of parents or sibling interaction?, *Political Psychology*, 3, 605-623.
- Antonakis, J. & Dalgas, O. (2009), Predicting elections: Child's play!, *Science*, 323, 1183.
- Antonakis, J. (in press), Predictors of leadership: The usual suspects and the suspect traits, In: Bryman, A., Collinson, D, Grint, K., Jackson, B. & Uhl-Bien, M. (Eds). *Sage Handbook of Leadership*. Thousand Oaks: Sage Publications.
- Armstrong, J. S. (1985). *Long-range forecasting: From crystal ball to computer*, New York: John Wiley.
- Armstrong, J. S. & Cuzán, A. G. (2006). Index methods for forecasting: An application to the American Presidential Elections, *Foresight*, Issue 3, 10-13.
- Armstrong, J. S., Green, K. C., Jones, R. J. & Wright, M. (2009), *Predicting elections from politicians' faces*, Available at <http://mpra.ub.uni-muenchen.de/9150>.
- Bartolini, T., Kresge, J., McLennan, M., Windham, B., Buhr, T.E., & Pryor, B. (1988), Perceptions of personal characteristics of men and women under three conditions of eyewear, *Perceptual and Motor Skills*, 67, 779-782.
- Berg, J., Nelson, F. & T. A. Rietz (2008). Prediction market accuracy in the long run. *International Journal of Forecasting*, 24, 285-300.
- Burgess, E. W. (1939). *Predicting success or failure in marriage*, New York: Prentice-Hall.
- Campbell, J. E. (1996). Polls and votes: the trial-heat presidential election forecasting model, certainty, and political campaigns, *American Politics Research*, 24, 408-443.
- Campbell, J. E. (2008). The trial-heat forecast of the 2008 presidential vote: Performance and value considerations in an open-seat election, *PS: Political Science & Politics*, 41, 697-701.
- Cuzán, A. G. & Bundrick, C. M. (2009). Predicting presidential elections with equally-weighted regressors in Fair's Equation and the Fiscal Model. *Political Analysis*, 17, 333-340.
- Cuzán, A. G. & Heggen, R. J. (1984). A fiscal model of presidential elections in the United States, 1880-1980, *Presidential Studies Quarterly*, 14, 98-108.

- Czerlinski, J., Gigerenzer, G. & Goldstein, D. G. (1999). How good are simple heuristics? In: G. Gigerenzer & Todd, P. M. (Eds.), *Simple heuristics that make us smart*. Oxford University Press, pp. 97-118.
- Dana, J. & Dawes, R. M. (2004). The superiority of simple alternatives to regression for social science predictions, *Journal of Educational and Behavioral Statistics*, 29, 317-331.
- Einhorn, H. J. & Hogarth, R. M. (1975). Unit weighting schemes for decision-making, *Organizational Behavior & Human Performance*, 13, 171-192.
- Erikson, R. S. & C. Wlezien (2008). Are political markets really superior to polls as election predictors? *Public Opinion Quarterly*, 72, 190-215.
- Fair, R. C. (1978). The effect of economic events on votes for president, *Review of Economics and Statistics*, 60, 159-173.
- Gough, H. G. (1962). Clinical versus statistical prediction in psychology. In: L. Postman (Eds.), *Psychology in the making*. New York; Knopf, pp. 526-584.
- Hogarth, R. M. (2006). When simple is hard to accept. In: P. M. Todd & Gigerenzer, G. (Eds.), *Ecological rationality: Intelligence in the world (in press)*. Oxford; Oxford University Press, pp.
- Jones, R. J. & Cuzán, A. G. (2008). Forecasting U.S. presidential elections: A brief review, *Foresight*, Issue 10, 29-34.
- Judge, T. A. & Cable, D. M. (2004). The effect of physical height on workplace success and income: Preliminary test of a theoretical model, *Journal of Applied Psychology*, 3, 428-441.
- Judge, T. A., Colbert, A. E. & Ilies, R. (2004). Intelligence and leadership: A quantitative review and test of theoretical propositions, *Journal of Applied Psychology*, 89, 542-552.
- Lichtman, A. J. (2006). The keys to the white house: Forecast for 2008, *Foresight*, Issue 3, 5-9.
- Lord, R. G., De Vader, C. L. & Alliger, G. M. (1986). A meta-analysis of the relation between personality traits and leadership perceptions: An application of validity generalization procedures, *Journal of Applied Psychology*, 71, 402-410.
- Newman, J. & Taylor, A. (1994). Family training for political leadership: Birth order of United States state governors and Australian prime ministers, *Political Psychology*, 15, 435-442.
- Rhode, P. W. & K. S. Strumpf (2004). Historic presidential betting markets, *Journal of Economic Perspectives*, 18, 127-142.
- Schmidt, F. L. (1971). The relative efficiency of regression and simple unit predictor weights in applied differential psychology, *Educational and Psychological Measurement*, 31, 699-714.
- Simonton, D. K. (2006). Presidential IQ, Openness, Intellectual Brilliance, and Leadership: Estimates and Correlations for 42 U.S. Chief Executives, *Political Psychology*, 27, 511-526.
- Simonton, D. K. (1999), *Origins of genius*, Oxford: Oxford University Press.
- Snowberg, E., Wolfers, J. & E. Zitzewitz (2007). Partisan impacts on the economy: Evidence from prediction markets and close elections, *Quarterly Journal of Economics*, 122, 807-829.
- Thornton, G. R. (1944). The effect of wearing glasses upon judgments of personality traits of persons seen briefly, *Journal of Applied Psychology*, 28, 203-207.
- Todorov, A., Mandisodza, A. N., Goren, A. & Hall, C. C. (2005). Inferences of Competence from Faces Predict Election Outcomes, *Science*, 308, 1623-1626.
- Wlezien, C. (2001). On forecasting the presidential vote, *PS: Political Science and Politics*, 34, 25-31.

## Appendix 1: The cues

Cue No.	Cue	Coded as 1 if candidate (otherwise: 0)	Additional explanation
<b>YES / NO VARIABLES</b>			
<b>Family</b>			
1	Adopted children	Has adopted children	See children. Voters might favor candidates who adopted children.
2	Ancestry	Descends from a presidential family	Descent from renowned families has been shown to have a positive impact on an individual's career chances (Simonton 1984).
3	Children	Has children	Being the social norm to have children, we expected that voters favor candidates who have children.
4	Divorce	Has <i>not</i> been divorced	Although divorces are common, they violate the social norm.
5	Father (political office)	Has a father who held a political office	The role of a candidate's father may have an impact of a candidate's chances to be elected. Similar to Simonton (1981), we quantified whether a candidate's father held one of the offices listed from questions 19 to 32.
6	First born	Is the first-born child in his family	Simonton (1984) summarized research showing that first-born children tend to achieve more than later-born children. Newman and Taylor (1994) analyzed samples of 45 male U.S. Governors and 24 Australian prime ministers. Compared to the population at large, the politicians in both samples were more likely to be first-born and less likely to be middle-born.
7	Single child	Is the single child	Single children were found to have an advantage over children from larger families. For example, Simonton (1981) found a negative correlation between family size and political performance for the 38 U.S. presidents up to Jimmy Carter. Andeweg and Van Den Berg (2003) analyzed birth-order data for almost 1,200 Dutch politicians. Compared to the general population, they found single children to be overrepresented, whereas middle-children were underrepresented.
8	Marriage	Is married	It is the social norm to get married.
<b>Education</b>			
9	College	Went to college	
10	College graduate	Graduated from college	
11	Law degree	Has a Law (J.D.) degree	
12	Master's degree	Has a Master's degree	Similar to Simonton (1981), we quantified the level of formal education and assigned values of 1, if a candidate went to college, graduated from college, obtained a Master's degree, obtained a PhD degree, obtained a Law (J.D.) degree, or worked as a university professor.
13	PhD	Has a PhD / doctoral degree	
14	Professor	Has been a college or university professor	
15	Phi beta kappa	Is member of Phi beta kappa	Similar to Simonton (1981), we quantified whether a candidate was an in-course (not alumnus or honorary) member of Phi Beta Kappa to measure scholastic performance.
16	Prestigious college	Attended an Ivy-League college	
17	U.S. Naval / Military Academy	Went to U.S. Naval / Military Academy	To have an objective and unambiguous criterion for the reputation of a college, we considered all Ivy-League colleges as well as the U.S. Naval and Military Academies as prestigious.
18	Fraternity	Was member of a fraternity	Fraternities often promote the development of leadership as a benefit of membership.
<b>Political life</b>			
19	Attorney General	Is / was U.S. or State Attorney General	
20	City major	Is / was a city major	
21	Election defeat	Has <i>not</i> been defeated in a political election	
22	Governor	Is / was a state governor	Similar to Simonton (1981), we assessed prior political experience by assigning values of 1 if a candidate had occupied one of the offices listed on the left.
23	Judge	Is / was a judge	
24	Lieutenant Governor	Is / was Lieutenant Governor	
25	Solicitor General	Is / was U.S. Solicitor General	
26	State Representative	Is / was a state representative	
27	State Senator	Is / was a state senator	

Cue No.	Cue	Coded as 1 if candidate (otherwise: 0)	Additional explanation
28	U.S. President	Is / was U.S. president	
29	U.S. Representative	Is / was a U.S. representative	
30	U.S. Secretary	Is / was a U.S. Secretary	
31	U.S. Senator	Is / was a U.S. senator	
32	Vice President	Is / was Vice President of the U.S.	
<b>Traumatic / adverse experiences</b>			
33	Disability	Suffers from physical or sensory disability	Simonton (1999, p.115) reported empirical evidence that supports the idea that the development of genius may be enforced by traumatic experiences, particularly in childhood or adolescence. He referred to literature that found people who lost a parent during childhood to be more likely to achieve more in life. Following Simonton (1981), we considered a candidate as an orphan if one (or both) of his parents died before the age of 30. Similarly, we quantified whether a candidate lost one (or more) children, siblings, or a spouse.  Other traumatic experiences that may have a positive impact on leader emergence may be the survival of a major life-threatening disease, physical or sensory disability, or chronic illness in childhood (Simonton 1999, p.115). For the latter, we quantified whether a candidate suffered from chronic illness before the age of 30.
34	Disease survivor	Survived a major life-threatening disease	
35	Chronic illness	Has suffered from chronic illness in childhood or adolescence	
36	Loss of children	Has lost one or more children	
37	Loss of sibling	Has lost one or more siblings	
38	Loss of spouse	Has lost a spouse	
39	Orphanhood	Is an orphan	
<b>Other</b>			
40	Age	Is between 47 and 64 years old	We expected that candidates might have a disadvantage if they are either too young or too old. Prior research supports this assumption for high-level positions in large public firms. In analyzing a sample of more than 10,000 CEOs, Nelson (2005) found that the median age was 57 years, the 10th percentile 47 years, and the 90th percentile 64 years. Interestingly, these numbers conform to the ages of the elected presidents in the 29 elections in our sample. Here, the median was 55 years, the 10th percentile 46.7 years, and the 90th percentile 64.6 years.
41	Athlete	Is known as athletic	In his review of the literature, Stogdill (1948) summarized several studies that found a positive relationship between leadership and athletic ability.
42	Book author	Has authored one or more books	The number of books that a president published prior to be elected has been found to have a positive impact on his political performance (Simonton 1981). Moreover, a publishing record should have a positive impact on the wide recognition of a candidate among voters.
43	Celebrity	Is / was a celebrity in a field other than politics	Being a famous person in a field other than politics should have a positive impact on the wide recognition of a candidate among voters. This can include being a famous actor, athlete, artist, or TV (radio) moderator.
44	Facial hair	Is clean-shaved	Several studies have examined how of facial hair (i.e. clean-shaved, mustache, goateed, beard) influenced perception of people. However, we are not aware of a meta-analysis on this topic. For example, in their experimental study, Terry and Krantz (1993) found beards to be associated with lessened competence. Another experiment, conducted by Shannon and Stark (2003) found the rate of bearded applicants that were selected for management positions to be lower compared to non-bearded applicants. By comparison, an experiment by Reed and Blunk (1990) found consistently more positive perceptions of social/physical attractiveness, personality, competency, and composure for men with facial hair. Given that most politicians, especially in recent years (note that William Taft was the last U.S. president with facial hair), are clean shaved, we expected facial hair to have a negative effect on the evaluation of candidates.
45	Glasses	Wears glasses	In analyzing results from a lab experiment, Thornton (1944) found people wearing eyeglasses to be perceived more industrious, dependable, and honest. Another lab experiment found that eyeglasses enhanced an individual's perceived authority (Bartolini et al.1988). Terry and Krantz (1993) found eyeglasses to be associated with heightened competence but also diminished forcefulness. We expected eyeglasses to have a positive impact on the evaluation of candidates.
46	Hair	Is not bald	Although not identifying a voter bias, Sigelman et al. (1990) found that bald and balding men are underrepresented among governors and Congress members as compared to the general public.
47	Military experience	Has military experience	Similar to Simonton (1981), we quantified whether a candidate has some military experience as wartime recruit, professional soldier, or military general.
48	Military honors	Has been awarded with military honors	We quantified whether a candidate was awarded with military honors.
49	Gender	Is male	In their meta-analysis, Eagly & Karau (1991) found men to emerge more often than women as leaders from initially leaderless groups. This goes back to the fact that leadership is perceived in terms of male stereotypical characteristics, which makes it more difficult for women to

Cue No.	Cue	Coded as 1 if candidate (otherwise: 0)	Additional explanation
			emerge as leaders.
<b>MORE / LESS VARIABLES</b>			
50	Facial competence	Is more competent	Several studies have measured competence ratings based on people's assessments of candidates' headshots (Todorov et al. 2005, Antonakis & Dalgas 2009). It was found that candidates with higher ratings of 'facial competence' were more likely to win elections. Evaluations of facial competence were available for the 2004 (Little et al. 2007) and 2008 elections (Armstrong et al. 2008).
51	First name	Has the more common first name	We expected the candidate with the more common first name to have an advantage. Name popularity was obtained from 1990 U.S. census ( <a href="http://names.mongabay.com">http://names.mongabay.com</a> ).
52	Height	Is taller	Height is a well-known predictor for leadership emergence and performance. In their meta-analysis, Judge & Cable (2004) found physical height to be positively correlated to esteem ( $r=.41$ ), leader emergence ( $r=.24$ ), performance ( $r=.18$ ), and income ( $r=.26$ ). In estimating factors to predict presidential greatness, both McCann (1992) and Simonton (1981) found a positive correlation between height and political performance.
53	Home state	Is from the state with more electoral votes	We expected that a candidate would win the votes of his home state. Thus, the candidate coming from the state with more electoral votes was assumed to have an advantage. The numbers for electoral votes by states in each election were derived from <a href="http://www.archives.gov/federal-register/electoral-college/votes/votes_by_state.html">http://www.archives.gov/federal-register/electoral-college/votes/votes_by_state.html</a> .
54	IQ	Is more intelligent	Results from a meta-analysis showed that intelligence predicts leader emergence (Lord et al. 1986). Simonton (2006) correlated IQ scores for all 42 U.S. Presidents before Barack Obama with evaluations of presidential leadership performance. He found that intelligence is positively correlated with political success. IQ scores for 42 presidents were obtained from Simonton (2006). Where available, we used information from polls by searching the iPoll Databank of the Roper Center.
55	Physical attractiveness	Is more attractive	King & Leigh (2009) assessed the beauty of political candidates from major political parties and then estimated the effect of beauty on vote share for candidates in the 2004 Australian election. They found that beautiful candidates are more likely to win elections. A similar effect was reported by Berggren et al. (2002). In analyzing more than 10,000 visual assessments of almost 2,000 Finnish political candidates, the authors a positive relationship between attractiveness and the received vote share of candidates. Attractiveness scores for 39 presidents were obtained from Simonton (1986). The coding for the 1920 election race between Harding and Cox was based on Gladwell (2005). Where available, we used information from polls by searching the iPoll Databank of the Roper Center.
56	Race	Represents the larger race	We expected that voters would tend to endorse a candidate that represents their race. Thus, we assumed the candidate that represents the larger race to have an advantage. Also, in analyzing ballot photographs for low-information elections, Banducci et al. (2008) found that the probability of winning for white candidates is 38% greater than for nonwhite candidates.
57	Religious affiliation	Is affiliated with the larger religion	We expected that voters would tend to endorse a candidate that identifies with their religious beliefs. Thus, we assumed the candidate that identifies himself with the larger religion to have an advantage.
58	Surname	Has the more common surname	We expected the candidate with the more common surname to have an advantage. Name popularity was obtained from 1990 U.S. census ( <a href="http://names.mongabay.com">http://names.mongabay.com</a> ).
59	Voice	Has the more dominant voice	Gregory & Gallagher (2002) analyzed the acoustic frequency of candidates' voices in presidential debates. They found that this nonverbal vocal communication reveals social dominance and thus can be helpful to predict the popular vote. We used data from the eight elections in their sample for our analysis.
60	Weight	Is heavier	In his review of the literature, Stogdill (1948) provided evidence that weight is positively correlated with leadership ( $r = .23$ ): seven studies found leaders to be heavier, whereas two studies found leaders to be lighter; two studies found no difference.

## Appendix 2: References for Appendix 1, not cited in the main paper

- Banducci, S. A., Karp, J. A., Thrasher, M. & Rallings, C. (2008), Ballot photographs as cues in low-information elections, *Political Psychology*, 29, 903-917.
- Berggren, N., Jordahl, H., & Poutvaara, P. (2002), The looks of a winner: Beauty, gender, and electoral success, *CESifo working paper*.
- Eagly, A. H. & Karau, S. J. (1991), Gender and the emergence of leaders: A meta-analysis, *Journal of Personality and Social Psychology*, 60, 685-710.
- Gladwell, M. (2005), *Blink: The power of thinking without thinking*, London: Little, Brown and Company.
- Gregory, S. W. & Gallagher, T. J. (2002), Spectral analysis of candidates' nonverbal vocal communication: Predicting U.S. Presidential Election outcomes, *Social Psychology Quarterly*, 65, 298-308.
- King, A. & Leigh, A. (2009), Beautiful politicians, *Kyklos*, 62, 579-593.
- Little, A. C., Burriss, R. P., Jones, B. C. & Roberts, S. C. (2007), Facial appearance affects voting decisions, *Evolution and Human Behavior*, 28, 18-27.
- McCann, S. J. (1992). Alternative formulas to predict the greatness of U.S. presidents: Personological, situational, and zeitgeist factors, *Journal of Personality and Social Psychology*, 62, 469-479.
- Nelson, J. (2005), Corporate governance practices, CEO characteristics and firm performance, *Journal of Corporate Finance*, 11, 197-228.
- Reed, J. A. & Blunk, E. M. (1990), The influence on facial hair on impression formation, *Social Behavior and Personality*, 18, 169-175.
- Shannon, M. L. & Stark, C. P. (2003), The influence of physical appearance on personnel selection, *Social Behavior and Personality*, 31, 613-623.
- Sigelman, L., Dawson, E., Nitz, M. & Whicker, M. L. (1990), Hair loss and electability: The bald truth, *Journal of Nonverbal Behavior*, 14, 269-283.
- Simonton, D. K. (1981), Presidential greatness and performance: Can we predict leadership in the White House?, *Journal of Personality*, 49, 306-322.
- Simonton, D. K. (1984), *Genius, creativity, and leadership*, Cambridge, MA: Harvard University Press.
- Simonton, D. K. (1986), Presidential personality: biographical use of the Gough Adjective Check List, *Journal of Personality and Social Psychology*, 51, 149-160.
- Stogdill, R. M. (1948), Personal factors associated with leadership: A survey of the literature, *Journal of Psychology*, 25, 35-71.
- Terry, R. L. & Krantz, J. H. (1993), Dimensions of Trait Attributions Associated with Eyeglasses, Men's Facial Hair, and Women's Hair Length, *Journal of Applied Social Psychology*, 23, 1757-1769.