Elections, Ability, and Candidate Honesty

Jonathan Woon*
University of Pittsburgh

Kristin Kanthak†
University of Pittsburgh

February 27, 2016‡

Abstract

An important function of elections is to select the best representative, a task facilitated when candidates are honest about their qualifications. But are they? To what extent do candidates' claims depend on the alignment of incentives between themselves and voters? We conduct an incentivized laboratory experiment in which candidates can choose to be honest or to exaggerate, varying the benefits of winning office. We find that strong office motives clearly induce exaggeration and, surprisingly, that only about half of laboratory candidates tell the truth even when incentives are completely aligned. We show that the prevalence of lying in elections results not from impure (e.g., Machiavellian) motives, but rather as a rational response to the expectation that other candidates will lie. Although honesty and integrity are desirable virtues in elected officials, our experiment suggests that the nature of electoral processes can make dishonesty endemic to the democratic selection of leaders.

*Associate Professor, Department of Political Science, Department of Economics (secondary), and Pittsburgh Experimental Economics Laboratory, woon@pitt.edu. (Corresponding author)

†Associate Professor, Department of Political Science, kanthak@pitt.edu

‡This research was supported by the National Science Foundation under Grant Number SES-1154739. We thank Danielle Thomsen, Lise Vesterlund, Alistair Wilson, PEEL group members, and seminar participants at the Columbia Political Economy Seminar for helpful comments and suggestions, P.J. Healy for sharing his belief elicitation instructions, and David Klinowski and Shane Redman for their capable research assistance. Previous versions were presented at the 2013 North American Economic Science Association Conference and the 2015 Midwest Political Science Association Conference

Governing is a skill much like any other: Politicians and public officials vary in how much of that skill they possess. Some have great honesty, integrity, and competence. Some are effective at solving problems and building consensus, while others are committed to the public good and inspire trust and confidence. Great leaders have all of these characteristics, and bad ones lack them. An important function of democratic elections is to select the best, most virtuous, leaders (Besley 2005). And indeed, democratic elections seem to do just that (Alt, Bueno de Mesquita and Rose 2011, Besley, Montalvo and Reynal-Querol 2011, Besley and Reynal-Querol 2011, Buttice and Stone 2012, De Paola and Scoppa 2010, Galasso and Nannicini 2011, Hirano and Snyder 2014, Jones and Olken 2005, Mondak 1995, McCurley and Mondak 1995, Stone, Maisel and Maestas 2004).

The extent to which elections facilitate the selection of the best leaders depends, in large part, on the ability of voters to discern good candidates from bad.[1] The voter's problem is made complicated, however, by the fact that candidates do not always communicate their qualifications or accomplishments honestly. Candidates in the United States, for example, have a tendency to embellish their military records. In one notable case, Representative Bruce Caputo withdrew from challenging Daniel Patrick Moynihan's re-election to the Senate when the media exposed inconsistencies about his claim to have served in Vietnam.[2] However, similar exaggerations by Richard Blumenthal did not impede his election to the Senate in 2010.[3] Governors, too, inflate their accomplishments by varying degrees: Jeb Bush "cut taxes by $19 billion" in Florida, Mike Huckabee "raised average family income by 50%" in Arkansas, and John Kasich "took the state of Ohio from an $8 billion hole...to a $2 billion surplus."[4]

[1]Career decisions and self-selection into politics are also important factors that affect quality selection. For example, see Caselli and Morelli (2004), Gagliarducci, Nannicini and Naticchioni (2010), and Messner and Polborn (2004) for theoretical models. Also see Kanthak and Woon (2015) for an experimental investigation of self-selection in the context of gender and diversity.

[2]Caputo claimed to have been drafted and to have served as lieutenant in the Army, when he had instead avoided military service by working as a civilian at the Pentagon. The inconsistency was pointed out to the media by Tim Russert, then Moynihan's campaign manager. See Maurice Carroll, "Caputo Quitting Race for Senate Over Inaccuracy" (*New York Times*, March 9, 1982) and Richard Stengel, "Tim Russert: A Man This Good is Hard to Find" (*Rolling Stone*, February 14, 1985).

[3]http://www.nytimes.com/2010/05/18/nyregion/18blumenthal.html

[4]The *Washington Post* gave Bush's claim 3 Pinocchios (http://www.washingtonpost.com/news/fact-checker/wp/2015/10/07/jeb-bushs-misleading-claim-that-i-cut-taxes-by-19-billion-over-eight-years-

Formal models of campaigns and elections demonstrate that office-seekers frequently have incentives to obfuscate: about their policy positions (Callander and Wilkie 2007), their opponents (Davis and Ferrantino 1996), and their plans for office (Corazzini et al. 2014).[5] Empirically, however, candidate honesty has received relatively little attention in the political science literature, largely because of the difficulty of collecting large-scale data on moments of obfuscation.[6] In overlooking candidates' strategic incentives to lie, scholars have ignored an important reason why voters would not bother to become informed. If there are no real means by which voters can discern truth from candidates' statements, then there is little reason to know what those statements are.

The central issue we address is one of strategic communication in which the electoral environment itself, not the moral failings of candidates, may be a significant source of lying behavior. Of course, real world politicians may be intrinsically dishonest, as many people perceive them to be, but by turning to laboratory experiments we can minimize intrinsic dishonesty and focus specifically on investigating how electoral context affects the strategic choices of candidates to hew to the truth or exaggerate. Under what conditions do candidates in elections campaign honestly? To what extent do elections and electoral incentives encourage candidates to exaggerate their qualifications?

We answer these questions using an incentivized laboratory experiment in which candidates can choose to be honest, to obfuscate, or to exaggerate. In our experiment, subjects complete a problem-solving task, and their performance on this task serves as a measure of their quality, which is privately known. Subjects then choose campaign messages

fact-checker-biography/), while the website Politifact rated Huckabee's claim as "mostly false" (http://www.politifact.com/truth-o-meter/statements/2015/may/04/mike-huckabee/mike-huckabee-says-he-raised-average-family-income/) and Kasich's claim as "mostly true." (http://www.politifact.com/truth-o-meter/statements/2015/aug/06/john-kasich/kasich-i-took-state-ohio-8-billion-hole-2-billion-/).

[5]In Corazzini et al. (2014), campaign promises serve as commitments that may actually alter actions once in office despite being cheap talk. In Callander and Wilkie (2007), the extent to which some candidates' willingness to lie undermines elections is limited by the fact that voters recognize this problem and make appropriate inferences. Other models of lying and lying aversion in candidate positioning include Banks (1990), Kartik and McAfee (2007), and Huang (2010).

[6]The few observational studies of lying in politics we are aware of include Pfiffner (1999), Bucciol and Zarri (2013), Ringquist and Dasse (2004), and Armstrong-Taylor (2012).

that either honestly convey information about their performance or distort their performance through ambiguity or exaggeration. Because these messages are the only information that voters have about candidates, the information transmitted in the election depends crucially on candidates' message strategies. We find that a substantial proportion of candidates lie, even when there is no private incentive to do so. Our analysis suggests, however, that such lies are not attempts to mislead voters. Instead, they are rational responses to their beliefs that other candidates will be lying. These findings allow us to reconceptualize lying on the campaign trail not as a moral failing of an unsavory political class, but as a strategic response to electoral politics.

The details of our experiment are as follows: In the baseline condition, candidates and voters are rewarded only for the quality of the elected leader. The strategic problem they face is purely informational and revolves around coordination. Given the complete alignment of incentives, honest campaigning transmits information accurately and facilitates the common goal of identifying the best representative. However, honesty is not necessarily the best policy—because our election game admits multiple equilibria, honesty is rational only if all other candidates are also honest. Exaggeration can be individually and socially rational for highly qualified candidates (given our communication technology) if they believe that other candidates, especially low quality candidates, are likely to exaggerate. Although theory provides some basis for forming expectations in our setting, experimental analysis is ideal for understanding how real people actually behave when there is tension between the virtue of honesty and the social good of selecting the most qualified leader.

We manipulate the benefit of winning the election to include a private bonus from winning in our treatment conditions. In these settings, the incentives for lying are much clearer. If messages are taken at face value, then lower quality candidates personally gain by exaggerating their qualifications at the expense of voter welfare. In addition, we vary whether the benefits of winning accrue solely to the individual candidate or to the candidate's group. This manipulation varies other-regarding or group-based interests while holding individual

incentives constant, and it allows us to test whether candidates who may be reluctant to lie for selfish reasons might instead lie for social (albeit parochial) ones.

Our experimental study yields two key findings. First, we find that the degree of honesty in laboratory elections indeed varies with office incentives. When candidates stand to earn individual benefits of office, they are more likely to exaggerate than they are in the baseline condition. We find no difference between the individual and group bonuses in terms of overall exaggeration, but we do find a stronger relationship between task ability and exaggeration in the group bonus condition than in the individual bonus condition.

Second, and more surprisingly, we find that in the absence of office benefits—when candidates and voters share the common goal of selecting the best representative—only half of candidates are honest, while the other half exaggerates. In other words, candidates lie when deception provides no obvious, direct benefit. This finding is striking because in many laboratory experiments on communication, people exhibit an aversion to lying even if they are rewarded for doing so (Gneezy 2005, Gneezy, Rockenbach and Serra-Garcia 2013, Hurkens and Kartik 2009, Lundquist et al. 2009, Sanchez-Pages and Vorsatz 2007, Wang, Spezio and Camerer 2010).[7] Our results also stand in contrast to laboratory studies of elections that find that candidates and political decisionmakers refrain from pandering to voters (Landa and Duell 2014, Woon 2012, 2014). In these studies, laboratory politicians can gain sizable electoral rewards from exploiting voters' concern about selection under conditions of incomplete information, but instead they forgo opportunistic gains and act in voters' best interests. Whereas previous studies find that experimental subjects are reluctant to lie when they would be rewarded for doing so, we find the opposite. In our baseline elections, there is no obvious incentive for candidates to lie, yet they do so anyway.

Given that there are multiple equilibria in our game, a plausible interpretation of our finding is that the level of candidate exaggeration we observe is consistent with instrumental

[7]Other studies of experimental sender-receiver games find that while senders are more than willing to exaggerate, they end up overcommunicating because they fail to conceal their information optimally (Cai and Wang 2006, Minozzi and Woon 2013).

rationality. That is, when candidates agree that elections should be informative, exaggeration is a rational best response when others are exaggerating as well. Hence, lying behavior need not betray the baser instincts of candidates engaged in politics as usual, but possibly reveals honest intentions tinged with a degree of strategic distrust of others. To investigate more carefully the source of candidate exaggeration, we take advantage of the measurement tools that the laboratory provides us. Specifically, we use an incentivized belief elicitation procedure and contrast beliefs against a personality-based explanation for lying. We find that exaggeration in the baseline can be explained primarily by beliefs: The probability of exaggeration increases significantly with the strength of a candidate's belief that the opposing candidate exaggerated. In contrast, "Machiavellian" personality traits (such as the desire for control or status, e.g., Dahling, Whitaker and Levy 2008) are unrelated to lying behavior. That we find a greater degree of lying than in the existing literature may be due to the specific strategic and competitive environment of electoral politics that features widespread expectations of dishonesty.

Experimental Procedures

In the experiment, subjects first completed a skill-based problem-solving task, received individual feedback about their performance on the task, and then made a series of political decisions. Task performance can be thought of as a laboratory analogue of policy-making skill or competence. This skill is valuable to all group members, but it is privately known only to each subject. The political decision in this experiment is the subject's choice of campaign message, and the type of message that a subject chooses affects the information conveyed to other subjects. Importantly, we vary the incentives for winning the election, and subjects choose their campaign messages under different conditions.

We conducted the experiment at the ***** in 12 sessions with a total of 120 participants (61 women and 59 men, 10 subjects per session). Subjects were assigned an ID number

Figure 1: Sample screenshot of the anagram task

and were identified in all their interactions with other subjects only by the ID number. After reading the instructions aloud, subjects were randomly placed in one of two groups. They did not know which other subjects were in their group, nor did they have any information about their fellow group members at the start of the experiment. All interactions between subjects took place through a computer interface programmed in z-tree (Fischbacher 2007). Subjects were also informed that the experiment consisted of several parts, that we would randomly select one part for payment, and that they would receive written instructions before beginning each part. We did this so that the anticipation of later parts of the experiment would not influence choices in earlier parts.

The task we used in the experiment is a simple anagram task, which involves unscrambling letters to form words. Subjects received instructions about how to perform the anagram task and were allowed to practice with the graphical interface (shown in Figure 1).[8] Each anagram consisted of five jumbled letters that could be rearranged to form a common word, with each set of letters having a unique solution (e.g., the solution to the anagram in Figure 1 is "walks"). In Part 1 of the experiment, subjects had four minutes to solve as many of the anagrams as they could and were paid 75 cents for each anagram puzzle they could solve in the four minutes. (The payment incentivizes effort.) There were 40 puzzles in total

[8]Detailed instructions can be found in the Appendix.

and all subjects saw the puzzles in the same order.[9] Subjects could move to a new puzzle by submitting an incorrect guess and incorrect guesses did not count against their winnings. At the end of Part 1, subjects received individual feedback about the number of anagrams they solved correctly but did not learn how well any of the other subjects performed on the task.

In Part 2 of the experiment, subjects participated in a series of elections to select a representative and knew that they would complete the anagram task again after completing their election choices. Each set of elections involved a different set of incentives (explained below), and in each set, every subject would be a candidate in at most one election and a voter in every election. We elicited choices in a way analogous to the "strategy method" used in many incentivized laboratory experiments (Morton and Williams 2010, pp. 87-90). Specifically, we explained the rules for the first set of elections (Elections 1-5), subjects made their choices as the candidate in these elections, then we explained the rules for a second set (Elections 6-10), and then subjects made their choices for the second set; only after all candidate made their choices did we hold the elections, and then subjects completed the four-minute anagram task again.[10] This order of events is crucial for preventing subjects from learning anything about the distribution of anagram scores as candidates or from learning anything more about their own task ability from completing the task again between elections. If we selected Part 2 for payment, then we also randomly selected one of the elections to count. Thus, each subject made choices as a candidate under different conditions, but in the end we would only select (at most) one representative and subjects completed the anagram task only once after the elections.

In each of the first five elections (Elections 1-5), one member from each group was randomly selected to face off as the candidates. As a candidate, each subject chose one

[9]We chose 40 puzzles because the maximum score in a pilot study was 20 correct puzzles. While there might be some potential for censoring because subjects know the maximum possible score, such concerns are negigible. Only one subject had a score higher than 34 (it was 35) such that telling the big lie would reveal the message to be a lie.

[10]There was an additional election, Election 11, but we do not analyze it in this paper.

of four possible campaign messages and knew that the message was all that other subjects would know about them during their election.[11] The message could be:

Ambiguous – "I solved a lot of puzzles"

Truthful – "I solved X puzzles"

Small Lie – "I solved $X + 3$ puzzles"

Big Lie – "I solved $X + 6$ puzzles"

where X is the true number of puzzles the subject solved in Part 1 of the experiment.[12] We did not use any of the terms "ambiguous," "truthful," "small lie," or "big lie" to describe the messages, but instead referred to them in abstract terms as "Message 1," "Message 2," and so on. Note that *all* subjects made a choice about which campaign message to use because every subject was a candidate in one of the elections.[13]

If we selected one of Elections 1-5 to count, we paid subjects 50 cents for each correct answer the representative provided and 25 cents for each of their own correct answers. Subjects received no bonus for winning the election, and we refer to this first set of elections as the *Performance Pay* (baseline) condition. In this condition, subjects have common incentives, regardless of their group membership, to identify and select the most competent representative.

In Elections 6-10, subjects chose from the same set of campaign messages, but with one important difference in the incentive structure. In addition to the same performance incentives as in Elections 1-5 (50 cents for each of the representative's correct answers and 25

[11] Candidates were randomly assigned letters, which were uncorrelated with group membership or ID numbers. In this way, subjects had no information about the identity or group membership of each candidate

[12] For example, if the subject solved 10 puzzles and chose the Big Lie, voters would see the message "I solved 16 puzzles."

[13] It is important to note that our experiment does not include any means by which voters can confirm the veracity of candidates' claims. This is by design, which reflects the fact that in actual elections, the claims candidates make are, by their very nature, largely unverifiable. Such claims include, for example, that the candidate is "working hard to bring government money to the district" or promises of "new leadership" on complex issues (Tomz and Van Houweling 2009, Grimmer 2013, Grimmer, Westwood and Messing 2015).

cents for each of the subject's own correct answer), the winning candidate received a bonus of $2.50 if we selected their election for payment. Thus, in Elections 6-10, there is a bonus for winning the election, which we can think about as an office-holding benefit.

We varied the exact nature of the bonus across sessions. In the *Individual Bonus* condition only the candidate who won received the $2.50 while in the *Group Bonus* condition all members of the winning candidate's group received the $2.50 bonus. We designed the different bonus conditions so that a candidate's individual benefit is constant across all sessions while the amount of the group's benefit varies by condition. From the perspective of a purely selfish, individual payoff-maximizing candidate, the two conditions are identical, but from the perspective of a candidate with pro-social motivations, the Group Bonus condition provides a stronger incentive for winning the election. By varying the bonuses in this way, we could see how tapping into other-regarding preferences might affect candidates' message choices.

Hypotheses

We begin our analysis by considering the strategic incentives candidates face in the Performance Pay condition (our baseline). In this condition, subjects earn twice as much for the elected representative's performance as for their own, so if we reasonably assume that subjects care only about their own monetary payoffs, there is a common incentive for members of both groups in the experiment to elect the best possible representative (i.e., the highest performing member regardless of group membership). That is, individual and group incentives are aligned. For purposes of exposition, let us also assume that a subject's score in the individual anagram task (Part 1) is a perfect indicator of their underlying task ability. Because these abilities are privately known, subjects can use campaign messages to share and transmit information, thereby making abilities public knowledge and allowing the electorate to correctly identify and elect the best representative.

Despite the fact that campaign messages have literal meanings, they remain cheap talk, so a message's correct interpretation depends critically on subjects' expectations of the behavior of others (e.g., Crawford and Sobel 1982, Crawford 2003, Kartik and Van Weelden 2014). In our setup, sharing information entails solving a coordination problem: All candidates must choose exactly the same type of message for elections to be perfectly informative. For example, it is reasonable to believe that since incentives are aligned, there is no incentive to lie. Thus, if every subject chooses the truthful message, then it will be obvious to everyone else which of the two candidates is the highest performer in every election, as the ranking of candidates according to their messages reflects their true ranking. In game theoretic terms, such behavior constitutes a Nash equilibrium because as long as everyone else is being truthful, there is no incentive to do otherwise.

Although there is no incentive to lie, there is no real incentive to tell the truth either. That is, the incentive to tell the truth is not a strict one, as there are multiple equilibria. There is also a Nash equilibrium in which everyone tells a small lie and another in which everyone tells a big lie. This is because the actual content of the message does not matter, so long as campaign messages correctly reveal the relative ordering of candidates' ability. This ordering is preserved if everyone lies by the same amount. Thus, when everyone shares the common goal of selecting and identifying the best possible representative, any campaign message (other than the ambiguous message) can be supported as "rational" behavior. Of course, the most natural or focal belief arises when subjects reason that since everyone has a common incentive to select the best representative, then there is no reason for lying and therefore everyone is likely to tell the truth.

Hypothesis 1. *Preference alignment and common norms of honesty imply that an equilibrium with truthful messages will be focal. Truthful messages will be the most likely type of message in the Performance Pay condition.*

Truthful campaign messages depend on the expectation that others will be truthful, but what if subjects harbor doubts about the honesty or motives of other candidates? For

example, there might be subjects who care about and enjoy winning the election, even if doing so prevents the highest ability candidate from becoming the representative, leading to lower monetary payoffs for everyone. The expectation that other candidates might exaggerate will lead rational—and otherwise honest—candidates to exaggerate as well.

To see this, consider the following example. Suppose that candidate A believes she has a higher score than others and that others lie with positive probability. Also suppose that the candidate who claims the highest score wins the election; since voters cannot observe the truth, they simply take the messages at face value. More concretely, suppose A's score is 10 and that B is the other candidate, whose score is 8. If A tells the truth and B exaggerates, then A will lose the election since her message will be 10 and B's message will be either 11 or 14. The inferior candidate would therefore win the election. But if A exaggerates, she is more likely to win the election—making everyone better off. Indeed, she has a stronger incentive to tell the big lie (a message of 16, which guarantees she wins the election) than she has to tell the small lie (a message of 13, which still allows the inferior candidate to win by telling the big lie). Furthermore, even if A lies but B tells the truth, then there is no harm since the better candidate wins the election, again making everyone better off. However, this incentive to lie, given the belief there is some chance that others lie, depends on whether A believes her score to be higher than that of B. If A's score is lower, then she is better off telling the truth to ensure that the higher-scoring B wins. Thus, exaggeration is "rational" only if a subject believes that others exaggerate and is sufficiently confident that she has a higher score than others.

Hypothesis 2. *Candidates who are more confident in their abilities than others, or who believe that other candidates are likely to exaggerate, are themselves more likely exaggerate their qualifications.*

As one might imagine, introducing an individual-level bonus for winning the election (as in the Individual Bonus and Group Bonus conditions) naturally provides a greater in-

centive for lying, even for low or moderate ability candidates. The availability of the bonus disrupts the alignment of preferences since some subjects would prefer to obtain the bonus even if doing so prevents the selection of the best ability representative. Specifically, since the \$2.50 bonus for winning is equivalent to a difference in 5 anagrams solved correctly, anyone who believes her own score is within 4 correct answers of the highest score has an incentive to tell the big lie in order to win the election. For example, if a subject's score is 10, she would earn $10 \times \$0.75 + \$2.50 = \$10$ for winning the election herself, whereas if the best score is 14, the payoff from electing the best representative would be $14 \times \$0.50 + 10 \times \$0.25 = \$9.50$. In other words, the individual bonus from winning the election provides an incentive to exaggerate simply to deceive others in the hopes of winning the election.

Hypothesis 3. *Candidates are more likely to exaggerate their qualifications in the Bonus Conditions than in the Performance Pay Condition.*

We conjectured that candidates who are averse to lying might be more willing to exaggerate when winning the election benefits others rather than just themselves. Such subjects are considered other-regarding in the sense that they put weight on the payoffs of other members of their group when evaluating their choices. We might expect these subjects to be more likely to exaggerate on behalf of others in the Group Bonus condition than they are to exaggerate solely on their own behalf in the Individual Bonus condition.

Hypothesis 4. *If there are candidates who are other-regarding in the sense that they seek to maximize the payoffs of their fellow group members, then such candidates are more likely to exaggerate in the Group Bonus condition than in the Individual Bonus condition.*

We should note that we are least confident about the prediction stated in Hypothesis 4 for several reasons. Other-regarding preferences in terms of group welfare is only one form of altruism or social preference, and groups in our setting are extremely minimal. There is no opportunity to communicate or form a group identiy or cohesion, so it may be that subjects don't particularly care about the payoffs of other group members in the experiment. Another

reason for our hesitation is the possibility that in activating or priming other-regardingness, the group bonus also activates other forms of social preference. For example, it might strengthen norms of honesty or enhance lying aversion in some candidates. Overall, our prior expectations regarding Hypothesis 4 are fairly weak.

Results

Performance on the anagram task varied sufficiently across subjects to make the selection of a quality representative a reasonably difficult coordination problem. Figure 2 shows kernel densities of performance on the anagram task in both parts of the experiment (before and after the elections). On average, subjects correctly solved 15.6 anagrams in the first task, with a minimum of 2, a maximum of 35, and a standard deviation of 6.4. In the second task, the average was 14.4, the minimum was 2, the maximum was 34, and the standard deviation was 5.7.

Our central research question involves subjects' choices of campaign message under different electoral incentives, and specifically, the effect of office-motives on the extent of

Figure 2: Task performance (kernel densities)

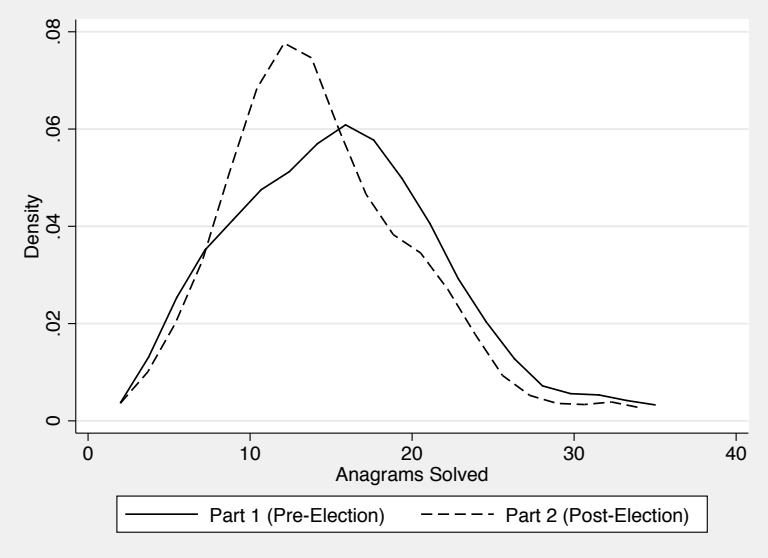

Figure 3: Messages in Performance Only Condition

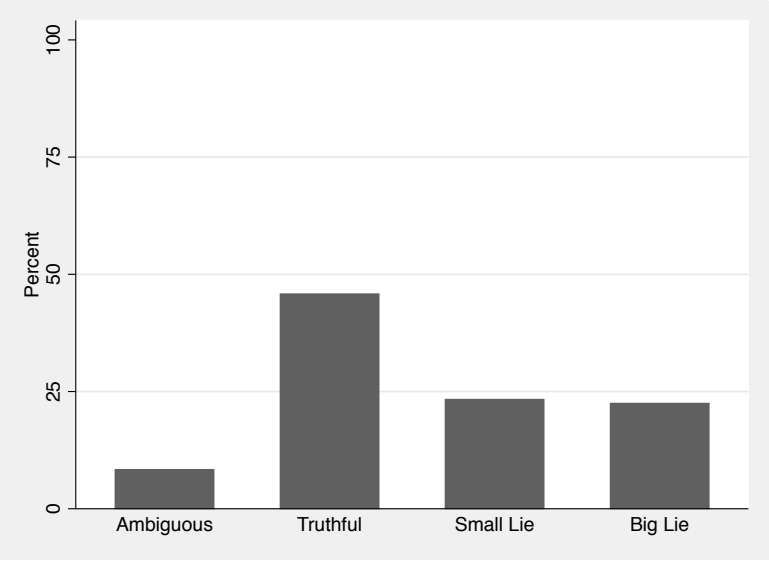

deception and exaggeration. Will subjects tell the truth? Or will they exaggerate their qualifications to try to win office? Figure 3 shows the distribution of messages in the Performance Pay condition. Recall that in this baseline condition, subjects' incentives are aligned to elect the highest-performing candidate. Consistent with Hypothesis 1, truthtelling is the modal message type. However, coordinating on the truth is by no means universal. While 46% of subjects told the truth about their scores, an equal proportion exaggerated (almost equally divided between small lies and big lies). Ambiguity was the least popular message type, at 8%. What is striking about this result is that even when candidates have common incentives to share information (and hence to refrain from using deception), a substantial proportion of them lie.

When we motivate subjects by giving them explicit bonuses for winning, we find, not surprisingly, a substantial shift in campaign message behavior from truthtelling to exaggeration. This result supports Hypothesis 3. Figure 4 shows that in both bonus conditions, the big lie is the modal message type (58% in the Individual Bonus condition and 40% in the Group Bonus condition). Indeed, exaggeration—in the form of either the small lie or big lie—accounts for a substantial majority of messages (80% and 61%, respectively) when there is a bonus for winning. At the same time, truthtelling declines (to 17% and 28%).

Figure 4: Messages in Bonus Conditions

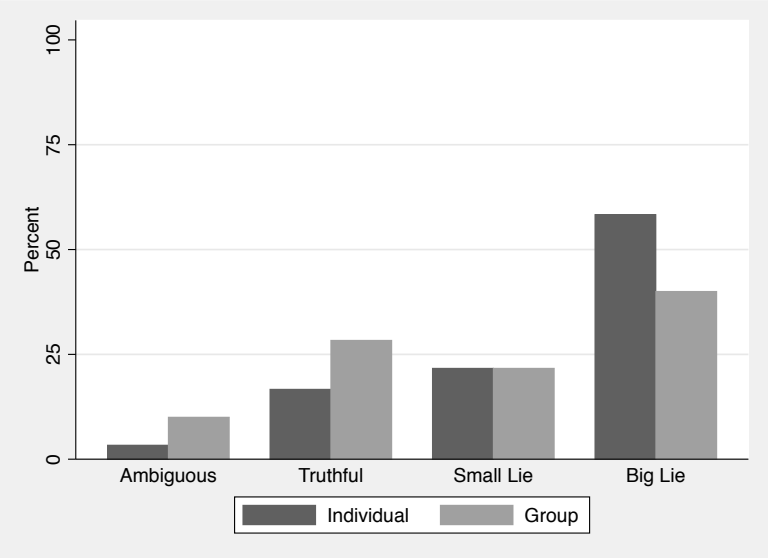

Recall that an individual subject's reward for winning the election is the same across both bonus conditions and that the group bonus condition was designed to activate other-regarding preferences. Figure 4 shows that there are more truthful messages and fewer big lies in the Group Bonus condition than in the Individual Bonus condition. Although Pearson's chi-squared test cannot reject the equality of the overall distributions at conventional levels of significance ($\chi^2 = 5.866, p = 0.12$), we do find that there is a significant difference between the proportion of big lies ($\chi^2 = 4.035, p = 0.05$). The data therefore contradict our expectation in Hypothesis 4 that the Group Bonus would activate social preferences and increase candidates' lying on behalf of their group. The data suggest instead that the Group Bonus condition prompts increased lying aversion, perhaps because subjects think more directly about the costs and benefits to others.

Turning now to the relationship between candidate quality and the honesty of their campaign messages, Figures 5 and 6 present box and whiskers plots summarizing the distributions of scores by message type the Performance Pay and Bonus conditions. Figure 5 suggests that message behavior may be weakly related to underlying task performance. Candidates who choose the big lie as their campaign message have a higher median task score

Figure 5: Scores in Performance Only Condition

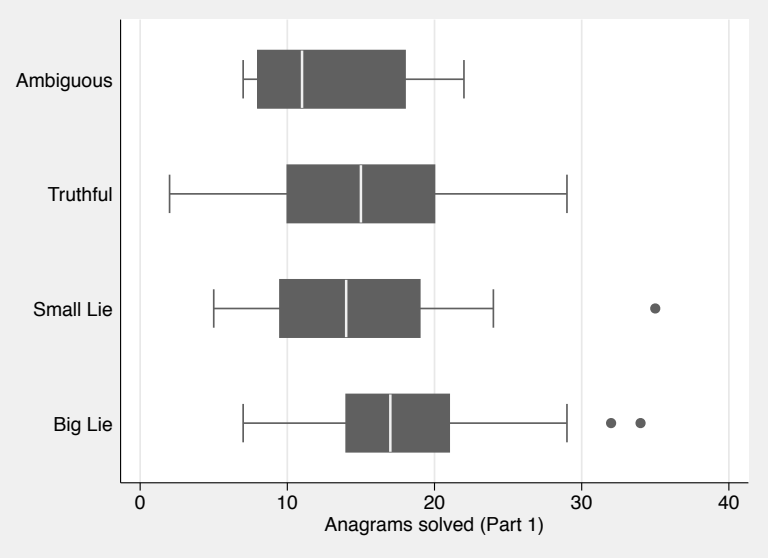

Figure 6: Scores in Bonus Conditions

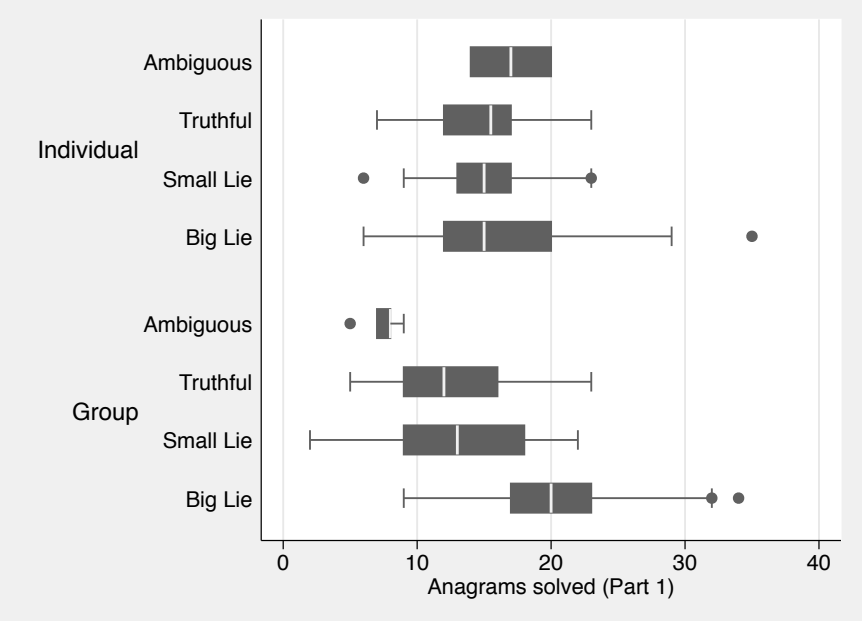

(17) than candidates who send truthful messages (15), and candidates who send ambiguous messages have a lower median score (11). The top part of Figure 6 shows that median scores do not vary much by message type in the Individual Bonus condition, suggesting a weaker relationship than in the baseline, which is not particularly surprising given that most subjects choose the big lie. In contrast, the bottom part of Figure 6 suggests an intriguing finding with respect to the relationship between scores and messages in the Group Bonus condition. The relationship between task performance and message choice appears to be much stronger when the bonus is given not just to the winning candidate, but also to every member of the candidate's group: The median score of candidates choosing the big lie is 20, which far exceeds the median score conditional on any other type of message (13 given a small lie, 12 given a truthful message, and 8 given an ambiguous message). The Group Bonus appears to suppress lying behavior by weaker candidates, perhaps lending credence to our interpretation that it increases the salience of lying aversion or norms of honesty.

To assess these relationships more systematically, we estimate a series of linear probability regression models. For each model, we dichotomize the dependent variable in a different way while including the same set of covariates: dummy variables for the bonus treatment (i.e., Part 2, regardless of whether the bonus is for the winning candidate or the entire winning group), the group bonus (to distinguish between the individual and group bonus effects), Part 1 score, and their relative beliefs about their scores. We measured relative beliefs (that is, confidence) with a questionnaire that followed the incentivized part of the experiment, asking each subject the degree to which they believed their score was higher than other subjects' scores. This measure takes on integer values from 0 (higher than no one else's score) to 4 (higher than everyone else's score).[14] Given that this measure of beliefs is not incentivized, we caution that we are not particularly confident that it has a high degree of measurement reliability. Nevertheless, we include it as control variable.[15]

[14]The intermediate values are "higher than about 1/4 of them" (1), "higher than 1/2 of them" (2), and "higher than 3/4 of them" (3).

[15]In a later section, we report additional results with incentivized beliefs.

Table 1: Linear probability model analysis of message behavior

	(1) Ambiguous b/se	(2) Truth b/se	(3) Small Lie b/se	(4) Big Lie b/se	(5) Non-Truth b/se	(6) Any Lie b/se
Bonus Treatments	-0.017	-0.233**	-0.017	0.267**	0.017	0.250**
	(0.03)	(0.06)	(0.05)	(0.06)	(0.03)	(0.06)
Group Bonus	-0.029	-0.014	0.076	-0.034	0.029	0.042
	(0.03)	(0.06)	(0.05)	(0.06)	(0.03)	(0.06)
Part 1 Score	-0.004	-0.003	-0.018**	0.026**	0.004	0.008
	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)
Relative Beliefs	0.029	0.031	-0.100**	0.040	-0.029	-0.060
	(0.02)	(0.04)	(0.03)	(0.04)	(0.02)	(0.04)
Constant	0.142	0.473**	0.587**	-0.202	0.858**	0.384*
	(0.11)	(0.17)	(0.16)	(0.14)	(0.11)	(0.18)
R^2	0.042	0.071	0.048	0.163	0.043	0.107
Observations	240	240	240	240	240	240

* $p < .05$, ** $p < .01$; robust standard errors clustered by subject

The first four columns of Table 1 present the results when we dichotomize the dependent variable for each type of message (whether the message was ambiguous, truthful, etc).[16] These results show that electoral incentives clearly decrease the probability of choosing truthful messages (column 2) while increasing the probability of choosing the big lie (column 4), providing further support for Hypothesis 3. In contrast, electoral incentives do not affect the prevalence of ambiguous messages (columm 1) or small lies (column 3). In column 5, the dependent variable is any non-truthful message (anything but the truth), while in column 6 it is any amount of exaggeration (pooling the small lies and big lies, excluding ambiguity); the difference between the last two columns suggests that incentives work by increasing the prevalence of exaggeration rather than ambiguity. No matter how we code the dependent variable, the coefficient for group bonus is not significant in any of the models. We can infer from this that there is no difference between the group and individual

[16]We obtain similar results when we estimate a multinomial probit model.

bonus in terms of message behavior. That is, the individual bonus for winning the election seems to be sufficient to encourage candidates to lie, while adding the group bonus neither significantly increases nor decreases such behavior.

We also find that lying behavior depends partly on ability and beliefs. Part 1 scores have a positive and statistically significant effect on the probability of big lies and a negative and statistically significant effect on the probability of small lies. These results mean that higher ability candidates are more likely to tell the big lie and less likely to tell the small lie, consistent with Hypothesis 2. The coefficient estimate for relative beliefs is negative and statistically significant only for small lies, which means that candidates who are more confident in their scores are less likely to use the small lie (which would be less effective). The coefficient is positive but not significant for big lies. This exaggeration by high ability candidates can be consistent with instrumental rationality if they believe low ability candidates might lie, thereby countering the distortions from low ability candidates and increasing the informativeness of the overall campaign environment. We cannot say definitively whether such lying behavior is consistent with instrumental rationality, however, because we do not have direct measures of candidates' beliefs about the behavior of their opponents.

To investigate whether the relationship between ability and confidence differs across electoral incentives, we report a series of separate ordered probits in Table 2. In this analysis, we can think of the type of message as being ordered in terms of the degree of precision and exaggeration. Since ambiguous messages are the least precise, we treat them as the lowest category, and then we order the other types of messages in terms of the degree of exaggeration.[17] Consistent with Figures 5 and 6, the results suggest that precision and exaggeration are weakly related to task scores in the Performance Pay condition, unrelated in the Individual Bonus condition, and strongly related in the Group Bonus condition. Thus, although the group bonus does not have an overall direct effect on campaign messages, we find that it does have an interactive effect with task ability.

[17]The results do not change appreciably if we omit ambiguous messages from the analysis and treat the dependent variable as being ordered in terms of the level of exaggeration.

Table 2: Ordered probit analysis of message behavior

	Performance		Individual		Group	
	b/se	b/se	b/se	b/se	b/se	b/se
Part 1 Score	0.03*	0.03	0.02	0.01	0.14**	0.11**
	(0.02)	(0.02)	(0.03)	(0.04)	(0.03)	(0.03)
Relative Beliefs		-0.02		-0.07		-0.25
		(0.14)		(0.20)		(0.21)
μ_1	-0.88**	-0.93	-1.59**	-1.84*	0.36	-0.47
	(0.28)	(0.55)	(0.53)	(0.94)	(0.40)	(0.81)
μ_2	0.64*	0.59	-0.60	-0.84	1.64**	0.86
	(0.27)	(0.54)	(0.46)	(0.88)	(0.42)	(0.77)
μ_3	1.31**	1.25*	0.03	-0.21	2.38**	1.61*
	(0.28)	(0.54)	(0.46)	(0.88)	(0.46)	(0.79)
Observations	120	120	60	60	60	60
Log likelihood	-146.4	-146.4	-63.31	-63.25	-62.46	-61.75

* $p < .05$, ** $p < .01$

Our experimental analysis shows that the degree of honesty in campaigns varies considerably across electoral contexts. Office motivations in the form of private benefits of winning office clearly induce candidates to exaggerate their qualifications. Strikingly, we also find that in an environment in which winning is not inherently beneficial (because candidates and voters have common incentives to elect the best representative), many candidates campaign honestly, but many nevertheless exaggerate. The relationship between messages and candidate quality (task ability) suggests that such exaggeration could be instrumental, perhaps reflecting a strategic response to distrust of others in political settings rather than an outright intention to deceive.

Beliefs and Motives

Cleanly differentiating between instrumental versus corrupt motives requires additional data on beliefs and motives. By beliefs, we mean probability beliefs in the standard sense used in decision theory and game theory. Specifically, to determine whether exaggeration is instrumental (a best response) or not, we need to know how likely a candidate believed that others exaggerated or lied. We therefore designed a set of incentivized belief elicitation procedures that we added to the end of the Individual Bonus treatment, and we were careful to elicit beliefs after subjects chose their campaign messages (to avoid influencing their choices) but before the elections or any of the results were revealed.[18] We refer to the modified treatment as the Belief Treatment, and we ran 10 sessions of this treatment (100 subjects). Incentivizing beliefs is important because it forces subjects to think about their beliefs in a way that is consequential and ensures that subjects understand beliefs in the same way, thereby enhancing the reliability of our belief measures.

[18]For full details of our procedure, see the text of the instructions in the Appendix. We ensured comprehension of the procedure with two comprehension checks; in addition, after the first comprehension check, we provided an additional explanation of the procedure. Because the belief elicitation requires a lengthy instruction period, we only elicited beliefs about the first election. Note that timing of the elicitation procedure means that the campaign message decisions are *identical* in the Belief Treatment and the Individual Bonus Treatment.

Our belief elictation procedure is based on probability matching (Healy 2011, Trautmann and van de Kuilen 2014). For each belief that we elicit, subjects face a series of binary choices between two gambles: one based on their subjective belief and one based on an objective probability. To illustrate, consider the belief held by candidate i that candidate j chose the big lie (i.e., that j exaggerated their score by 6). The subjective belief is represented by Option A, which pays a subject \$15 if candidate j actually chose the big lie and pays \$0 otherwise. The objective probability is represented by Option B, which is a binary lottery with a known probability of winning. We describe this probability to subjects in terms of a "virtual bingo cage" with red and white balls, where a subject wins \$15 if we draw a virtual red ball and wins \$0 if we draw a virtual white ball. Subjects then face a series of choices between Option A and Option B where we vary the probability that Option B wins, from 0% to 100% in increments of 5 percentage points. The probability belief that j chose the big lie is the objective probability for which subject i switches from preferring to be paid for Option A to preferring to be paid for Option B—that is, the proability that makes them *indifferent* between their subjective probability and the objective probability.

We used this procedure to elicit beliefs about relative ability and about campaign messages. Specifically, we elicited the following beliefs each subject held about his or her election opponent in the Performance Pay condition (the first set of elections):

$$\begin{aligned}
\text{Higher} &= \Pr(i\text{'s score higher than } j\text{'s score}) \\
\text{3+ Higher} &= \Pr(i\text{'s score more than 3 higher than } j\text{'s score}) \\
\text{6+ Higher} &= \Pr(i\text{'s score more than 6 higher than } j\text{'s score}) \\
\text{Truthful} &= \Pr(j \text{ chose truthful message}) \\
\text{Small lie} &= \Pr(j \text{ chose small lie}) \\
\text{Big lie} &= \Pr(j \text{ chose big lie})
\end{aligned}$$

Figure 7: Beliefs about others' scores

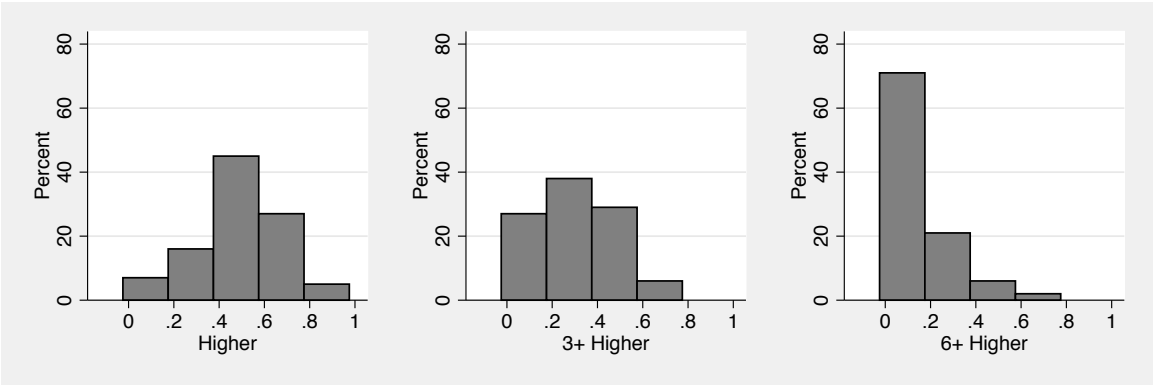

Figure 7 presents histograms of elicited beliefs about relative ability. Although we are not particularly interested in the accuracy of beliefs, it appears that in the aggregate, subjects hold reasonable beliefs about their relative abilities: 40% believe themselves to be below the median, 18% believe themselves to be at at the median, and 42% believe themselves to be above the median (see the left part of Figure 7). Few subjects are confident that their scores are extremely high (more than 6 higher than others' scores): 71% of subjects indicate such a probability belief less than 20%, with the highest such belief being 65%. For subjects who are extremely confident that their ability is in the upper tail of the distribution, there would be little cost to exaggerating their campaign message (given their level of confidence, it would be unlikely that exaggerating would prevent a higher ability subject from being elected).

More relevant to understanding whether exaggeration is rational are the histograms presented in Figure 8, which shows subjects' beliefs about the types of messages chosen by opposing candidates. The histogram on the left of Figure 8 suggests that subjects are not especially confident about the honesty of others: Over half of subjects (62%) hold a probability belief of less than 50% that the opposing candidate told the truth. Similarly, the histogram on the right shows that the belief that other candidates told the big lie is not uncommon: 19% of subjects have a probability belief greater than 50% while over half of subjects (56%) hold a probability belief of greater than 20%. Indeed, the distribution

of beliefs about the big lie is not dissimilar from the distribution of beliefs about truthful messages.

To what extent do these beliefs affect candidates' messages? Table 3 reports linear probability model regressions similar to the ones reported previously, in which we dichotomize the dependent variable in several ways. The results suggest that beliefs indeed account for candidate exaggeration. More specifically, beliefs about the lying behavior of others matter most, rather than beliefs about relative ability. Column 2 shows that as the beliefs that other candidates chose the small lie or big lie increase, the probability of choosing a truthful message decreases. Column 3 shows that the probability of a small lie is positively and significantly related to the belief that others chose the small lie. Similarly, column 4 shows that the probability of a big lie is positively and significantly related to the belief that others told the big lie. Overall, these results provide a strong indication that the level of exaggeration is consistent with an instrumentally rational best response: Candidates are more likely to choose the type of message that they believe others selected.

We also investigate whether we could account for differences in lying behavior by considering heterogeneity in personality traits or attitudes that might predispose some subjects to engage in either honest or dishonest behavior. We therefore included several attitudinal measures and personality scales in the questionnaire we administered after the incentivized portion of the Belief Treatment. The most relevant measure is the Machiavellian Person-

Figure 8: Beliefs about others' messages

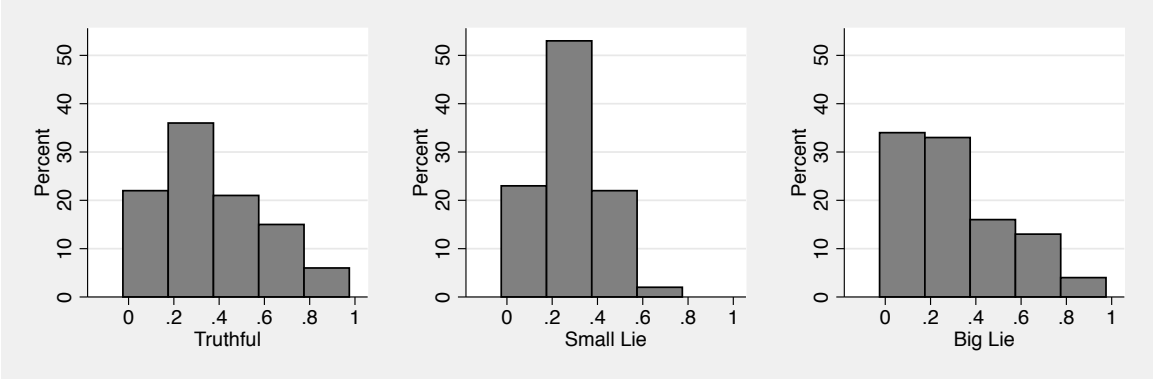

Table 3: Linear probability model analysis of beliefs and messages

	(1) Ambiguous b/se	(2) Truth b/se	(3) Small Lie b/se	(4) Big Lie b/se	(5) Non-Truth b/se	(6) Any Lie b/se
Part 1 Score	-0.005	-0.002	0.004	0.003	0.005	0.007
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Belief, Score 3+ higher	0.263	-0.063	0.074	-0.273	-0.263	-0.199
	(0.41)	(0.41)	(0.38)	(0.33)	(0.41)	(0.36)
Belief, Score 6+ higher	-0.439	-0.196	0.320	0.315	0.439	0.635
	(0.38)	(0.46)	(0.47)	(0.36)	(0.38)	(0.44)
Belief, Other Small Lie	-0.021	-0.823**	0.860**	-0.016	0.021	0.844**
	(0.16)	(0.28)	(0.28)	(0.21)	(0.16)	(0.28)
Belief, Other Big Lie	-0.057	-1.220**	0.073	1.204**	0.057	1.277**
	(0.14)	(0.15)	(0.13)	(0.14)	(0.14)	(0.15)
Constant	0.162+	1.133**	-0.171	-0.124	0.838**	-0.295*
	(0.09)	(0.14)	(0.14)	(0.09)	(0.09)	(0.13)
R^2	0.0371	0.335	0.139	0.418	0.0371	0.405
Observations	100	100	100	100	100	100

+ $p < .10$, * $p < .05$, ** $p < .01$; robust stanard errors

ality Scale (MPS, Dahling, Whitaker and Levy 2008), which attempts to measure several aspects of manipulative personalities: one's propensity to distrust others (Distrust), engage in amoral manipulation (Amoral), seek control over others (Control), and seek status for oneself (Status). If these personality traits drive lying behavior, we would expect that subjects scoring higher in each dimension of Machiavellianism to be more likely to exaggerate their qualifications for office.

In addition to MPS, we used a short version of the Big Five personality factors (TIPI, Gosling, Rentfrow and Swann 2003), which measures Extraversion, Agreeableness, Conscientiousness, (Emotional) Stability, and Openness. We have no prior expectations regarding which of these factors would be related to greater deceptive behavior. The questionnaire also included a generalized measure of trust, measures of risk aversion (Kam 2012), and basic demographics (e.g., gender).

In Table 4, we report various specifications of ordered probit analyses using the additional data from the Belief Treatment, with the incentivized beliefs and questionnaire measures as covariates. The factor that clearly stands out is a subject's belief that others are likely to exaggerate greatly (the probability belief that the other candidate chose the big lie). This coefficient is large, positive, and statistically significant across specifications. Beliefs about relative ability are positive and significant at the .05 level only when personality factors are excluded from the analysis (column 1). None of the Machiavellianism factors are significantly correlated with lying behavior, which is striking because these are the factors most relevant to the propensity to lie and cheat. Gender, generalized trust, and risk aversion play no role in accounting for candidate exaggeration as well. Interestingly, we find that subjects who perceive themselves as more emotionally stable as measured by the TIPI scale are less likely to exaggerate. Although this factor is negatively and significantly related to exaggeration, it does not affect our central conclusions about the importance of beliefs. If anything, it strengthens the importance of beliefs in the analysis, as the coefficient for beliefs about big lie are higher when we include TIPI covariates.

Table 4: Ordered probit analysis of beliefs, traits, and message behavior

	(1) b/se	(2) b/se	(3) b/se	(4) b/se	(5) b/se	(6) b/se
Part 1 Score	0.04*	0.02	0.02	0.02	0.03	0.04
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Belief, Score Higher		-0.44	-0.41	-0.31	-0.25	-0.26
		(0.79)	(0.80)	(0.82)	(0.83)	(0.88)
Belief, Score 6+ higher		1.84	1.80	1.73	1.05	0.48
		(1.11)	(1.12)	(1.13)	(1.14)	(1.20)
Belief, Other Small Lie		0.82	0.87	0.62	0.93	1.06
		(0.80)	(0.81)	(0.83)	(0.83)	(0.88)
Belief, Other Big Lie		3.61**	3.69**	3.68**	4.28**	4.61**
		(0.63)	(0.65)	(0.65)	(0.70)	(0.74)
Female			0.04			-0.07
			(0.24)			(0.28)
Trust			0.36			0.53
			(0.56)			(0.66)
Risk Aversion			0.12			-0.26
			(0.79)			(1.13)
Amoral				0.36		1.15
				(0.84)		(0.96)
Control				-0.08		-1.24
				(0.67)		(0.92)
Status				0.73		0.81
				(0.57)		(0.64)
Distrust				-0.83		-0.40
				(0.85)		(0.95)
Extraversion					0.37	0.30
					(0.50)	(0.62)
Agreeableness					0.63	0.34
					(0.57)	(0.69)
Conscientiousness					0.87	1.37*
					(0.53)	(0.66)
Stability					-1.15*	-1.45*
					(0.52)	(0.58)
Openness					0.95	1.35
					(0.77)	(0.87)
μ_1	-0.82**	-0.06	0.19	0.10	1.50	1.96
	(0.28)	(0.44)	(0.56)	(0.67)	(0.87)	(1.35)
μ_2	0.64*	1.66**	1.91**	1.84**	3.38**	3.91**
	(0.27)	(0.46)	(0.59)	(0.69)	(0.92)	(1.40)
μ_3	1.23**	2.51**	2.76**	2.71**	4.30**	4.88**
	(0.29)	(0.50)	(0.61)	(0.71)	(0.95)	(1.42)
Observations	100	100	100	100	100	100
Log likelihood	-122.8	-101.9	-101.6	-100.6	-96.53	-93.71

* $p < .05$, ** $p < .01$

Conclusion

Elections are the cornerstone of representative democracy. But elections also include a critical information problem: Candidates know their own quality but voters do not. We use laboratory elections to investigate the circumstances under which candidates either honestly convey (or conversely, obfuscate) their credentials. Several results are worth emphasizing.

First, we find that the private benefit of winning the election induces candidates to inflate their qualifications. Incentives clearly matter in ways that we would expect. When candidates have a stronger motivation to win, they are more likely to lie. We caution that because our experimental setup is specifically tailored to address the question of candidate honesty (by severely constraining the message space), we cannot say whether exaggeration necessarily undermines the informational quality of elections under more general circumstances.[19] For example, our framework does not address how lying varies with the risk of discovery (Boudreau 2009).

Second, we find a sizeable proportion of candidates who exaggerate when the only goal of the election is to select the highest ability representative. Interestingly, this finding contrasts with the experimental literature on strategic communication and lying aversion (e.g., Gneezy 2005, Gneezy, Rockenbach and Serra-Garcia 2013) as well as the literature on pandering and elections (e.g., Woon 2012, 2014) which finds that informed players are often reluctant to exploit their informational advantage by lying. At first glance, this result might seem troubling because it suggests that lying may be prevalent in even the most benign political settings. Upon further inspection, however, we show that exaggeration is driven primarily by candidates who hold strong beliefs that other candidates exaggerate. These beliefs are specific to the electoral environment and are unrelated to generalized trust or distrust of others or to power-seeking proclivities. Candidates in our experiment seem to have a sufficient degree of strategic sophistication to understand that conveying information

[19]Indeed, because greater coordination leads to greater information transmission in our game, the high degree of exaggeration in the Individual Bonus condition implies that it has the greatest informational content. We make no claims whatsoever that this logic applies to elections in general.

to voters depends not just on what they say, but what they say in relation to the claims that other candidates make.

Notably, our results indicate that the particular strategic and competitive context of elections may make untruthfulness endemic to democratic elections. One need not be a dishonest sort of person to grasp that lying in elections need not benefit only oneself, but provided that one is highly qualified, exaggerating may be in the public's best interest. The alternative to lying may be worse for the public at large if only the most competent candidates are honest while the least competent are deceitful. Noted American humorist Will Rogers famously said "If you ever injected truth into politics, you'd have no politics." Our results support Rogers' conclusion, but we would add that the fault may not lie solely with politicians, but with politics itself.

References

Alt, James, Ethan Bueno de Mesquita and Shanna Rose. 2011. "Disentangling Accountability and Competence in Elections: Evidence from U.S. Term Limits." *Journal of Politics* 73(01):171–186.

Armstrong-Taylor, Paul. 2012. "When Do Politicians Lie?" *The Berkeley Electronic Journal of Economic Analysis and Policy* 12(3):1–26.

Banks, Jeffrey S. 1990. "A model of electoral competition with incomplete information." *Journal of Economic Theory* 50(2):309–325.

Besley, Timothy. 2005. "Political Selection." *Journal of Economic Perspectives* 19:43–60.

Besley, Timothy, Jose G Montalvo and Marta Reynal-Querol. 2011. "Do Educated Leaders Matter?" *Economic Journal* 121(554):F205–227.

Besley, Timothy and Marta Reynal-Querol. 2011. "Do Democracies Select More Educated Leaders?" *American Political Science Review* 105(03):552–566.

Boudreau, Cheryl. 2009. "Closing the gap: When do cues eliminate differences between sophisticated and unsophisticated citizens?" *The Journal of Politics* 71(03):964–976.

Bucciol, Alessandro and Luca Zarri. 2013. "Lying in Politics: Evidence from the US." University of Verona Department of Economics Working Paper Series.

Buttice, Matthew K and Walter J Stone. 2012. "Candidates Matter: Policy and Quality Differences in Congressional Elections." *Journal of Politics* 74(03):870–887.

Cai, Hongbin and Joseph Tao-Yi Wang. 2006. "Overcommuniation in Strategic Information Transmission Games." *Games and Economic Behavior* 56:7–36.

Callander, Steven and Simon Wilkie. 2007. "Lies, Damned Lies, and Political Campaigns." *Games and Economic Behavior* 60(2):262–286.

Caselli, Francesco and Massimo Morelli. 2004. "Bad politicians." *Journal of Public Economics* 88(3-4):759–782.

Corazzini, Luca, Sebastian Kube, Michel Andre Marechal and Antonio Nicolo. 2014. "Elections and Deceptions: An Experimental Study on the Behavioral Effects of Democracy." *American Journal of Political Science* 58:579–592.

Crawford, Vincent and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50:1431–1451.

Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intention." *American Economic Review* 93(1):133–149.

Dahling, Jason J, Brian G Whitaker and Paul E Levy. 2008. "The development and validation of a new Machiavellianism scale." *Journal of Management* 35(2):219–257.

Davis, Michael L. and Michael Ferrantino. 1996. "Towards a Positive Theory of Political Rhetoric: Why Do Politicians Lie?" *Public Choice* 88(1):1–13.

De Paola, Maria and Vincenzo Scoppa. 2010. "Political competition and politician quality: evidence from Italian municipalities." *Public Choice* 148(3-4):547–559.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-Made Economics Experiments." *Experimental Economics* 10:171–8.

Gagliarducci, Stefano, Tommaso Nannicini and Paolo Naticchioni. 2010. "Moonlighting politicians." *Journal of Public Economics* 94(910):688 – 699.

Galasso, Vincenzo and Tomasso Nannicini. 2011. "Competing on Good Politicians." *American Political Science Review* 105(01):79–99.

Gneezy, Uri. 2005. "Deception: The Role of Consequences." *American Economic Review* 95(1):384–94.

Gneezy, Uri, Bettina Rockenbach and Marta Serra-Garcia. 2013. "Measuring Lying Aversion." *Journal of Economic Behavior and Organization* 93:293–300.

Gosling, Samuel D, Peter J Rentfrow and William B Swann. 2003. "A very brief measure of the Big-Five personality domains." *Journal of Research in Personality* 37(6):504–528.

Grimmer, Justin. 2013. *Representational Style in Congress: What Legislators Say and Why It Matters*. New York: Cambridge University Press.

Grimmer, Justin, Sean J. Westwood and Solomon Messing. 2015. *The Impression of Influence: Legislator Communication, Representation, and Democratic Accountability*. Princeton, NJ: Princeton University Press.

Healy, Paul J. 2011. "Epistemic Foundations for the Failure of Nash Equilibrium." Unpublished manuscript.

Hirano, Shigeo and James M Snyder, Jr. 2014. "Primary elections and the quality of elected officials." *Quarterly Journal of Political Science* 9(4):473–500.

Huang, Haifeng. 2010. "Electoral Competition When Some Candidates Lie and Others Pander." *Journal of Theoretical Politics* 22(3):333–358.

Hurkens, Sjaak and Navin Kartik. 2009. "Would I Lie to You? On social preferences and lying aversion." *Experimental Economics* 12:180–92.

Jones, Benjamin F. and Benjamin A. Olken. 2005. "Do leaders matter? National leadership and growth since World War II." *Quarterly Journal of Economics* 120:835–864.

Kam, Cindy D. 2012. "Risk Attitudes and Political Participation." *American Journal of Political Science* 56(4):817–836.

Kanthak, Kristin and Jonathan Woon. 2015. "Women Don't Run? Election Aversion and Candidate Entry." *American Journal of Political Science* .

Kartik, Navin and R. Preston McAfee. 2007. "Signaling character in electoral competition." *American Economic Review* 97(3):852–870.

Kartik, Navin and Richard Van Weelden. 2014. "Informative Cheap Talk in Elections." Unpublished mansucript.

Landa, Dimitri and Dominik Duell. 2014. "Social Identity and Electoral Accountability." *American Journal of Political Science* .

Lundquist, Tobias, Tore Ellingsen, Erik Gribbe and Magnus Johannesson. 2009. "The Aversion to Lying." *Journal of Economic Behavior and Organization* 70(1):81–92.

McCurley, Carl and Jeffery J. Mondak. 1995. "Inspected by# 1184063113: The influence of incumbents' competence and integrity in US House elections." *American Journal of Political Science* 39(4):864–885.

Messner, Matthias and Mattias K. Polborn. 2004. "Paying politicians." *Journal of Public Economics* 88(12):2423 – 2445.

Minozzi, William and Jonathan Woon. 2013. "Lying aversion, lobbying, and context in a strategic communication experiment." *Journal of Theoretical Politics* 25(3):309–337.

Mondak, Jeffery J. 1995. "Competence, integrity, and the electoral success of congressional incumbents." *The Journal of Politics* 57(4):1043–1069.

Morton, Rebecca B. and Kenneth C. Williams. 2010. *Experimental Political Science and the Study of Causality: From Nature to the Lab*. Cambridge University Press.

Pfiffner, James P. 1999. "The Contemporary Presidency: Presidential Lies." *Presidential Studies Quarterly* 29(4):903–917.

Ringquist, Evan J. and Carl Dasse. 2004. "Lies, Damned Lies, and Campaign Promises? Environmental Legislation in the 105th Congress." *Social Science Quarterly* 85(2):400–419.

Sanchez-Pages, Santiago and Marc Vorsatz. 2007. "An Experimental Study of Truth-Telling in a Sender-Receiver Game." *Games and Economic Behavior* 61(1):86–112.

Stone, Walter J, L. Sandy Maisel and Cherie D. Maestas. 2004. "Quality counts: Extending the strategic politician model of incumbent deterrence." *American Journal of Political Science* 48(3):479–495.

Tomz, Michael and Robert P. Van Houweling. 2009. "The Electoral Implications of Candidate Ambiguity." *American Political Science Review* 103(1):83–98.

Trautmann, Stefan T and Gijs van de Kuilen. 2014. "Belief Elicitation: A Horse Race among Truth Serums." *The Economic Journal* .

Wang, Joseph Tao-yi, Michael Spezio and Colin F. Camerer. 2010. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games." *American Economic Review* 100(3):984–1007.

Woon, Jonathan. 2012. "Democratic Accountability and Retrospective Voting: A Laboratory Experiment." *American Journal of Political Science* 56(4):913–930.

Woon, Jonathan. 2014. "An Experimental Study of Electoral Incentives and Institutional Choice." *Journal of Experimental Political Science* 1(2):181–200.