Campaign messages, polling, and elections: theory and experimental evidence

Abstract

We analyse and test a model of politicians' behaviour and voting. Politicians choose how much of an economy's resources to keep for themselves and how much to allocate to the citizenry, then face re–election against a challenger. Both incumbents and challengers have private information about their own quality which determines the economy's level of resources. We vary (i) whether communication from candidates to voters (campaigning) is possible, and (ii) whether candidates' quality heterogeneity is high or low. We observe that both higher quality heterogeneity and allowing campaigning benefit citizens at the expense of officials. We also find that campaigns have an impact on elections and that challengers' negative campaigning (criticising the incumbent) increases, and incumbents' positive campaigning (emphasising their own strengths) decreases, when incumbents performed poorly and when quality heterogeneity is high.

1 Introduction

In elections, voters rely on information to choose candidates. Two important sources of this information are candidates' track records and their election campaigning. The interaction between these is important: the incumbent's track record influences her campaigning and that of the challenger, while campaign promises may constrain the winner's subsequent policies. In this paper, we introduce a theoretical framework to investigate these factors. In each period, an incumbent official has access to an endowment of resources that depends on her intrinsic quality. The official chooses how much to keep for herself, with the remainder divided amongst the other citizens. After this decision is announced, one of the citizens is nominated to challenge the incumbent in an election. Candidates' qualities are their own private information, meaning that after a disappointing outcome, voters need to determine whether it was due to low official quality or because the official took too large a share.

We consider four versions of this setting, differing along two dimensions. The first concerns *quality heterogeneity*. Candidates can have high or low quality, but the difference between them is small in one treatment and large in the other (though average quality is fixed). We aim to compare a setting where evaluating quality is of secondary importance to one where it is crucial. When quality variance is low, the incumbent's track record on its own is a good indicator of her future behaviour. When it is high, the signal extraction problem becomes more important.

The second dimension involves *campaigning*. Candidates either can or cannot send messages to the other voters, before the election but after results from a non–binding "straw poll" are known. By comparing outcomes across this dimension, we can identify the effects of campaigning.

We test our theoretical findings using a laboratory experiment. The lab is the most appropriate test–bed for several reasons. First, researchers can induce individuals' preferences using monetary rewards, instead of having to infer or assume preferences as in observational studies and field experiments. Second, data collection is more complete (e.g., we record every campaign message, and each individual's votes). Finally, and most importantly, the lab allows control over features of the environment (such as the options available to decision makers, and the information they receive)

that would be impossible or unethical to vary in the field. Combined with random assignment to treatments, this removes or minimises many issues of endogeneity and selection.

Our study contributes to the literature in three important ways. First, our model and experiment allow us to consider the whole process of policy–making and elections – with and without the possibility of campaigning – so that, differently from observational studies or field experiments, we can study the effect the possibility of campaigning has on the entire process. We find that when campaigning is allowed, citizens benefit because candidates bid for voters' attention. Without campaigning, challengers cannot bring this additional pressure to bear.

We also study whether voters' uncertainty about the candidates affects their ability to hold politicians accountable. Theoretical and empirical work suggests that accountability improves when voters are better able to judge politicians' performance (see Ashworth 2012 for a survey). We find that voters benefit when quality variance is high because voters find it easier to screen high–quality from low–quality candidates via their performance. Interestingly, we find that campaigns can achieve the same effect because voters can screen out low–quality candidates who cannot compete with the promises of high–quality candidates.

Finally, our analysis allows us to examine certain features of the candidates' campaigns. In particular, we find: (i) effort spent campaigning is greater when the "straw poll" indicates a competitive election (Erikson and Palfrey 1998); (ii) only challengers tend to resort to negative campaigning, and more so the worse the incumbent's performance; (iii) incumbents tend to campaign more – and more positively – the better their performance was; and (iv) positive campaigning benefits candidates while negative campaigning is less effective (Lau and Pomper 2002).

2 Literature review

Our study speaks to several strands of the literature. First, our voters can vote retrospectively based on observed performance (accountability), and prospectively (selection), since candidates may differ in quality. Empirical tests of accountability versus selection often use voter sophistication as

a proxy for prospective voting, with the rationale that only future performance is relevant (Clarke and Stewart 1994; Erikson et al. 2000; MacKuen et al. 1992). Laboratory experiments have found that retrospective voting is more important than predicted (Azfar and Nelson 2007; Konrad and Sherif 2019; Landa 2010; Woon 2012). Ashworth et al. (2017) argue that accountability and selection reinforce each other when candidates' quality and actions are complements, as in our setting. Our results support this conclusion, but we also show that campaigning may provide an alternative mechanism, since voters can evaluate candidates based not only on whether they have fulfilled their past promises but also on their promises for the future.[1]

Second, our analysis is relevant to the literature on political campaigning.[2] This literature argues that campaigns benefit voters because they help them understand the fundamental issues (Arceneaux 2006; Vavreck 2009), while we show a direct effect on voters' welfare. With respect to the determinants of effort spent on campaigning, our use of pre–election polling allows candidates to assess where the election stands. Erikson and Palfrey (2000) find that campaign spending is higher in close elections; this is confirmed in our setting, with more effort spent on campaigning when polling indicates a close election.

Finally, a related issue is that of negative versus positive campaigning. Theory suggests that negative campaigning is less likely from front–runners (Harrington and Hess 1996; Skaperdas and Grofman 1995) and more likely when candidates possess negative information about their opponent (Li and Li 2013; Polborn and Yi 2006) and when "micro–targeting" particular groups of voters is possible (Schipper and Woo 2017). Mattes and Redlawsk (2014) argue that negative campaigning allows voters to learn about the candidates. Sigelman and Shiraev (2002) find that in Russian presidential elections, challengers are more likely to campaign negatively than incumbents. Walter et al. (2014) study the issue in a European context and conclude that negative campaigning is more likely from challengers, but not when elections are close or from candidates who are behind

[1]Corazzini et al. (2014) show that promises can be credible even in a static setting. Feltovich and Giovannoni (2015) consider the impact of campaign promises, but since their candidates have homogeneous qualities, they cannot account for selection.

[2]Stratmann (2005) and Jacobson (2015) provide two important surveys.

in the polls. Lau and Rovner (2009) survey the literature and find relatively little evidence of the effectiveness of negative campaigning. We find that negative campaigning is done mainly by challengers and that incumbent performance does matter.[3] We do not find that negative campaigning is more effective than positive campaigning, but voting in our setting is compulsory, whereas much of the empirical literature argues that an important rationale for negative campaigning is to reduce turnout by the opponent's supporters (e.g. Ansolabehere et al. 1994; Ansolabehere and Iyengar 1995; Finkel and Geer 1998; Freedman and Goldstein 1999; Wattenberg and Brians 1999).

3 The model and hypotheses

There are infinitely many discrete rounds, where an official and $n-1$ ordinary citizens interact, with $n \geq 5$. The official is chosen randomly in the first round, and by elections thenceforth. At the beginning of round t, the economy receives an endowment $n\theta_t$, where $\theta_t \in \left\{\theta^l, \theta^h\right\}$ represents the official's *quality* (with $\theta^h > \theta^l > 0$), which is her private information (citizens only know that both types are equally likely) and remains constant throughout her tenure as official. The official chooses how much *income* $\pi_t^c = x_t \in \left[0, \frac{n}{n-1}\theta_t\right]$ to give to each citizen (equally) from the endowment, keeping the remainder $\pi_t^o = n\theta_t - (n-1)\, x_t$ as her *salary*.

After citizens are informed of their incomes (but not the endowment or official's salary), a challenger is nominated, equally likely to be any of the $n-1$ citizens, and with quality equally likely to be θ^h or θ^l.[4] The challenger's quality is private information, and no other information about the challenger is disseminated to the other voters. It is costless to run for office.

Next, a pre–election *straw poll* is run. Everyone, including the candidates, simultaneously and costlessly votes with no possibility of abstention, after which the results are announced. We include a straw poll because it allows us to identify the treatment effects of campaigning, by comparing how voting changes between the straw poll and the election when campaigning is possible versus when it is not. Our objective is not to examine whether the presence of polling changes behaviour;

[3] In our setting, there is no track record for challengers, so the former result is not surprising.

[4] Thus, our setting does not deal with the issue of candidate self–selection.

hence we do not vary whether there is a straw poll. Our implementation of polling abstracts in a few ways from reality. Everyone is polled, rather than a sample of the electorate. We do this for the identification purpose discussed above, but note that it is common in the experimental literature (e.g., Agranov et al. 2018). Also, unlike real–world polling, our polling is costless, removing the main reason for using sampling in real polling. Finally, we remove the possibility of abstention, in polling and in the election. This simplifies the setting, though we acknowledge that most elections involve voluntary rather than compulsory voting.[5]

In the game *without* campaigning, the election takes place immediately after the straw poll. Again, everyone simultaneously and costlessly votes. The candidate with more votes wins the election and will be the official in the next round (with quality unchanged from the current round), with ties broken in the incumbent's favour. The election loser goes back to being an ordinary citizen.

The game *with* campaigning allows both incumbent and challenger to simultaneously broadcast a natural–language message *after* the straw poll results are announced, but *before* the election. Messages incur no costs other than the effort required to compose them, and do not formally bind future decisions.

The stage game (shown in Figure 1) is infinitely repeated. All players have the same lifetime utility function:

$$U(\pi_1, \pi_2, ...) = \sum_{t=1}^{\infty} \delta^{t-1} u\left(\pi_t\right), \tag{1}$$

where $\delta \in [0, 1)$ is the common discount factor and $u(\cdot)$ is a common differentiable, strictly increasing, weakly concave utility function with $u(0) = 0$.[6]

[5]Australia is a notable exception. Bhattacharya et al. (2014) find that while compulsory and voluntary voting differ in how voters behave strategically (via insincere voting and abstention respectively), the associated electoral outcomes are very similar, suggesting that allowing abstention would not materially affect our results.

[6]Thus, preferences can be either risk neutral or risk averse. We do not include an explicit effort cost of sending campaign messages, but in the appendix we discuss campaigning when costs are positive but small (comparable to the effort needed to type messages in the experiment). Finally,

Figure 1: Sequence of events in a round

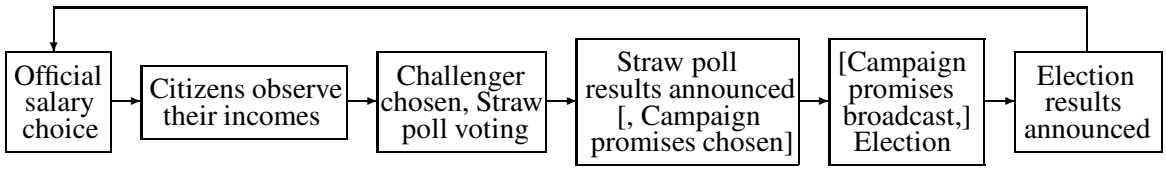

Both games are symmetric from a non–candidate citizen's standpoint. It is therefore natural to focus on symmetric equilibria where these citizens utilise the same voting strategy, though many asymmetric equilibria also exist. Following the literature, we further refine these into stationary pure–strategy equilibria. In the game without campaign promises, this, combined with our restriction that $n \geq 5$ (so the candidates are never pivotal), implies that voters will re–elect an incumbent if and only if their income is above a given threshold. In the game with campaign promises, re–election also depends on campaign promises; citizens consider not only their incomes, but also (i) whether the incumbent behaved in accordance with her previous–round campaign promises, and (ii) the incumbent's and challenger's (credible) current–round promises.[7]

We define three classes of such equilibria. In *pooling equilibria*, both high– and low–quality incumbents choose a salary that leads to re–election. In *separating equilibria*, only high–quality incumbents choose salaries that lead to re–election, while low–quality incumbents take a maximal salary and give up on re–election. In *kleptocratic equilibria*, both types forgo re–election and take a maximal salary.

---

our experiment implements the theoretical model in the standard fashion, with indefinite repetition (the game continues to the next round with probability δ, but without discounting) substituting for infinite repetition using the same value of δ as the discount factor. This technique is well–accepted in experimental research (see, e.g., Camera and Casari 2014; Cason and Mui 2014; Duffy and Puzzello 2014; Engle-Warnick and Slonim 2006), and preserves all of the theoretical results (Mailath and Samuelson 2006, p. 106).

[7]See the appendix for specifics. We focus on these equilibria because Feltovich and Giovannoni (2015) provide empirical support for them.

In the appendix, we show that *without* campaign promises, pooling equilibria obtain when citizens condition re–election on low levels of income, separating equilibria when they require intermediate incomes, and kleptocratic equilibria when they require high incomes. If we keep average quality constant but increase its variance, then for any risk–averse (or –neutral) preferences, the range of citizens' incomes for which separating equilibria should obtain increases, making these equilibria more likely. Under weak additional assumptions on preferences, we also show citizens receive higher incomes in separating equilibria.[8] By contrast, *with* campaign promises, if the cost of sending messages is positive but negligible, and if voters reward candidates who promise higher incomes and keep these promises, then only separating and kleptocratic equilibria exist (pooling equilibria cannot exist since since only high–quality types can credibly promise higher incomes).

Since citizens' incomes in a separating equilibrium increase as quality variance increases, and since separating equilibria yield higher citizen incomes than pooling equilibria, we have:

Hypothesis 1 *Citizens' incomes are higher when quality variance is higher.*

Hypothesis 2 *Citizens' incomes are higher when campaigning is possible.*

Similar to the reasoning behind Hypothesis 1, citizens will raise the threshold for re–election when quality variance is high, making it less likely that a given observed income will be above the bar:

Hypothesis 3 *Conditional on citizens' income, re–election is more likely in the low–variance treatment than in the high–variance treatment.*

When campaigning is impossible, separating equilibria become more likely as quality variance increases. Hence:

[8]This result is true for a wide range of preferences. In the appendix, we show that the result holds if we assume standard HARA (hyperbolic absolute risk aversion) preferences with decreasing absolute risk aversion and either constant (CRRA) or increasing (IRRA) relative risk aversion. Theory (e.g. Arrow 1971) and empirical work (Holt and Laury 2002; Meyer and Meyer 2005; Harrison et al. 2007) support the assumption that preferences are either CRRA or IRRA.

Hypothesis 4 *When campaigning is not possible, re–election will be more likely for high–quality officials, and less likely for low–quality officials, when quality variance is higher.*

(When campaigning is possible, pooling equilibria cannot obtain, but higher quality variance makes separating equilibria relatively more likely relative to kleptocratic equilibria, so an increase in quality variance should increase the possibility that high–quality incumbents are re–elected.)

The importance of the straw poll in our setting is that it allows us to determine the impact of campaign messages. Because the non–candidate citizens have identical preferences and constitute a majority, there always exists an equilibrium where they vote in the straw poll in the same way they would vote, absent messages, in the actual election.[9] Thus, if voting changes between the straw poll and the election when campaigning is possible, then we can attribute these differences to the campaign messages. Due to its non–zero effort cost, campaigning should be more prevalent when the straw poll indicates a competitive election.

Our analysis in the appendix shows that when messages have positive but vanishingly small costs, high–quality incumbents will campaign and low–quality incumbents will not, while challengers will campaign if promises are broken by the incumbent, and not otherwise. We classify messages according to the following (non–exclusive, non–exhaustive) categories:

- A message contains *positive* campaigning if it refers to the sender's behaviour and/or quality. It contains *negative* campaigning if it refers to the opponent's behaviour and/or quality.[10]

- A message is about *quality* if it refers to either candidate's quality. It is *retrospective* if it refers to either candidate's previous performance and *prospective* if it refers to either's (anticipated) future performance. In these last two cases, we also say the message is about *performance*.

[9]The straw poll is cheap talk, so other equilibria exist, including babbling equilibria where the result is ignored.

[10]Candidates could, if they wished, campaign positively about the opponent or negatively about themselves. Probably unsurprisingly, this nearly never happened in the experiment.

In our model, message content is cheap talk, and thus outside the scope of game–theoretic predictions. However, our focus on equilibria where candidates make promises and, if elected, are held accountable for them, suggests some natural message strategies. When citizens' incomes are low, promises are less likely to have been kept, and challengers will emphasise that through negative and retrospective messages. When incomes are high, challengers' messages will avoid mentioning the incumbent's performance and thus be of a positive and prospective kind. Also, when quality variance is higher, a given income is less likely to be satisfactory to citizens (recall Hypothesis 3), implying more retrospective and negative campaigning from challengers in the high–variance than in the low–variance treatment. Hence:

Hypothesis 5 *When campaign messages are possible, the probability of negative and retrospective campaigning by challengers is (i) decreasing in the income given to citizens, and (ii) higher in the high–variance treatment conditional on citizens' income.*

For incumbents, the reverse predictions apply:

Hypothesis 6 *When campaign messages are possible, the probability of positive and retrospective campaigning by incumbents is (i) increasing in the income given to citizens, and (ii) lower in the high–variance treatment conditional on citizens' income.*

4 Experimental procedures

The experimental sessions took place at [...] (see Table 1), and were programmed in z–Tree (Fischbacher, 2007). Subjects were primarily undergraduate students (no exclusion criteria). Subjects received written instructions, which were also read aloud.[11] Questions were answered individually by the experimenter.

[11]Sample instructions are in the appendix. Additional materials are available from the corresponding author.

Table 1: Experiment information

Cell	Campaign messages?	Quality heterogeneity	Sessions	Groups of size 5	Groups of size 6	Groups of size 7	Subjects
MH (Messages, High variance)	Yes	Low	5	7	1	6	83
ML (Messages, Low variance)	Yes	High	5	8	2	4	80
EH (Election only, High variance)	No	Low	4	6	1	4	64
EL (Election only, Low variance)	No	High	4	6	2	3	63

Subjects were randomly assigned to groups at the beginning of each supergame; these stayed fixed throughout the supergame. Subjects were visually isolated and asked not to communicate, and no identifying information was provided about other group members.

A round began with subjects' screens displaying their group's population and whether they were the official (see Figure 1). The official was shown the endowment and prompted to choose a salary. Citizens were informed of their income, but not the endowment, official's quality or salary. One citizen was privately informed of nomination as challenger and his/her quality.[12] Then the straw poll was held.

In the Messages cells, candidates could write messages ("campaign announcements") after the straw–poll vote counts and winner were announced. Messages had to be 0–140 characters, in English, and without identifying information. Messages were displayed to all group members, then the election took place (in the Election–only treatment, this happened after the straw poll

[12]Quality was randomly drawn for challengers, even those who had lost an earlier election. This is intended to keep our setting comparable to the existing literature (e.g., Persson et al. 1997, 2000), where infinite population means a zero probability of being re–nominated after an election loss. Our experiment necessarily uses finite populations, but re–drawing new candidates' qualities, and other design features like maintaining subject anonymity, let us simulate infinite populations.

results were announced). After the election winner and vote counts were announced, subjects were informed if the supergame would continue or end (so all decisions were made before knowing whether the supergame would continue).

After the sixth supergame, subjects completed a lottery task (Eckel and Grossman, 2008) and a demographic questionnaire. Then, subjects were privately paid the sum of payoffs (salary for officials, income for other citizens) from the last round of each supergame plus the lottery task, in Australian dollars (AUD), and a 5 AUD (≈ 3.75 USD) participation payment. Sessions typically lasted 60–120 minutes. Total earnings averaged just over 47 AUD, ranging from 23.05 to 112.05.

5 Experimental results

Many of the results reported below relate to hypotheses; we note when this is the case. (This is especially important in experiments – like ours – that were not pre–registered.) Other results do not relate to hypotheses, but are interesting in their own right. Because of the risk of false positives, our emphasis is on results connected to hypotheses, but we include the additional results for completeness.

5.1 Economic conditions and election outcomes

First, we observe (Table 2) substantial rents to being in office; on average, officials receive more than six dollars per round (over 20 percent of the endowment) more than individual citizens. This difference is confirmed by a Wilcoxon signed–ranks test (pooled MH, ML, EH and EL cells, $p < 0.001$).[13]

Second, low–quality officials take a significantly larger fraction of the economy's endowment

[13]Our non–parametric tests err conservatively by using session–level data and two–tailed rejection regions. When not stated explicitly, it should be understood that we continue to use the test mentioned most recently. See Siegel and Castellan (1988) for descriptions of these tests, and Feltovich (2005) for critical values of the robust rank–order test.

Table 2: Aggregate official salary choices

Cell	Official	Salary choice		Citizen income	
	quality	in dollars	frac. of endowment	in dollars	frac. of endowment
MH	8	12.19	0.273	7.06	0.156
	3	8.13	0.488	1.86	0.108
	pooled	11.06	0.333	5.61	0.143
ML	6	8.72	0.247	5.45	0.156
	5	9.64	0.332	4.07	0.145
	pooled	9.00	0.273	5.03	0.152
EH	8	13.78	0.303	6.78	0.151
	3	10.41	0.617	1.40	0.083
	pooled	12.88	0.386	5.35	0.133
EL	6	9.91	0.284	5.22	0.155
	5	12.14	0.416	3.56	0.128
	pooled	10.79	0.336	4.56	0.144

than high–quality officials in every cell (pooled MH, ML, EH and EL, $p < 0.001$). Since the endowment is smaller when the official is low–quality, the effect of officials' quality on salary choices in dollars depends on the quality variance (and hence the difference between low and high endowments): low–quality officials' dollar salaries are higher under low variance and lower under high variance ($p \approx 0.012$ and $p \approx 0.055$ for pooled MH/EH and pooled ML/EL respectively) compared to high–quality officials. The effect of officials' quality on citizens' incomes is unambiguous: higher (in dollars or fractions of the endowment) when the official is high–quality (pooled cells, $p < 0.001$ for both comparisons).

Third, both officials and citizens earn more under high variance, though the differences are insignificant (robust rank–order test, pooled E and M cells, $p \approx 0.080$ and $p \approx 0.067$ for salaries

and incomes respectively). Campaigning appears to benefit citizens at the expense of officials, consistent with Hypothesis 2, with lower salaries in the M treatment than the E treatment, though again these differences are insignificant (pooled H and L cells, $p \approx 0.079$ and $p \approx 0.059$ for dollar salaries and shares of the endowment respectively).

We further examine these results using regressions. As non–parametric tests err conservatively, it is necessary to consider those results in conjunction with corresponding regression results (which tend to err liberally) to get the full picture. Our dependent variables are the four from Table 2, but we focus on two: officials' salaries as fractions of the endowment (rather than dollars, to allow comparison across quality levels and treatments), and citizens' incomes in dollars (rather than fractions of the endowment) because this is what they observe and care about.

We estimate Tobit models, using the group–round as the unit of observation. Our main explanatory variables are dummies for the Messages treatment, the High–variance treatment, high official quality, and all interactions of these. We include a constant term, the supergame and round numbers (controlling for learning), dummies for group sizes of 5 and 7 (baseline = 6) and the official's post–experiment lottery choice (controlling for risk attitudes).

Table 3 shows the most important results. We focus on marginal effects – either averaged or conditional on particular variables' values – accompanied by standard errors (clustered by session).[14] The main results from Table 2 are confirmed, and all significant. Officials' salaries are decreased by allowing campaign messages and increased when quality variance is high, while high–quality officials take less of the endowment as salary. Citizens' incomes are higher when the official is high–quality, when quality variance is high, and when campaign messages are available. These last two differences provide support for our Hypotheses 1 and 2.

[14]Clustering by group instead of session, using linear models instead of probits, and leaving out the official's risk tolerance do not materially change the results. We omit some variables from the tables to save space, because they are insignificant or their effects are unimportant to us; however, we leave them in the regressions themselves. Details of all regressions mentioned in the paper, including those not reported in the main text, are available from the corresponding author.

Table 3: Tobit average marginal effects (standard errors in parentheses)

	[1]	[2]	[3]	[4]
	Official's salary		Citizen income	
	in $	as frac. of endowment	in $	as frac. of endowment
Messages (dummy)	-1.642^{**}	-0.051^{**}	0.402^{**}	0.012^{**}
	(0.731)	(0.022)	(0.189)	(0.006)
High variance (dummy)	2.186^{***}	0.068^{***}	0.331^{*}	-0.018^{***}
	(0.693)	(0.022)	(0.181)	(0.006)
High quality (dummy)	–0.282	-0.186^{***}	3.667^{***}	0.046^{***}
	(0.600)	(0.025)	(0.149)	(0.006)
Supergame	-0.488^{***}	-0.010^{**}	0.132^{***}	0.003^{**}
	(0.157)	(0.004)	(0.040)	(0.001)
Round	0.169^{*}	0.007^{**}	-0.047^{*}	-0.002^{**}
	(0.099)	(0.003)	(0.027)	(0.001)
Official's risk tolerance	0.514^{***}	0.015^{***}	-0.148^{***}	-0.005^{***}
	(0.144)	(0.005)	(0.040)	(0.001)
Group–size dummies?	Yes	Yes	Yes	Yes
N	1555	1555	1555	1555

* (**,***): *Marginal effect significantly different from zero at the 10% (5%, 1%) level.*

Result 1 *Citizens' incomes are higher in the high– than the low–variance treatment.*

Result 2 *Citizens' incomes are higher in the Messages treatment than the Election–only treatment.*

We next examine voting behaviour. Table 4 shows two measures of election outcomes: incumbents' vote share and re–election frequency. We observe:

Table 4: Aggregate election outcomes, by cell and official quality

	Official	Incumbent vote share			Re–election probability		
	quality	High var.	Low var.	Sig. diff.?	High var.	Low va.	Sig. diff.?
Message	High	0.718	0.632	$p \approx 0.044$	0.787	0.742	$p > 0.20$
treatment	Low	0.203	0.379	$p \approx 0.008$	0.103	0.281	$p \approx 0.206$
Election–only	High	0.614	0.557	$p \approx 0.114$	0.755	0.648	$p \approx 0.114$
treatment	Low	0.181	0.323	$p \approx 0.028$	0.031	0.255	$p \approx 0.028$
Pooled	High	0.669	0.597	$p \approx 0.025$	0.772	0.698	$p \approx 0.086$
	Low	0.193	0.348	$p < 0.001$	0.071	0.266	$p \approx 0.003$

Notes: p–values based on robust rank–order tests. "Var."=variance.

Result 3 *Low–quality officials fare worse in elections, and high–quality officials fare better, under high variance, though differences are not always significant.*

The differences in the Election–only treatment are consistent with Hypothesis 4, and as expected, the effect is similar in the Messages treatment. If we pool treatments, differences are significant in three of the four cases, and just miss significance in the fourth.

Figure 2 displays relationships between citizens' incomes and incumbents' electoral success in each cell. Both measures improve with citizens' incomes, and have roughly S–shaped curves: consistent with retrospective voting, and specifically threshold voting strategies (see the appendix). A positive correlation between incomes and incumbent election success is seen in every session of every cell, and in both the straw poll and the election; the probability of this occurring if there were actually zero or negative correlation is vanishingly small.

Table 5 shows associated probit regression results. The dependent variables are (a) an individual's straw–poll vote for the incumbent, (b) the incumbent winning the poll, (c) an individual's election vote for the incumbent, and (d) the incumbent winning the election. The unit of observation is the subject–round for (a) and (c) (dropping candidates from the sample), and the

Figure 2: Election outcomes by cell and citizen income

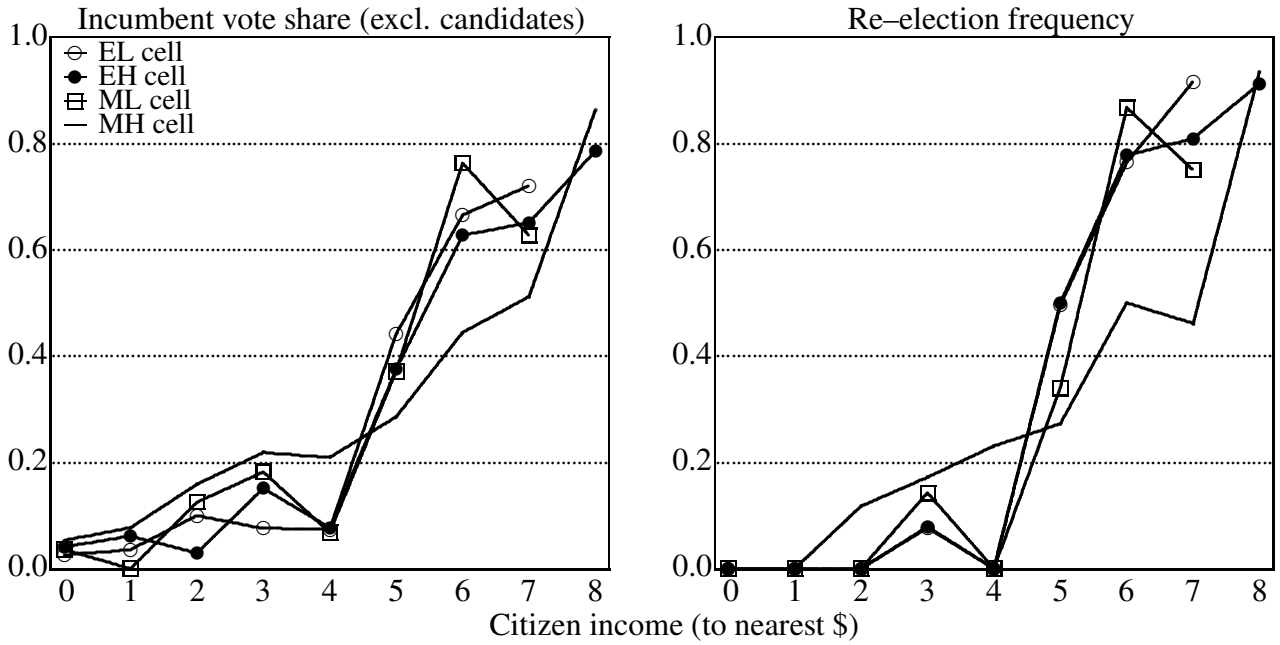

group–round for (b) and (d). Our main explanatory variables are the two treatment dummies used previously, the income of a representative citizen, and all interactions, along with the supergame and round numbers, group–size dummies, and a constant term.

We observe an increasing relationship between citizens' income and incumbents' elections success: not connected to our hypotheses, but clearly consistent with retrospective voting, previous experimental results, and Figure 2. We also observe that, for a given citizens' income, allowing campaigning favours incumbents as they receive significantly higher shares of both the straw vote and the election vote; however, this success yields only an insignificant boost in poll and election victories.[15] Finally, high variance favours the challenger, consistent with Hypothesis 3:

Result 4 *Incumbents' electoral outcomes are worse under high– than under low–variance, controlling for citizens' incomes.*

[15] Finding a significant effect in vote counts but not election wins is not paradoxical. Messages could potentially have a large effect in uncompetitive elections but a smaller effect in close ones. Alternatively, it could be a consequence of using subject–level data for individual votes and group–level data (and hence smaller sample sizes) for outcomes.

Table 5: Probit average marginal effects (standard errors in parentheses)

	[5] Straw vote for incumbent (excl. candidates)	[6] Incumbent wins straw poll	[7] Vote for incumbent (excl. candidates)	[8] Re–election
Messages (dummy)	0.056^{***}	0.024	0.057^{***}	0.019
	(0.017)	(0.023)	(0.018)	(0.024)
High variance (dummy)	-0.037^{**}	-0.055^{**}	-0.045^{***}	-0.051^{**}
	(0.014)	(0.024)	(0.015)	(0.021)
Citizen income ($)	0.111^{***}	0.151^{***}	0.118^{***}	0.162^{***}
	(0.004)	(0.010)	(0.006)	(0.012)
Supergame	0.014^{**}	0.001	0.015^{***}	0.017^{***}
	(0.006)	(0.008)	(0.005)	(0.006)
Round	0.008^{***}	0.009^{***}	0.008^{***}	0.007^{***}
	(0.003)	(0.002)	(0.002)	(0.002)
Group–size dummies?	Yes	Yes	Yes	Yes
N	5910	1555	5910	1555

** (***): *Marginal effect significantly different from zero at the 5% (1%) level.*

Table 6 shows re–election frequencies broken down by cell and incumbent–challenger quality pair. Voters appear successful at inferring incumbents' quality, with high–quality incumbents typically 50 percentage points more likely than low–quality incumbents to be re–elected. These differences are insignificant for individual cells due to the use of session–level data (Wilcoxon test, $p = 0.125$ for incumbents facing either high– or low–quality challengers in either EH or EL, $p = 0.0625$ for the same comparisons in either MH or ML), though we note that these represent the lowest two–sided p–values possible given the sample sizes of 4 (EH and EL cells) and 5 (MH

and ML).

Table 6: Re–election percents, by cell and candidate qualities

Candidate qualities	Cell			
(incumbent, challenger)	MH	ML	EH	EL
(H,L)	82.1	80.1	78.2	63.3
(H,H)	75.2	67.7	72.4	66.4
(L,L)	19.6	40.9	2.5	18.7
(L,H)	1.7	12.7	3.6	32.1

Table 6 highlights an important difference between increasing variance and allowing for campaign messages. In the former, voters are better able to evaluate the incumbent's performance and to discern their quality, but they still have no information about challengers. In the latter, voters may discern challengers' quality via their campaign messages. As evidence of this, voters are more likely to return incumbents facing low–quality challengers: the difference is small and insignificant when the incumbent herself is high–quality ($p > 0.20$ and $p = 0.125$ in the MH and ML cells, respectively), but is larger when the incumbent is low–quality (the minimum p–value of 0.0625 in either MH or ML). By contrast, voters in the Election–only cells are not more likely to re–elect when the challenger is low–quality.[16]

A final result comes from comparing elections between two high–quality candidates versus elections between two low–quality candidates. In both MH and ML cells, incumbents are more

[16]In our theoretical analysis, we argue for equilibria with credible communication, which imply that low–quality incumbents are replaced regardless of the opponent's quality. Table 6 suggests that this usually happens, but under low variance (ML cell), voters sometimes do re–elect low–quality incumbents who face a low–quality challenger. Importantly, however, the overwhelming majority of voters replace low–quality incumbents in favour of *high*–quality challengers when given the opportunity.

often re–elected in (H,H) elections than in (L,L) elections ($p = 0.0625$ for both comparisons).[17] Thus, while campaigning matters for electoral success, observed track records matter more.

5.2 Polling and election outcomes

Straw poll votes are cheap talk, with no strict incentive to reveal one's preference.[18] Also, voters may change their minds between the straw poll and the election. Nonetheless, straw polls forecast the election result reasonably well, with the straw–poll winner also winning the subsequent election 92.1 percent of the time in the Election–only treatment, and 88.7 percent of the time in the Messages treatment.

Figure 3 shows, for every election, a plotted point whose coordinates are the incumbent's straw–vote share and her election vote share, along with least–squares (OLS) trend lines for each cell, and 45–degree lines where the poll is exactly correct.[19] We define *polling bias* as the challenger's straw–poll vote share minus her election vote share, and *polling error* as the absolute value of the polling bias. Each trend line lies below the 45–degree line, indicating pro–incumbent polling bias that is small (2.7 and 2.1 percentage points in the Election–only and Messages treat-

[17]The earlier remark about 0.0625 being the minimum possible p–value holds here as well. Also, if we pool low– and high–variance cells, the results in Table 6 become stronger. For a given challenger's quality, comparing high– and low–quality incumbents yields p–values of 0.008 and 0.002 (Election–only and Messages treatments respectively). Comparing incumbents of a given quality facing low– versus high–quality challengers in the Messages treatment yields p–values of 0.027 and 0.002, while in the Election–only treatment the corresponding differences are still insignificant ($p \approx 0.16$ and $p > 0.20$ for low– and high–quality incumbents). Comparing (H,H) and (L,L) elections in the Messages treatment yields a p–value of 0.002.

[18]However, unlike other settings (e.g., Agranov et al. 2018), there is also no strict incentive to vote insincerely.

[19]For visibility, individual points were "jittered" by adding noise (uniform over [-0.05, 0.05]) to both coordinates. The trend line and statistical analysis are based on the original, un–jittered points.

ments respectively) but significant (Wilcoxon test, $p \approx 0.0078$ and $p \approx 0.0098$). It does not differ significantly by cell ($p > 0.20$ for all comparisons). By contrast, polling error is substantially and

Figure 3: Incumbent vote share, polling and elections (individual groups, jittered)

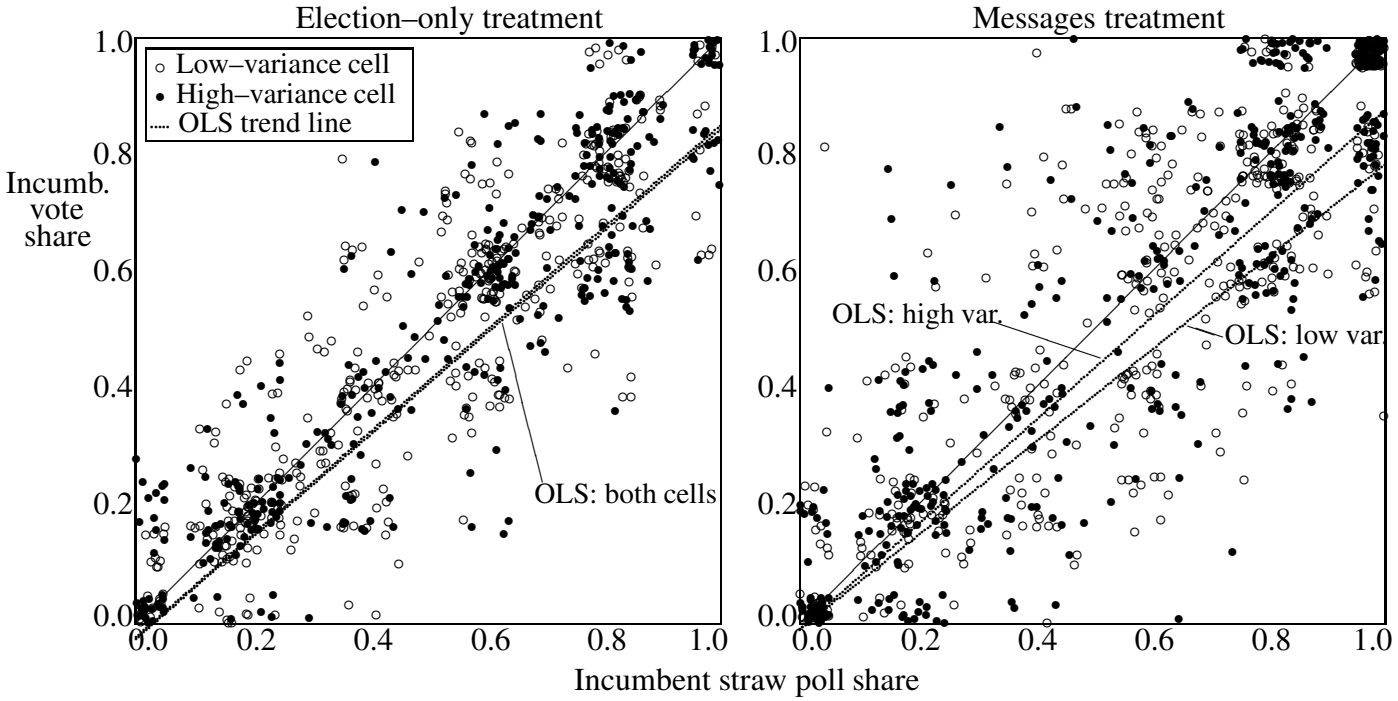

significantly higher in the Messages treatment: 11.6 percentage points, versus 7.8 in the Election–only treatment ($p \approx 0.049$ for MH versus EH or ML versus EL, $p \approx 0.004$ for pooled Messages versus pooled Election–only).

5.3 Campaign messages – quantitative analysis

Differences in polling error between the Message and Election–only treatments (previous section), and in challenger quality's effect on re–election rates (Table 6), suggest a role for campaigning: candidates convey information in messages that voters view as informative. We first examine the former. Figure 4 shows how message lengths and frequencies of blank (zero–character) messages depend on polling. Messages take effort to compose, with their length a proxy for the effort required. If the poll vote is close, both candidates have a chance of winning, so both should

expend more effort, sending more and longer messages.

Figure 4: Message effort

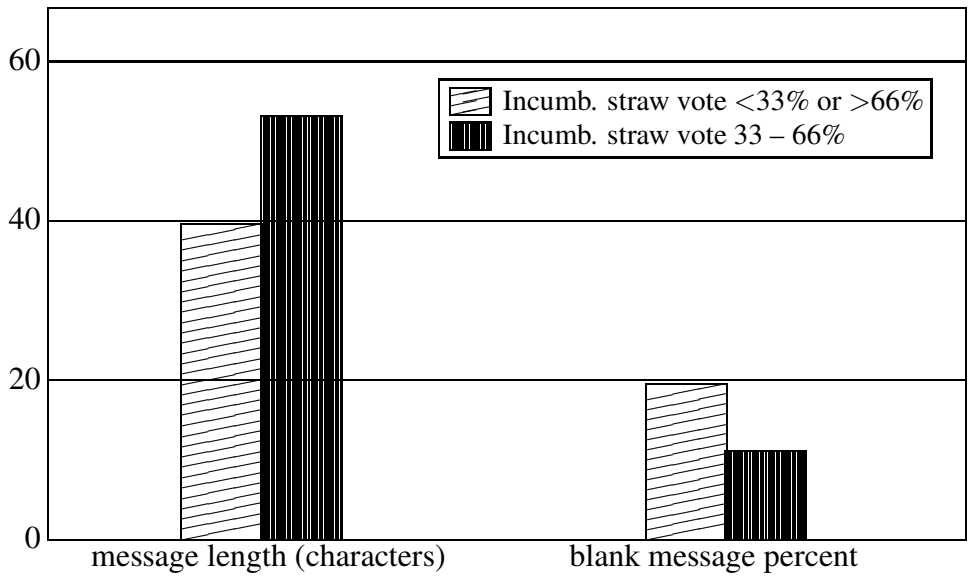

Candidates behave consistently with this logic. Pooling the two kinds of uncompetitive election (favouring incumbents and favouring challengers), we find messages are significantly shorter and more often blank than in close elections (Wilcoxon test, $p \approx 0.001$ for message lengths and 0.002 for blank messages). Also, challengers send longer messages than incumbents. This difference is significant whether we pool close and uncompetitive elections ($p \approx 0.014$ for length, $p \approx 0.008$ for blank message frequency) or examine them separately (p–values of 0.027–0.084 for the four comparisons).

Table 7 reports Tobit regression results with either polling bias or polling error as the dependent variable. In models 9 and 11, explanatory variables are the Message and High–variance dummies, citizen income, and all interactions, along with a constant term, the supergame and round numbers, and group–size dummies. Models 10 and 12 drop the Message dummy and its interactions, add the candidates' message lengths and their products with the High–variance dummy, and use only the Messages–treatment data. Consistent with our earlier discussion (e.g., Table 5 and Figure 3), allowing campaigning benefits the incumbent, and raises polling error overall. Message length also matters: longer messages increase the candidate's likelihood of winning (though for challengers,

Table 7: Effects of messages on polling outcomes (marginal effects, standard errors in parentheses)

	[9]	[10]	[11]	[12]
	Polling bias		Polling error	
Messages treatment (dummy)	0.107^{***}		0.068^{***}	
	(0.022)		(0.016)	
High–variance setting (dummy)	–0.023	–0.016	-0.035^{**}	–0.037
	(0.021)	(0.040)	(0.016)	(0.021)
Citizen income	–0.007	–0.009	0.010^{***}	0.007^{**}
	(0.006)	(0.006)	(0.003)	(0.003)
Incumbent message		0.120^{***}		0.035^{*}
length (in 100s)		(0.042)		(0.020)
Challenger message		–0.075		0.091^{***}
length (in 100s)		(0.047)		(0.027)
Supergame	–0.009	–0.008	–0.008	–0.007
	(0.007)	(0.008)	(0.004)	(0.005)
Round	–0.001	–0.005	–0.003	–0.006
	(0.005)	(0.009)	(0.002)	(0.003)
Sample	All groups	Messages treatment	All groups	Messages treatment
N	1555	808	1555	808

** (***): *Marginal effect significantly different from zero at the 5% (1%) level.*

the effect is insignificant, with a p–value of 0.11), and increase polling error.

5.4 Campaign messages – content analysis

We now examine the *content* of these messages. We hired three research assistants (RAs) who were not otherwise involved with this study. Each was given the experiment instructions and a data file that included – for each group–round – the session, supergame and round numbers, the group size and (representative) citizen's income, the straw–poll result, and the incumbent and challenger message texts. They were not told about our research questions, nor did they have access to candidates' private information (e.g., qualities or salary choices). They were paid a fixed wage and worked independently.

The RAs coded each message as (a) positive about the candidate; (b) negative about the opponent; (c) about own or opponent quality; (d) about past performance (including the official's current salary); and (e) about future decisions the candidate would take if elected.[20] We refer to these here as "positive campaigning", "negative campaigning", "quality statement", "performance statement", and "campaign promise" respectively, but these terms were not used in discussions with the coders. We classified a message as satisfying a criterion if at least two of the coders classified it that way.[21]

Figure 5 shows that associations between message content and straw–poll results are often non–monotonic. Exceptions exist; notably, challengers' negative campaigning uniformly decreases as the incumbent's straw vote improves (as bad polling likely reflects low incumbent quality or opportunistic salary choices). But of the 10 message types, 4 are significantly more likely in close elections (positive campaigning and quality statements by the challenger, promises by both candidates, Wilcoxon test, $p < 0.05$), and for one more the difference just misses significance (incumbent performance statements, $p \approx 0.084$).

Differences across candidate and message types are also illuminating. Positive campaigning

[20] Additional features of campaign messages may matter to voters: e.g., the level and tone of a message, or the use of particular wording. Advances in machine learning may facilitate more thorough analyses of natural–language messages and their effects in the lab and in real–life elections.

[21] There was moderate agreement among coders, with Fleiss's (1971) kappa equal to 0.496.

Figure 5: Message types in close elections and landslides

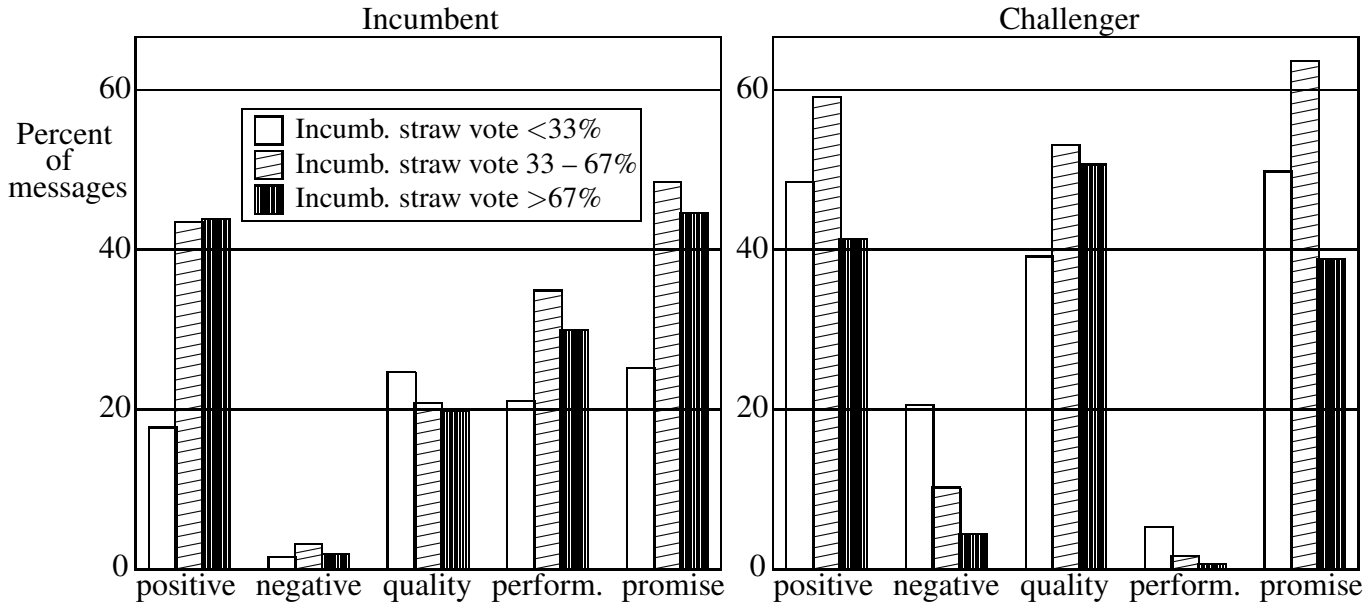

is far more common than negative campaigning, even for challengers.[22] Challengers rarely make statements about performance (though negative campaigning may implicitly concern incumbent performance), but compensate by making more quality statements and more promises. Incumbents make statements about quality and performance roughly equally often, though performance messages become more common as their polling improves.

Table 8 reports probit results with these message classifications as dependent variables. The main explanatory variables are dummies for the high–variance treatment and a high–quality candidate, their product, the incumbent straw–poll share and its square (to allow for non–monotonic relationships), and citizens' incomes. Additional variables include a dummy for incumbency and its product with these variables (to allow differing effects on incumbents and challengers), the supergame and round numbers, group–size dummies, and a constant. The sample was all candidate–

[22]Negative campaigning is extremely rare for incumbents. Given the lack of scope for such messages (the incumbent does not know the challenger's quality, and the challenger has no track record to criticise), this result is not surprising, but it is perhaps reassuring as verification that subjects understood the setting.

rounds in the Messages treatment.

The table shows the average marginal effects of the incumbency dummy, supergame and round, and those of other variables separately for incumbents and challengers.[23] Nearly all of the incumbency dummy's marginal effects are negative, with the exception of a positive effect on performance messages, suggesting that incumbents campaign less than challengers across–the–board. While we had not anticipated this result, it has a natural explanation: when messages are costly, incumbents rely on track records, which challengers cannot do. (The exception – performance messages – is probably also a consequence of incumbents running on their record.)

Also, the negative marginal effect of income for challengers, in the "negative campaigning" column, indicates (supporting part (i) of Hypothesis 5):

Result 5 *As citizens' incomes increase, challengers' negative campaigning decreases.*

By contrast, there is no evidence that citizens' income impacts positive campaigning, or whether challengers' messages are retrospective or prospective. Additionally, high–quality challengers are more likely to campaign positively, mention quality, and make promises than low–quality challengers.[24]

[23]Our quadratic specification for the straw vote makes its overall marginal effect hard to interpret, but Figure 6 in the appendix shows how poll results influence the messages candidates send.

[24]In our experiment, messages coded as "campaign promises" are arguably more precise about future intended behaviour than other messages. Previous research has found that while precise messages can be viewed as signals of candidate competence (Frenkel 2014), ambiguity is common and often not punished by voters (Kartik et al. 2017; Tomz and Van Houweling 2009). We find that higher–quality candidates send more precise messages (i.e., are more likely to make a campaign promise), consistent with Frenkel's result. Nonetheless, both precise and imprecise messages are sent (Figure 5), and the next section provides only equivocal evidence that voters prefer precise messages (e.g., Table 9), consistent with Kartik et al. and Tomz and Van Houweling. We additionally observe more precision under low variance, when polls indicate a close election, and from challengers, suggesting that candidates expect precision to be rewarded by voters, and that

The natural converse of Result 5 is that higher incomes should lead to more positive campaigning by the incumbent. Indeed, consistent with part (i) of Hypothesis 6:

Result 6 *As citizens' incomes increase, positive campaigning by the incumbent increases.*

Additionally, high–quality incumbents are *less* likely to mention quality or performance than low–quality incumbents.

We can also address the remaining parts of Hypotheses 5 and 6.

Result 7 *Challengers' negative and performance messaging are unaffected by quality variance.*

Result 8 *Incumbents are less likely to campaign positively in the high–variance setting.*

Result 7 does not support part (ii) of Hypothesis 5, while Result 8 strongly confirms Hypothesis 6's part (ii).[25]

The results of this section parallel campaign patterns in modern US presidential elections. Challengers attack incumbents who have performed poorly, while successful incumbents focus on their positives. For example, in the 1980 campaign, challenger Ronald Reagan famously attacked incumbent Jimmy Carter's economic record in a presidential debate, asking viewers "Are you better off than you were four years ago?". In the 1984 election, Reagan was a popular president, who

precision can partially substitute for other information such as incumbents' performance. These findings, while unhypothesised, are worthy of further investigation.

[25]In additional regressions, we examined the effect of the incumbent's tenure (i.e., the number of rounds the incumbent has been in office) on message choices. Regressions like those in Table 8, but including tenure and its product with the incumbency dummy, suggest that as incumbency lengthens, incumbents are less likely to send each of the five message types (and additionally send shorter messages), while there are no significant effects on challengers' messages. We similarly adapted the regressions in Table 9 (see next section), but observed no significant changes in the effects of candidates' message types on election outcomes according to incumbents' tenure, though we did find a general incumbency advantage as incumbents' tenure increased. So, incumbents campaign less the longer they have been in office, but this campaigning is arguably more efficient.

mostly campaigned positively, focussing on his strengths and his record on the economy and in the Cold War, as exemplified respectively by his "Morning in America" and "Bear in the Woods" television advertisements. By contrast, in 2020 the incumbent president Donald Trump – burdened by a poorly–handled pandemic response – concentrated his campaign messaging on claims about challenger Joe Biden's quality ("sleepy Joe").

5.5 Voter responses to campaign messages

We examine the impact of campaign messages with four probit regressions. The dependent variable is either a vote for the incumbent or re–election, so the sample is all subject–rounds excluding candidates, or all group–rounds, respectively. We estimate separate models for the Election–only and Messages treatments. For the Election–only probits, the right–hand–side variables are a constant, the incumbent straw–vote share and its square (to allow for non–linear effects), a dummy for the incumbent winning the straw vote, group–size dummies, and when applicable, a dummy for a straw–poll vote for the incumbent. For the Messages probits, we include those variables, message–characteristic dummies, their products with the incumbency dummy, and the products of all of these with the incumbent straw–vote share and its square (to allow these effects to depend on the polling).

The main results are shown in Table 9 (see also Figure 7 in the appendix). We observe that individuals typically vote in the election as they did in the straw poll. There is also an apparent bandwagon effect, with voters more likely to vote for a better–polling candidate, though winning the poll outright carries little additional benefit.

Some kinds of campaign messages work. Positive campaigning improves that candidate's vote share and election result, even after controlling for the poll (which should incorporate other information voters had from before the poll, like income). Campaign promises work for challengers, not significantly so for incumbents. Negative campaigning also does not work for incumbents, while for challengers it has a significant effect on the vote share.

Although unhypothesised, we summarise these regularities in Result 9.

Result 9 *Positive campaigning benefits the candidate, both for incumbents and challengers. There is weak evidence that negative campaigning and campaign promises benefit the candidate.*

6 Discussion

We introduce a dynamic model in which each period allows for the entire cycle of policy–making, campaigning and elections. We analyse the model and test its predictions in the lab. By considering two otherwise identical political settings, one which allows candidates to campaign and one which does not, we can ask the question of whether political campaigns benefit voters; we find that they do. The vast literature on political campaigns (Jacobson 2015) has largely sidestepped this question – focussing instead on whether and how campaigns change electoral outcomes – because by relying on observational data or field experiments, it cannot make the controlled other–things–equal comparison of campaigning versus no–campaigning.

A large literature has argued that political campaigning benefits challengers because it allows them to make themselves known to voters (e.g., Kahn and Kenney 1999; Jacobson 1990, 2006a, 2006b). Overall, we find the opposite – allowing campaigning gives incumbents a small advantage – but its impact depends on the candidates' qualities. Low–quality incumbents facing high–quality challengers are less likely to survive when campaigning is allowed, but do better when facing low–quality challengers. Conditionally on campaigns being allowed, we find that effort spent on campaigning benefits the candidate expending it, and is taken when polling indicates a close election; this is consistent with previous studies (Squire 1995; Jacobson 2013). However, we do not find that challengers have a higher return on campaigning than incumbents (Erikson and Palfrey 1998 observe a similar result to ours). Our results also address positive and negative campaigning. The literature does not provide decisive arguments in favour of negative campaigning (see Lau et al. 1999 and 2007). We similarly find that negative campaigning is infrequent compared to positive campaigning, and tends to be done by challengers when the incumbent has performed poorly.

As emphasised by Ashworth (2012), past performance may signal what the incumbent is likely

to do in future. Our treatment–level manipulation of quality variance allows us to verify this prediction directly (higher variance makes it easier to assess performance), complementing the previous field–experimental literature (Berry and Howell 2007; Snyder and Stromberg 2010). However, our manipulation of the possibility of messaging extends this literature: campaigning can facilitate selection even when accountability is difficult because quality variance is low.

We close by discussing some limitations of our study. One set of limitations is shared with most laboratory experiments: the features of the lab that provide its advantages also lead to questions about external validity. For example, our subject population of university students is probably fairly representative of voters.[26] However, it is less representative of politicians, who face multiple rounds of selection, and may therefore differ in economic traits (e.g., risk attitudes) and psychologically (e.g., extroversion, overconfidence, willingness to lie) from the general population. Some of these may affect levels of behaviour (e.g., if real officials are more risk–seeking, they should keep more of the endowment in all treatments, leading to lower citizens' incomes), though it is unclear whether qualitative differences between treatments would be affected. Some do have the potential to impact treatment effects (e.g., if real officials are less reluctant to lie, the effect of allowing campaigning may be weakened), though the ability of voters to punish lies in future elections may discipline officials' worst excesses.

Recent research suggests that the lab's abstractness on its own does not reduce external validity.[27] However, other differences between lab and field may matter. While our money stakes are high for a lab experiment, they are far below the stakes faced by real politicians. Thus, real

[26]Snowberg and Yariv (2021) report evidence that while students, other convenience samples (Mechanical Turk participants) and representative samples sometimes behave differently on average, differences across treatments are similar across these groups.

[27]For example, Duch and Stevenson (2013) find that subjects' ability to recognise agenda–setting power in the lab positively correlates with their ability to recognise it in a real legislature. Ambuehl et al. (2021) find that subjects' tastes for paternalism in an experimental setting (removing the opportunity for others to make extremely myopic decisions) correspond to their views about real–world paternalistic policies (e.g., taxes on sugary drinks).

officials may behave in a more risk–averse way than they would have in our lab setting. Higher stakes might also reduce noise in behaviour (Camerer and Hogarth 1999; Hertwig and Ortmann 2001). Finally, subjects in lab experiments do not have the experience and know–how of real–life campaign operations, nor is the lab suitable for distinguishing between campaigning through different media, suggesting that a more granular analysis of campaign messages is better suited to observational or field work.

A second set of limitations arises from design choices we made, but could be chosen differently in future studies. Here, all citizens receive the same income, so that their preferences over the official's choice are perfectly aligned. If officials made ideological decisions (e.g., on a political left–right spectrum), or could favour some citizens over others, this "solidarity" among citizens would be lost. Also, we have a unidimensional policy space, but multiple policy dimensions would allow officials to satisfy voters on some dimensions but not on others. Further, in our setting we do not allow for abstention, so we cannot address the issue of turnout. Changing these three factors could have many effects, but one would be a likely increase in negative campaigning, with candidates seeking to motivate their own supporters, and suppress turnout by the opponent's base. Similarly, challengers would be able to focus on those dimensions where incumbents disappointed. Future work concentrating on the effects of negative campaigning might include one or more of these features.

Our setting also had no communication to voters except for one–shot messages. In real elections, campaigns have a dynamic component and candidates react to each other's campaigns; they can use different methods; voters can talk to each other (in pairs or larger groups) and they receive messages from third–parties (media, partisan groups, etc.). For example, negative campaigning directly from candidates may be perceived as distasteful (we thank a referee for noting this), but the same message may be more palatable coming from a third–party (e.g. the 2004 "swift–boat" campaign against John Kerry), or in response to attacks from the opponent. Allowing for this diversity in communication may produce an increase in negative campaigning, but also a greater consensus amongst voters.

A final extension concerns polling. The impact of polling itself was not one of our research questions (our straw poll was intended for identification of campaigning's effects), but this is a topic worth investigating. A future experimental study could vary whether voters are polled before the election. We conjecture that compared to the no–polling case, allowing polling serves to coordinate voters' voting strategies, and to reduce noise in candidates' campaigning strategies (without polling, candidates cannot tailor their messages to how they are faring in the polls). Future work could also vary what fraction of voters are polled: all of them (as in our study) or a randomly– or self–selected sample. We speculate that polling a sample rather than the population would reduce the effects of polling compared to no polling.

Table 8: Factors affecting message characteristics – marginal effects (MEs) and std. errors

	[13] Positive campaigning	[14] Negative campaigning	[15] Quality message	[16] Performance message	[17] Campaign promise
Average ME					
Incumbent	-0.139^{***}	-0.078^{***}	-0.277^{***}	0.282^{***}	–0.085
	(0.030)	(0.011)	(0.050)	(0.027)	(0.048)
Supergame	–0.023	0.007^{***}	0.021^{**}	–0.006	-0.018^{**}
	(0.014)	(0.002)	(0.009)	(0.004)	(0.009)
Round	-0.009^{**}	0.007^{***}	-0.009^{***}	–0.006	-0.013^{***}
	(0.004)	(0.002)	(0.003)	(0.004)	(0.005)
ME conditional on incumbent					
High quality (dummy)	–0.041	–0.018	-0.120^{**}	-0.118^{***}	–0.018
	(0.038)	(0.019)	(0.053)	(0.043)	(0.041)
High–variance	-0.113^{**}	–0.009	0.043	–0.035	–0.055
setting (dummy)	(0.056)	(0.011)	(0.079)	(0.062)	(0.101)
Citizen income	0.061^{***}	0.000	0.018	0.017	0.012
	(0.017)	(0.003)	(0.013)	(0.019)	(0.014)
ME conditional on challenger					
High quality (dummy)	0.359^{***}	–0.022	0.211^{***}	–0.003	0.119^{***}
	(0.022)	(0.014)	(0.036)	(0.009)	(0.029)
High–variance	0.034	–0.024	0.108	0.002	–0.059
setting (dummy)	(0.049)	(0.026)	(0.057)	(0.012)	(0.079)
Citizen income	–0.003	-0.015^{**}	0.008	–0.001	–0.002
	(0.010)	(0.007)	(0.010)	(0.002)	(0.013)
Other variables: *Incumbent straw vote share (Figure 6), group–size dummies (not significant)*					

** (***): *Marginal effect significantly different from zero at the 5% (1%) level.* $N = 1616$.

Table 9: Non–candidates' responses to messages – marginal effects and standard errors

	[19] Vote for incumbent (Election–only)	[20] Re–election (Election–only)	[21] Vote for incumbent (Message)	[22] Re–election (Message)
Voted for incumbent	0.630^{***}		0.252^{***}	
in straw poll (dummy)	(0.023)		(0.021)	
Incumbent straw	0.248^{***}	0.998^{***}	0.442^{***}	0.692^{***}
vote share	(0.037)	(0.134)	(0.059)	(0.056)
Incumbent wins	0.018	–0.016	0.005	0.006
straw vote (dummy)	(0.018)	(0.030)	(0.025)	(0.030)
Incumbent messages				
Positive (dummy)			0.033^{**}	0.055^{**}
			(0.015)	(0.022)
Negative (dummy)			0.020	0.044
			(0.058)	(0.041)
Promise (dummy)			0.025	0.002
			(0.018)	(0.018)
Challenger messages				
Positive (dummy)			-0.088^{***}	-0.106^{***}
			(0.025)	(0.025)
Negative (dummy)			-0.059^{**}	–0.051
			(0.030)	(0.047)
Promise (dummy)			–0.053	-0.096^{***}
			(0.031)	(0.035)
Group–size dummies?	Yes	Yes	Yes	Yes
Other variables	—	—	*Quality, performance, message length (both candidates), group sizes*	
N	2809	747	3101	808

** (***): *Marginal effect significantly different from zero at the 5% (1%) level.*

References

Agranov, Marina, Jacob K. Goeree, Julian Romero and Leeat Yariv. 2018. "What makes voters turn out: the effects of polls and beliefs". *Journal of the European Economic Association* 16, 825–856.

Ambuehl, Sandro, B. Douglas Bernheim and Axel Ockenfels. 2021. "What motivates paternalism? An experimental study", *American Economic Review* 111, pp. 787–830.

Ansolabehere, Stephen, Shanto Iyengar, Adam Simon and Nicholas Valentino. 1994. "Does attack advertising demobilize the electorate?". *American Political Science Review* 88, 829–838.

Ansolabehere, Stephen and Shanto Iyengar. 1995. *Going Negative: How Attack Ads Shrink and Polarize the Electorate*. Free Press.

Arceneaux, Kevin. 2006. "Do campaigns help voters learn? A cross–national analysis". *British Journal of Political Science* 36, 159–173.

Arrow, Kenneth J. 1971. Essays in the Theory of Risk Bearing. New York: Markham Publishing.

Ashworth, Scott. 2012. "Electoral accountability: recent theoretical and empirical work". *Annual Review of Political Science* 15, pp. 183–201.

Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 2017. "Accountability and information in elections". *American Economic Journal: Microeconomics* 9, 95–138.

Azfar, Omar and William Nelson Jr. 2007. "Transparency, wages, and the separation of powers: an experimental analysis of corruption". *Public Choice* 130, 471–493.

Bhattacharya, Sourav, John Duffy and Sun-Tak Kim. 2014. "Compulsory versus voluntary voting: an experimental study". *Games and Economic behaviour* 84, 111–131.

Berry, Christopher and William Howell. 2007. "Accountability and local elections: rethinking retrospective voting". *Journal of Politics* 69, pp. 844–858.

Camera, Gabriele and Marco Casari. 2014. "The coordination value of monetary exchange: experimental evidence". *American Economic Journal: Microeconomics* 6, 290–314.

Camerer, Colin F. and Robin M. Hogarth. 1999. "The effects of financial incentives in experiments: a review and capital-labor-production framework". *Journal of Risk and Uncertainty* 19, pp. 7–42.

Cason, Timothy N. and Vai-Lam Mui. 2014. "Coordinating resistance through communication and repeated interaction". *Economic Journal* 124, F226–F256.

Clarke, Harold and Marianne Stewart. 1994. "Prospections, retrospections, and rationality: the 'bankers' model of presidential approval". *American Journal of Political Science* 38, 1104–1123.

Corazzini, Luca, Sebastian Kube, Michel A. Maréchal and Antonio Nicolò. 2014. "Elections and deceptions: an experimental study on the behavioural effects of democracy". *American Journal of Political Science* 58, 579–592.

Duch, Raymond and Randolph Stevenson (2013). "Voter perceptions of agenda power and attribution of responsibility for economic performance", *Electoral Studies* 32, pp. 512–516.

Duffy, John and Daniela Puzzello. 2014. "Gift exchange versus monetary exchange: theory and evidence". *American Economic Review* 104, 1735–1776.

Eckel, Catherine C. and Philip J. Grossman. 2008. "Forecasting risk attitudes: an experimental study of actual and forecast risk attitudes of women and men", *Journal of Economic behaviour & Organization* 68, 1–17.

Engle-Warnick, James and Robert L. Slonim. 2006. "Learning to trust in indefinitely repeated games". *Games and Economic Behavior* 54, 95–114.

Erikson, Robert and Thomas Palfrey. 1998. "Campaign spending and incumbency: an alternative simultaneous equations approach". *The Journal of Politics.* 60, pp. 355–373.

Erikson, Robert and Thomas Palfrey. 2000. "Equilibria in campaign spending games. Theory and data". *American Political Science Review* 94, 595–609.

Erikson, Robert, Michael MacKuen and James Stimson. 2000. "Bankers or peasants revisited: economic expectations and presidential approval". *Electoral Studies* 19, 295-312.

Feltovich, Nick. 2005. "Critical values for the robust rank–order test". *Communications in Statistics—Simulation and Computation* 34, 525–547.

Feltovich, Nick and Francesco Giovannoni. 2015. "Selection vs. accountability: an experimental investigation of campaign promises in a moral–hazard environment". *Journal of Public Economics* 126, 39–51.

Ferejohn, John. 1986. "Incumbent performance and electoral control". *Public Choice* 50, 5–25.

Finkel, Steven and John Geer. 1998. "A spot check: casting doubt on the demobilizing effect of attack advertising". *American Journal of Political Science* 42, 573–595.

Fischbacher, Urs. 2007. "z–Tree: Zurich toolbox for ready–made economic experiments". *Experimental Economics* 10, 171–178.

Fleiss, Joseph L. 1971. "Measuring nominal scale agreement among many raters". *Psychological Bulletin* 76, pp. 378–382.

Fréchette, Guillaume R. and Sevgi Yuksel. 2017. "Infinitely repeated games in the laboratory: four perspectives on discounting and random termination". *Experimental Economics* 20, 279–308.

Freedman, Paul and Ken Goldstein. 1999. "Measuring media exposure and the effects of negative campaign ads". *American Journal of Political Science* 43, pp. 1189–1208.

Frenkel, Sivan. 2014. “Competence and ambiguity in electoral competition”. *Public Choice* 159, 219–234.

Gerber, Alan S. 2004. “Does campaign spending work? Field experiments provide evidence and suggest new theory”. *American behavioural Scientist* 47, 541–574.

Harrington, Joseph and Gregory Hess. 1996. “A spatial theory of positive and negative campaigning”. *Games and Economic behaviour* 17, 209–229.

Harrison, Glenn W., Morten I. Lau and E. Elisabet Rutström. 2007. “Estimating risk attitudes in Denmark: a field experiment”. *Scandinavian Journal of Economics* 109, 341–368.

Hertwig, Ralph and Andreas Ortmann. 2001. “Experimental practices in economics: a methodological challenge for psychologists?” *Behavioral and Brain Sciences* 24, pp. 383–451.

Holt, Charles A. and Susan K. Laury. 2002. “Risk Aversion and Incentive Effects”. *American Economic Review* 92, 1644–1655.

Jacobson, Gary. 1990. “The effects of campaign spending in House elections: new evidence for old arguments”. *American Journal of Political Science* 34, pp. pp. 334–362.

Jacobson, Gary. 2006a. “Measuring campaign spending effects in U.S. House elections”. in Brady, Henry and Richard Johnson, eds. *Capturing Campaign Effects,* Ann Arbor, University of Michigan Press.

Jacobson, Gary. 2006b. “Campaign spending effects in U.S. Senate elections: evidence from the National Annenberg Election Survey”. *Electoral Studies*25, pp. 195—226.

Jacobson, Gary. 2013. *The Politics of Congressional Elections*. New York, Pearson.

Jacobson, Gary. 2015. “How do campaigns matter?”. *Annual Review of Political Science* 18, pp. 31–47.

Kahn, Kim Friedkin and Patrick Kenney. 1999. *The spectacle of U.S. Senate campaigns.* Princeton, Princeton University Press.

Kartik, Navin, Richard Van Weelden and Stephane Wolton. 2017. "Electoral ambiguity and political representation". *American Journal of Political Science* 61, 958–970.

Konrad, Kai A. and Raisa Sherif. 2019. "Sanctioning, selection, and pivotality in voting: theory and experimental results". *Constitutional Political Economy* 30, 330–357.

Landa, Dimitri. 2010. "Selection incentives and accountability traps: a laboratory experiment". Paper presented at the annual meeting of the American Political Science Association, Washington, DC.

Lau, Richard and Gerald Pomper. 2002. "Effectiveness of negative campaigning in US Senate elections". *American Journal of Political Science* 46, 47–66.

Lau, Richard and Ivy Brown Rovner. 2009. "Negative campaigning". *Annual Review of Political Science* 12, pp. 285–306

Lau, Richard, Lee Sigelman and Ivy Brown Rovner. 2007. "The effects of negative political campaigns: a meta–analytic reassessment". *Journal of Politics* 69, pp. 1176–1209.

Lau, Richard, Lee Sigelman, Caroline Heldman and Paul Babbitt. 1999. "The effects of negative political advertisements: a meta–analytic assessment". *American Political Science Review* 93, pp. 851–875.

Li, Hao and Wei Li. 2013. "Misinformation". *International Economic Review* 54, 253–277.

MacKuen, Michael, Robert Erikson, and James Stimson. 1992. "Peasants or bankers? The American electorate and the US economy". *American Political Science Review* 86, 597–611.

Mailath, George J. and Larry Samuelson. 2006. *Repeated Games and Reputations: Long–Run Relationships*. Oxford: Oxford University Press.

Mattes, Kyle and David P. Redlawsk. 2014. *The positive case for negative campaigning.* Chicago: University of Chicago Press.

Meyer, Donald J. and Jack Meyer. 2005. "Relative risk aversion: what do we know?" *Journal of Risk and Uncertainty* 31, 243–262.

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

Persson, Torsten, Gerard Roland, and Guido Tabellini. 1997. "Separation of powers and political accountability". *Quarterly Journal of Economics* 112, 1163–1202.

Persson, Torsten and Guido Tabellini. 2000. *Political Economics – Explaining Economic Policy*. Cambridge, MA: MIT Press.

Polborn, Matthias and David Yi. 2006. "Informative positive and negative campaigning". *Quarterly Journal of Political Science* 1, 351–371.

Rankin, Frederick W. (2006), "Requests and social distance in dictator games", *Journal of Economic behaviour & Organization* 60, 27–36.

Schipper, Burkhard and Hee Woo. 2019. "Political awareness, microtargeting of voters, and negative electoral campaigning", *Quarterly Journal of Political Science* 14, 41–88.

Siegel, Sidney and N. John Castellan, Jr. 1988. *Nonparametric Statistics for the behavioural Sciences*. New York: McGraw–Hill.

Sigelman, Lee and Eric Shiraev. 2002. "The rational attacker in Russia? Negative campaigning in Russian presidential elections". *Journal of Politics* 64, pp. 45–62.

Skaperdas, Stergios and Bernard Grofman. 1995. "Modeling negative campaigning". *American Political Science Review* 89, 49–61.

Snowberg, E. and L. Yariv. 2021. “Testing the waters: behavior across participant pools”. *American Economic Review* 111, 687–719.

Snyder, James and David Strömberg. 2010. “Press coverage and political accountability”. *Journal of Political Economy* 118, pp. 355–408.

Squire, Peverill. 1995. “Candidates, money, and voters: assessing the state of congressional elections research”. *Political Research Quarterly* 48, pp. 891–97.

Stratmann, Thomas. 2005. “Money in politics. A (partial) review of the literature”. *Public Choice* 124, pp. 135–156.

Tomz, Michael and Robert P. Van Houweling. 2009. “The electoral implications of candidate ambiguity”. *American Political Science Review* 103, 83–98.

Vavreck, Lynn. 2009. *The Message Matters: The Economy and Presidential Campaigns*. Princeton, Princeton University Press.

Walter, Annemarie S., Wouter van der Brug and Philip van Praag. 2014. “When the stakes are high: party competition and negative campaigning”. *Comparative Political Studies* 47, 550–573.

Wattenberg, Martin and Craig Brians. 1999. “Negative campaign advertising: demobilizer or mobilizer?” *American Political Science Review* 93, 891–899.

Wilson, Richard K. 2011. “The contribution of behavioural economics to political science”. *Annual Review of Political Science* 14, 201–223.

Woon, Jonathan. 2012. “Democratic accountability and retrospective voting: a laboratory experiment”. *American Journal of Political Science* 56, 913–930.

Appendix: Campaign messages, polling, and elections: theory and experimental evidence

Intended for online publication only

Appendix: table of contents

A Additional analysis

A.1 Detailed theoretical analysis and hypotheses

We have two versions of the game, one with campaign promises and one without. We begin with the latter. From a citizen's perspective, this is a symmetric game, so we start by restricting attention to equilibria where non–candidate voters select identical voting strategies.[28] In these equilibria, no citizen is pivotal, so *any* symmetric deterministic voting rule is sequentially rational, meaning that verifying whether a corresponding strategy profile is an equilibrium boils down to verifying whether the official responds optimally to the voting rule.

We further restrict attention to stationary equilibria in which citizens vote to re–elect whenever their current–round income meets some threshold level – though as we have noted, there are many additional equilibria. The equilibria we examine correspond to those described in the canonical accountability model of elections (Ferejohn 1986; Persson and Tabellini 2000). Though we do not require voters to use threshold strategies in the experiment, Feltovich and Giovannoni (2015) find that in similar settings, citizens tend to vote this way. Further, we have a straw poll before the election. This, formally, is a cheap–talk phase where each citizen has the opportunity to send a message to the others, and as in all cheap–talk games, there exists a plethora of equilibria, including a babbling equilibrium where voters randomise in the straw poll and the results are ignored, and many others. However, as it is common knowledge that all non–candidate voters have the same preferences and constitute a majority, there also exists an equilibrium where these citizens vote in the straw poll in the same way they intend to vote in the actual election, given their current information. We will assume that this is the equilibrium being played. As a consequence, in the treatment without campaign promises, where nearly nothing happens between the straw poll and the actual election, we expect non–candidate citizens to vote the same way. In the treatment

[28] This implies that the candidates themselves will never be pivotal. Moreover, since there are positive rents for being in office in all of our equilibria, it is weakly dominant to vote for oneself, so candidates' votes will cancel each other out. We therefore focus on non–candidates' voting.

with campaign messages, we attribute any difference in voting between straw poll and election to campaign messages.

Stationary equilibria come in three classes. In *pooling equilibria*, the citizens' income threshold is sufficiently small that both high– and low–quality officials are willing to *comply* (specifically, they choose their salary so that citizens' income is exactly the threshold), and hence the official is always re–elected. Let $x_t = x$ be the equilibrium income for all t. The official must prefer to comply in each round t, rather than taking the entire endowment, losing the subsequent election, and getting the citizen's income thenceforth (because the new official would comply in each round).[29] We have

Proposition 1 *There exists a value $x^*(\theta^l) \in \left[0, \frac{n}{n-1}\theta^l\right]$ such that for any $x \leq x^*(\theta^l)$ there is an equilibrium where citizens re–elect if and only if their income is at least x, and vote in the same way in the straw poll. Officials of both quality always provide an income x to citizens.*

Proof As discussed above, we only need to check the official's incentives to comply with the equilibrium. For any $\theta \in \left\{\theta^l, \theta^h\right\}$ and any period t, providing an income x to the citizens is weakly preferred to deviating by taking all the endowment in this period and then getting x as a citizen thenceforth whenever

$$u\left(n\theta - x\left(n-1\right)\right)\frac{1}{1-\delta} \geq u\left(n\theta\right) + \frac{\delta}{1-\delta}u(x)$$

or, equivalently,

$$u\left(n\theta - x\left(n-1\right)\right) - \delta u(x) \geq \left(1-\delta\right)u\left(n\theta\right) \tag{2}$$

Since u is strictly increasing, the left hand side of (2) is strictly decreasing in x. Also, (2) strictly holds for $x = 0$ and cannot hold for $x = \frac{n}{n-1}\theta$, the maximum citizen income. By continuity, there must therefore be a value $x^*\left(\theta\right)$ such that (2) is satisfied with equality for this income and strictly

[29]Note that since this latter option is always available to an official, any equilibrium must yield positive rents from being in office. That is, officials will always earn higher payoffs than ordinary citizens in equilibrium.

satisfied for any $x < x^*(\theta)$. In addition, implicit differentiation yields

$$\frac{\partial x^*(\theta)}{\partial \theta} = \frac{n\left[u'(n\theta - x^*(\theta)(n-1)) - (1-\delta)u'(n\theta)\right]}{(n-1)u'(n\theta - x^*(\theta)(n-1)) + \delta u'(x^*(\theta))}$$

which is positive for any (weakly) concave u. So, we have that

$$x^*\left(\theta^h\right) > x^*(\theta^l)$$

and which implies that a pooling equilibrium exists whenever the citizen's re–election strategy asks for an income $x \leq x^*\left(\theta^l\right)$.■

In *separating equilibria*, the citizens' threshold is sufficiently high that only high–quality officials will comply; low–quality officials take the whole endowment and are voted out. So, unlike in pooling equilibria, we may observe initial sequences of rounds where citizens get zero income and vote out the incumbent. However, once a high–quality person becomes the official, every round will have citizens obtaining the threshold income and incumbents being re–elected.

Proposition 2 *There exists a value* $x^{**}(\theta^l, \theta^h) \in \left(x^*(\theta^l), \frac{n}{n-1}\theta^h\right]$ *such that for any* $x^*(\theta^l) < x \leq x^{**}(\theta^l, \theta^h)$ *there exists an equilibrium where, in each round, citizens re–elect if their income is at least* x*, and vote in the same way in the straw poll. High–quality officials always provide income* x *to citizens, and low–quality officials take the entire endowment.*

Proof Similarly to pooling equilibria, we only need to determine the highest citizen income at which high–quality officials will comply. Formally, we need that for any re–election income x,

$$u\left(n\theta^h - x(n-1)\right)\frac{1}{1-\delta} \geq u\left(n\theta^h\right) + \delta V(x) \tag{3}$$

where $V(x)$ is the continuation value for a high–quality incumbent, if she deviates just this once from compliance with the demand for x from these citizens. In actual fact, we can write

$$V(x) = \frac{1}{2}\left(\frac{1}{1-\delta}\right)u(x) + \frac{1}{2}\delta W(x)$$

In period $t+1$, the official at time t will no longer be the official because of her deviation, so she will definitively be a citizen. In that case, either the new incumbent will be of high quality (and

the old incumbent will get x in each period) or not, in which case, the old incumbent will get zero at $t+1$ and a continuation value $W(x)$ from $t+2$. So, $W(x)$ is the $t+2$ continuation value for a high–quality incumbent who deviates at t, if her replacement is of low quality and is herself replaced at $t+1$. We have

$$\begin{aligned} W(x) &= \frac{1}{n-1}\left(\frac{1}{2}u\left(n\theta^h - x(n-1)\right)\frac{1}{1-\delta} + \frac{1}{2}\left(u\left(n\theta^l\right) + \delta V(x)\right)\right) \\ &\quad + \frac{n-2}{n-1}\left(\frac{1}{2}\left(\frac{1}{1-\delta}\right)u(x) + \frac{1}{2}\left(\delta W(x)\right)\right) \\ &= \frac{1}{n-1}\left(\frac{1}{2}u\left(n\theta^h - x(n-1)\right)\frac{1}{1-\delta} + \frac{1}{2}\left(u\left(n\theta^l\right) + \delta V(x)\right)\right) + \frac{n-2}{n-1}V(x) \end{aligned}$$

With probability $\frac{1}{n-1}$,the old incumbent is selected as the challenger and wins since the new incumbent (who has low quality and will take all the endowment) is replaced. Then if the old incumbent is of high quality again she expects to get $u\left(n\theta^h - x(n-1)\right)$ in each future period, whereas if she is of low quality, she will get $\left(u\left(n\theta^l\right) + \delta V\right)$ because the continuation value $V(x)$ for an high–quality type who has deviated and taken everything and a low quality type who has taken everything (which is equilibrium behaviour) is the same. With probability $\frac{n-2}{n-1}$, the old incumbent remains a citizen and the continuation value is $V(x)$. Solving for $W(x)$ and $V(x)$ gives us

$$V(x) = \frac{\left[2u(x)(n-1) + \delta u\left(n\theta^h - x(n-1)\right) + \delta(1-\delta)u\left(n\theta^l\right)\right]}{(1-\delta)(2-\delta)(2n+\delta-2)}$$

and so (3) becomes

$$u\left(n\theta^h - x(n-1)\right)\frac{1}{1-\delta} \geq u\left(n\theta^h\right) + \delta\frac{\left[2u(x)(n-1) + \delta u\left(n\theta^h - x(n-1)\right) + \delta(1-\delta)u\left(n\theta^l\right)\right]}{(1-\delta)(2-\delta)(2n+\delta-2)} \tag{4}$$

or equivalently,

$$\begin{aligned} &2\frac{(2n+\delta-2) - \delta(n-1+\delta)}{(1-\delta)(2-\delta)(2n+\delta-2)}u\left(n\theta^h - x(n-1)\right) - 2\delta\frac{n-1}{(1-\delta)(2-\delta)(2n+\delta-2)}u(x) \\ &\geq u\left(n\theta^h\right) + \frac{\delta^2}{(2-\delta)(2n+\delta-2)}u\left(n\theta^l\right) \end{aligned} \tag{5}$$

Clearly, the left hand side of (5) is strictly decreasing in x. Notice that the right hand side is increasing in θ^l and let $x=0$. Then satisfying (5) becomes equivalent to satisfying

$$\left(4n + 3\delta - 2n\delta - \delta^2 - 4\right)u\left(n\theta^h\right) - 2(n-1)u(0) - \delta(1-\delta)u\left(n\theta^l\right) \geq 0$$

but if $\theta^l = \theta^h$ then we have that the left hand side becomes

$$2\left(2-\delta\right)\left(n-1\right)u\left(n\theta^h\right) > 0$$

so that (5) holds strictly for $x = 0$. Suppose now that $x = \frac{n}{n-1}\theta^h$. For this value, (4) becomes

$$u\left(0\right)\frac{1}{1-\delta} \geq u\left(n\theta^h\right) + \delta\frac{\left[2u\left(\frac{n}{n-1}\theta^h\right)\left(n-1\right) + \delta u\left(0\right) + \delta\left(1-\delta\right)u\left(n\theta^l\right)\right]}{\left(1-\delta\right)\left(2-\delta\right)\left(2n+\delta-2\right)}$$

which holds if and only if

$$0 \geq (1-\delta)u\left(n\theta^h\right) + \delta\frac{\left[2u\left(\frac{n}{n-1}\theta^h\right)\left(n-1\right) + \delta\left(1-\delta\right)u\left(n\theta^l\right)\right]}{\left(2-\delta\right)\left(2n+\delta-2\right)} \tag{6}$$

holds. Now, it is easy to see that the right hand side of (6) is strictly positive so for $x = \frac{n}{n-1}\theta^h$ (4) cannot hold. These results and the fact that u is continuous and strictly decreasing imply that there exists a $x^{**}\left(\theta^l, \theta^h\right) \in \left(0, \frac{n}{n-1}\theta^h\right)$ such that (4) holds with equality, and so the high–quality types will comply with any re–election rule that demands an income $x \leq x^{**}\left(\theta^l, \theta^h\right)$.

To compare $x^{**}\left(\theta^l, \theta^h\right)$ and $x^*(\theta^l)$ we begin by assuming that $\theta^h = \theta^l$, thus first comparing $x^{**}\left(\theta^l, \theta^l\right)$ and $x^*(\theta^l)$. While $x^*(\theta^l)$ is the solution to the equation

$$f\left(x, \theta^l\right) = u\left(n\theta^l - x\left(n-1\right)\right) - \left(1-\delta\right)u\left(n\theta^l\right) - \delta u(x) = 0$$

it can be shown that $x^{**}\left(\theta^l, \theta^l\right)$ is the solution to the equation

$$\frac{f\left(x, \theta^l\right)\left(2-\delta\right)\left(n-1\right) + \delta\left(1-\delta\right)\left[u\left(n\theta^l - x\left(n-1\right)\right) - u\left(n\theta^l\right) + u(x)(n-1)\right]}{\left(2-\delta\right)n + \left(2\delta - \delta^2 - 2\right)} = 0 \tag{7}$$

For $x = x^*(\theta^l)$, the left hand side of (7) becomes

$$\begin{aligned}&\frac{\delta\left(1-\delta\right)\left[\left(1-\delta\right)u\left(n\theta^l\right) + \delta u(x^*(\theta^l)) - u\left(n\theta^l\right) + u(x^*(\theta^l))(n-1)\right]}{\left(2-\delta\right)n + \left(2\delta - \delta^2 - 2\right)} \\ &= \frac{\delta\left(1-\delta\right)\left[-\delta u\left(n\theta^l\right) + \left(n-1+\delta\right)u(x^*(\theta^l))\right]}{\left(2-\delta\right)n + \left(2\delta - \delta^2 - 2\right)}\end{aligned} \tag{8}$$

Now, (8) is (weakly) positive if and only if

$$u(x^*(\theta^l)) \geq \frac{\delta u\left(n\theta^l\right)}{n-1+\delta}$$

When u is the identity function, the above holds with equality, so for risk–neutral preferences $x^{**}_{RN}\left(\theta^l,\theta^l\right)=x^{*}_{RN}(\theta^l)=\frac{\delta n}{n-1+\delta}\theta^l$. Now, assume u is strictly concave. Then we must have

$$u(x^*(\theta^l))>u(x^*_{RN}\left(\theta^l\right))=u\left(\frac{\delta}{n-1+\delta}n\theta^l+\frac{n-1}{n-1+\delta}\times 0\right)$$
$$>\frac{\delta}{n-1+\delta}u\left(n\theta^l\right)+\frac{n-1}{n-1+\delta}u(0)=\frac{\delta u\left(n\theta^l\right)}{n-1+\delta}$$

which implies that (7) will be satisfied at $x^{**}\left(\theta^l,\theta^l\right)>x^*(\theta^l)$. Now,

$$\frac{\partial x^{**}\left(\theta^l,\theta^h\right)}{\partial\theta^h}=$$
$$\frac{n\left[2\frac{(2n+\delta-2)-\delta(n-1+\delta)}{(1-\delta)(2-\delta)(2n+\delta-2)}u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)-u'\left(n\theta^h\right)\right]}{2\frac{(2n+\delta-2)-\delta(n-1+\delta)}{(1-\delta)(2-\delta)(2n+\delta-2)}(n-1)u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)+2\delta\frac{n-1}{(1-\delta)(2-\delta)(2n+\delta-2)}u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)}>0$$

which in turn implies that $x^{**}\left(\theta^l,\theta^h\right)>x^{**}\left(\theta^l,\theta^l\right)\geq x^*(\theta^l)$.■

In the third class of equilibria, voters set their income threshold high enough that both types of official take the entire endowment, ensuring they lose the election and are replaced by another official who does the same in the next round. Such *kleptocratic* equilibria always exist.

Our *low–variance* and *high–variance* treatments differ in that the latter has a higher θ^h and a lower θ^l, with average quality $(\theta^h+\theta^l)/2$ the same in both treatments. One of our objectives is to study the impact of an increase in quality variance: a measure of how much candidate selection matters. Given the existence of these three types of equilibria, we want to study how higher variance affects the way these equilibria are played. We do this in two ways. First, Proposition 3 shows that when variance increases, citizens' expected lifetime utility in the highest–income–threshold separating equilibrium increases relatively to that in the highest–threshold pooling equilibrium or in kleptocratic equilibria. This suggests that citizens will be more likely to play separating equilibria in a high–variance environment and provides predictions about re–election probabilities, both conditionally on income and unconditionally. Second, under (mild) additional assumptions on citizens' preferences, expected lifetime *citizen income* is higher in a separating equilibrium, allowing us to argue that in higher–variance environments we should expect higher incomes for citizens. To begin:

Proposition 3 *When we increase the variance in the type distribution, citizens lifetime expected utility in the highest–threshold separating equilibrium increases relatively to their lifetime expected utility in (i) the highest–threshold pooling equilibrium and (ii) in a kleptocratic equilibrium.*

Proof Suppose $n \geq 1$, as in our experiment, ensuring that the expressions below are well–defined and candidates' votes are never pivotal. Let u_c^S, u_c^P and u_c^K be the lifetime expected utility for citizens in the highest–threshold separating, pooling, and kleptocratic equilibria, respectively. The expected utilities from each equilibrium type are

$$\begin{aligned}
u_c^P &= \frac{1}{1-\delta} u\left(x^*(\theta^l)\right) \\
u_c^S &= \frac{1}{2}\left[\frac{1}{1-\delta} u\left(x^{**}\left(\theta^l,\theta^h\right)\right)\right] + \frac{1}{2}\delta W\left(x^{**}\left(\theta^l,\theta^h\right)\right) \\
&= \frac{1}{2}\left[\frac{1}{1-\delta} u\left(x^{**}\left(\theta^l,\theta^h\right)\right)\right] \\
&+ \frac{1}{2}\delta \frac{(2n+\delta-4)\, u\left(x^{**}\left(\theta^l,\theta^h\right)\right) + 2u\left(n\theta^h - x^{**}\left(\theta^l,\theta^h\right)(n-1)\right) + 2(1-\delta)\, u\left(n\theta^l\right)}{(1-\delta)(2-\delta)(2n+\delta-2)} \\
&= \frac{2(n-1)\, u\left(x^{**}\left(\theta^l,\theta^h\right)\right) + \delta u\left(n\theta^h - x^{**}\left(\theta^l,\theta^h\right)(n-1)\right) + \delta(1-\delta)\, u\left(n\theta^l\right)}{(1-\delta)(2-\delta)(2n+\delta-2)} \\
u_c^K &= \delta \widetilde{V} = \frac{\delta}{2(1-\delta)(n-1)}\left(u\left(n\theta^h\right) + u\left(n\theta^l\right)\right)
\end{aligned}$$

since

$$W\left(x^{**}\left(\theta^l,\theta^h\right)\right) = \frac{(2n+\delta-4)\, u\left(x^{**}\left(\theta^l,\theta^h\right)\right) + 2u\left(n\theta^h - x^{**}\left(\theta^l,\theta^h\right)(n-1)\right) + 2(1-\delta)\, u\left(n\theta^l\right)}{(1-\delta)(2-\delta)(2n+\delta-2)}$$

and

$$\widetilde{V} = \frac{1}{n-1}\left(\frac{1}{2} u\left(n\theta^h\right) + \frac{1}{2} u\left(n\theta^l\right)\right) + \delta \widetilde{V}.$$

Now, we can show that

$$\begin{aligned}
\frac{\partial\left[u_c^P - u_c^K\right]}{\partial \theta^l} &= \frac{1}{1-\delta}\left[u'\left(x^*\left(\theta^l\right)\right) \frac{n\left[u'\left(n\theta^l - x^*(\theta^l)(n-1)\right) - (1-\delta)\, u'\left(n\theta^l\right)\right]}{(n-1)\, u'\left(n\theta^l - x^*(\theta^l)(n-1)\right) + \delta u'\left(x^*\left(\theta^l\right)\right)} - \frac{\delta n}{2(n-1)} u'\left(n\theta^l\right)\right] \\
&= n \frac{(n-1)\, u'\left(n\theta^l - x^*\left(\theta^l\right)(n-1)\right)\left(2u'\left(x^*(\theta^l)\right) - \delta u'\left(n\theta^l\right)\right) - \left[2(n-1)(1-\delta) + \delta^2\right] u'\left(n\theta^l\right) u'\left(x^*(\theta^l)\right)}{2(n-1)\left[(n-1)\, u'\left(n\theta^l - x^*(\theta^l)(n-1)\right) + \delta u'\left(x^*(\theta^l)\right)\right]}
\end{aligned}$$

which has the same sign as

$$u^{'}\left(n\theta^{l}-x^{*}(\theta^{l})\left(n-1\right)\right)-u^{'}\left(n\theta^{l}\right)\frac{\left(2\left(1-\delta\right)\left(n-1\right)+\delta^{2}\right)u^{'}\left(x^{*}(\theta^{l})\right)}{\left(n-1\right)\left(2u^{'}\left(x^{*}(\theta^{l})\right)-\delta u^{'}\left(n\theta^{l}\right)\right)}$$
$$>u^{'}\left(n\theta^{l}-x^{*}(\theta^{l})\left(n-1\right)\right)-u^{'}\left(n\theta^{l}\right)>0$$

Also,

$$\frac{\partial\left[u_{c}^{S}-u_{c}^{K}\right]}{\partial\theta^{l}}$$
$$=\frac{2\left(n-1\right)u^{'}\left(x^{**}\left(\theta^{l},\theta^{h}\right)\right)\frac{\partial x^{**}\left(\theta^{l},\theta^{h}\right)}{\partial\theta^{l}}-\left(n-1\right)\delta u^{'}\left(n\theta^{h}-x^{**}\left(\theta^{l},\theta^{h}\right)\left(n-1\right)\right)\frac{\partial x^{**}\left(\theta^{l},\theta^{h}\right)}{\partial\theta^{l}}+n\delta\left(1-\delta\right)u^{'}\left(n\theta^{l}\right)}{\left(1-\delta\right)\left(2-\delta\right)\left(2n+\delta-2\right)}$$
$$-\frac{\delta n}{2\left(1-\delta\right)\left(n-1\right)}u^{'}\left(n\theta^{l}\right)$$

where

$$\frac{\partial x^{**}\left(\theta^{l},\theta^{h}\right)}{\partial\theta^{l}}=\frac{-n\frac{\delta^{2}}{(2-\delta)(2n+\delta-2)}u^{'}\left(n\theta^{l}\right)}{2\frac{(2n+\delta-2)-\delta(n-1+\delta)}{(1-\delta)(2-\delta)(2n+\delta-2)}\left(n-1\right)u^{'}\left(n\theta^{h}-x^{**}\left(\theta^{l},\theta^{h}\right)\left(n-1\right)\right)+2\delta\frac{n-1}{(1-\delta)(2-\delta)(2n+\delta-2)}u^{'}\left(x^{**}\left(\theta^{l},\theta^{h}\right)\right)}$$

which gives us

$$\frac{\partial\left[u_{c}^{S}-u_{c}^{K}\right]}{\partial\theta^{l}}$$
$$=-\frac{n\delta u^{'}\left(n\theta^{l}\right)}{2\left(n-1\right)\left(1-\delta\right)}\frac{\delta u^{'}\left(x^{**}\left(\theta^{l},\theta^{h}\right)\right)+\left[\left(n-1\right)+\delta\left(1-\delta\right)\right]u^{'}\left(n\theta^{h}-x^{**}\left(\theta^{l},\theta^{h}\right)\left(n-1\right)\right)}{\delta u^{'}\left(x^{**}\left(\theta^{l},\theta^{h}\right)\right)+\left[2\left(n-1\right)-\delta\left(n-2\right)-\delta^{2}\right]u^{'}\left(n\theta^{h}-x^{**}\left(\theta^{l},\theta^{h}\right)\left(n-1\right)\right)}$$
$$<0$$

and also implies $\frac{\partial\left[u_c^P-u_c^S\right]}{\partial\theta^l}>0$. Finally, $\frac{\partial u_c^K}{\partial\theta^h}>0$ and $\frac{\partial u_c^P}{\partial\theta^h}=0$ are obvious while

$$\frac{\partial\left[u_c^S-u_c^K\right]}{\partial\theta^h}$$

$$=\frac{1}{1-\delta}\left[\frac{2(n-1)u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)-(n-1)\delta u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}{(2-\delta)(2n+\delta-2)}\frac{\partial x^{**}\left(\theta^l,\theta^h\right)}{\partial\theta^h}+\frac{n\delta u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}{(1-\delta)(2-\delta)(2n+\delta-2)}-\frac{\delta n}{2(n-1)}u'\left(n\theta^h\right)\right]$$

$$=\frac{n\delta}{2(n-1)(1-\delta)}\left\{\begin{array}{c}\frac{u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)\left(2(1-\delta)(n-1)+\delta^2\right)\left[u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)-u'\left(n\theta^h\right)\right]}{\delta u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)+[2(n-1)-\delta(n-2)-\delta^2]u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}+\\ \frac{u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)\left[\left(\delta(2n-\delta-2)u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)-\delta\left(n+\delta-\delta^2-1\right)u'\left(n\theta^h\right)\right)\right]}{\delta u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)+[2(n-1)-\delta(n-2)-\delta^2]u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}\end{array}\right\}$$

$$>\frac{n\delta}{2(n-1)(1-\delta)}\left\{\begin{array}{c}\frac{u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)\left(2(1-\delta)(n-1)+\delta^2\right)\left[u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)-u'\left(n\theta^h\right)\right]}{\delta u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)+[2(n-1)-\delta(n-2)-\delta^2]u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}+\\ \frac{u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)\left[\left(\delta(2n-\delta-2)u'\left(n\theta^h\right)-\delta\left(n+\delta-\delta^2-1\right)u'\left(n\theta^h\right)\right)\right]}{\delta u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)+[2(n-1)-\delta(n-2)-\delta^2]u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}\end{array}\right\}$$

$$=\frac{n\delta}{2(n-1)(1-\delta)}$$

$$\times\left\{\frac{u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)\left(2(1-\delta)(n-1)+\delta^2\right)\left[u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)-u'\left(n\theta^h\right)\right]+u'\left(n\theta^h\right)u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)\delta\left(n-2\delta+\delta^2-1\right)}{\delta u'\left(x^{**}\left(\theta^l,\theta^h\right)\right)+[2(n-1)-\delta(n-2)-\delta^2]u'\left(n\theta^h-x^{**}\left(\theta^l,\theta^h\right)(n-1)\right)}\right\}>0$$

This establishes that

$$\begin{array}{rcl}\frac{\partial u_c^P}{\partial\theta^l} & > & \frac{\partial u_c^K}{\partial\theta^l}>\frac{\partial u_c^S}{\partial\theta^l}\text{ and}\\ \frac{\partial u_c^S}{\partial\theta^h} & > & \frac{\partial u_c^K}{\partial\theta^h}>\frac{\partial u_c^P}{\partial\theta^h}=0\end{array}$$

Now, to complete the proof, consider $\left(\theta^l,\theta^h\right)$ and $\left(\psi^l,\psi^h\right)$ such that $\psi^h>\theta^h$ and $\psi^l<\theta^l$. Then, from the above we have

$$u_c^K\left(\theta^l,\theta^h\right)-u_c^K\left(\psi^l,\psi^h\right)>u_c^K\left(\theta^l,\theta^h\right)-u_c^K\left(\psi^l,\theta^h\right)>$$
$$u_c^S\left(\theta^l,\theta^h\right)-u_c^S\left(\psi^l,\theta^h\right)>u_c^S\left(\theta^l,\theta^h\right)-u_c^S\left(\psi^l,\psi^h\right)$$

But

$$\begin{array}{rcl}u_c^K\left(\theta^l,\theta^h\right)-u_c^K\left(\psi^l,\psi^h\right) & > & u_c^S\left(\theta^l,\theta^h\right)-u_c^S\left(\psi^l,\psi^h\right)\Leftrightarrow\\ u_c^S\left(\psi^l,\psi^h\right)-u_c^S\left(\theta^l,\theta^h\right) & > & u_c^K\left(\psi^l,\psi^h\right)-u_c^K\left(\theta^l,\theta^h\right)\end{array}$$

An analogous argument holds with u_c^P replacing u_c^K.■

Proposition 3 tells us that increasing quality variance will increase the difference between x^{**} and x^* as the former increases and the latter decreases. More importantly, the increase on the

returns citizens can expect from the most favourable separating equilibrium in comparison to the most favourable pooling equilibrium or the kleptocratic equilibrium should increase the likelihood of a separating equilibrium. This in turn affects re–election probabilities, with low (high) quality officials less (more) likely to be re–elected in the high–variance treatment than in the low–variance treatment. This is because low–quality types are only re–elected in pooling equilibria (not in separating or kleptocratic equilibria, where they choose not to comply with voters' demands), and pooling equilibria are less likely when there is high variance. Similarly, it is only in kleptocratic equilibria that high–quality officials will *not* comply, and since these latter equilibria are less likely to obtain when there is high variance, the probability of observing high–quality officials being re–elected will also be higher. This yields Hypothesis 4.

Proposition 3 also entails a prediction about re–election probabilities conditional on incomes. In a kleptocratic equilibrium, the official takes the entire endowment, so any positive incomes imply that either a pooling or a separating equilibrium is being played, meaning that within a treatment, a higher observed income leads to a higher probability of winning. More interestingly, we can make predictions across treatments. Since the separating equilibrium is more likely in the high–variance treatment, it is less likely that a given observed income will comply with the citizens' threshold for re–election. Also, high–quality officials will comply with higher demands for citizens' incomes, so one can expect citizens to raise the bar for re–election (Hypothesis 3).

In order to make predictions about the relationship between quality variance and expected lifetime citizen income, we need to be able to argue that the latter is higher in a separating equilibrium than in a pooling equilibrium.[30] To do so, we first assume that citizens use x^* and x^{**} as their income thresholds for re–election in all pooling and separating equilibria respectively. Then, expected citizen lifetime income in the pooling equilibrium is

$$\sum_{t=0}^{\infty} \delta^t x^* = \frac{x^*}{1-\delta}$$

whereas in a separating equilibrium, given that types are equally likely and a low type is always

[30]Note that this is not the same as the expected lifetime income for a player, because ex–ante there is a positive probability that the player will begin as an official.

replaced, we have

$$\sum_{t=0}^{\infty} \delta^t \left[1 - \left(\frac{1}{2}\right)^{t+1}\right] x^{**} = \frac{x^{**}}{(1-\delta)(2-\delta)}$$

so that lifetime citizen income is higher in a separating equilibrium if and only if $x^{**} > (2-\delta)\, x^*$.

To check this condition, we parameterise the utility function:

$$u(x) = \frac{1-\gamma}{\gamma}\left(\frac{\beta x}{1-\gamma} + \eta\right)^{\gamma} \tag{9}$$

where $\gamma \in (0,1)$, $\beta > 0$ and $\eta > -\beta x/(1-\gamma)$. This *hyperbolic absolute risk aversion* form is very general, allowing for increasing, decreasing, or constant relative risk aversion as $\eta > 0$, $\eta < 0$, or $\eta = 0$; in the last case it converges to risk neutrality as $\gamma \to 1$. With (9) applied to (3) and (4), we can show that for the parameters in our experiment and when $\eta \geq 0$, the distance $x^{**} - x^*$ is minimised for $\gamma \to 1$, the risk–neutral asymptotic case. Thus, if $x^{**} > (2-\delta)\, x^*$ holds for the risk–neutral case, it will hold for any value of $\gamma \in (0,1)$ under either constant or increasing relative risk aversion. For the risk–neutral case,

$$x^{**} - (2-\delta)\, x^* = \frac{1}{2} n\delta \frac{\theta^h\left(4n + 3\delta - 2n\delta - \delta^2 - 4\right) + \theta^l\left(8n\delta - 9\delta - 8n + 3\delta^2 - 2n\delta^2 + 8\right)}{(2-\delta)\,(n-1)\,(n+\delta-1)}$$

which is positive for our experimental parameters $\delta = 0.8$, $n = 5, 6, 7$ and $\left(\theta^l, \theta^h\right) \in \{(5,6), (3,8)\}$.

We summarise the discussion with the following:

Remark 1 *If players have hyperbolic risk–averse preferences with increasing or constant relative risk aversion, then the lifetime citizen income in a separating equilibrium is higher than that in a pooling equilibrium.*

This, together with our previous argument from Proposition 3 that separating equilibria should be more likely in the high–variance treatment, leads us to Hypothesis 1.

We now turn to the game with campaign messages. These take place after the straw poll results are announced, and before the election. Adding communication does not necessarily change the set of equilibrium outcomes, since there always exist "babbling" equilibria in which voters ignore all communication. so that candidates can say whatever they like, making the game equivalent to one

without communication. However, communication may help citizens coordinate on particular re–election rules. Suppose that voters take messages from candidates seriously and candidates form correct expectations about voters' responses to their and their opponent's messages. Then, we can focus on stationary equilibria where voting rules depend on the quadruple $(x_t, m^I_{t-1}, m^I_t, m^C_t)$, where m^I_t and m^C_t are the incumbent's and challenger's period–t messages. That is, re–election depends on the official's choice of income for the citizens, her previous–round message m^I_{t-1} (e.g., whether a previous campaign promise was kept), and on both candidates' current messages.

We will focus on (*stationary) credible communication equilibria*. The incumbent will be re–elected if the salary she gives to citizens is compatible with the re–election rule as in Propositions 1 and 2, but she also needs that *either* (i) x_t does not contradict her past promise m^I_{t-1} and her current promise m^I_t gives citizens credible beliefs that future outcomes with her will be *no worse* than those they expect from the challenger given the latter's message m^C_t, *or* (ii) x_t may contradict her past promise m^I_{t-1} but her current promise m^I_t gives citizens credible beliefs that future outcomes with her will be *strictly better* than from the challenger given the latter's message m^C_t.[31] Simply put, an incumbent must keep her past promises and credibly promise *no less* than the challenger, or if she fails to keep her past promises, must credibly promise *more* than the challenger. It is important to underline that these promises must be credible, in the sense that they are compatible with sequential rationality for the official. In particular, given that the official's quality is her private information, any promise must imply a future income for citizens no larger than the maximal income a high–quality official would be willing to provide.

Feltovich and Giovannoni (2015) showed that voters largely take campaign promises at face value and punish officials who do not keep them, thus providing support for equilibria with credible promises. A significant difference from their setting, however, is that we allow for free–form

[31]Ultimately, citizens care about their incomes, so what matters are their beliefs over their future incomes generated by $(x_t, m^I_{t-1}, m^I_t, m^C_t)$. In our experiment, messages are free–form and so may not be explicit about future incomes, but what matters are the inferences citizens can draw about future incomes from such messages. We will return to the issue of message content below.

messages, instead of just asking candidates to put a number in a box, so that whether a promise has been kept or not may be less obvious. Because of this, we allow for the possibility that citizens will re–elect an incumbent that has not lived up to expectations, as long as they believe that they will be better off with the incumbent than with her opponent. Also, because messages have to be written down and their length is up to the candidate, it is reasonable to assume that there is a positive cost c to sending a message. This cost is assumed to be sufficiently small to never deter a candidate from sending a message that would be the difference between winning and losing, but large enough to stop candidates sending messages that would not affect the election outcome.

These additional assumptions make pooling equilibria, where both quality types are re–elected, impossible. To see this, assume to the contrary that a pooling equilibrium does exist. Since challengers always prefer to be elected, they will promise the highest credible amount possible (which is the highest income that a high type can credibly afford), whether or not they could or would deliver it. But low–quality incumbents also prefer to be re–elected, and since being honest would get them replaced, they will instead match their opponents' promises. But once elected, they will not be able to keep such promises, nor will they be able to promise more than their future challengers (at most they can promise the same). So, a low–quality incumbent will be always be replaced next period. This means that only separating equilibria or kleptocratic ones can exist in this setting. In the latter case, campaigns are irrelevant, but in the former we have:

Proposition 4 *Assume c is positive but arbitrarily small. In the game with communication, there exist separating credible communication equilibria where in each period t, conditional on the same campaign messages, citizens re–elect an official if their income x is such that $x^*(\theta^l) < x \leq x^{**}(\theta^l, \theta^h)$, if such income is compatible with the official's message in the previous period and if the official's message in this period implies a credible future income no lower than the challenger's message. The official is also re–elected, regardless of her choice of x, if her campaign message implies a credible future income higher than the challenger's message. In the straw poll, citizens vote in favour of the incumbent if she kept her promise from the previous period and in favour of the challenger otherwise. High–quality officials choose an income equal to $x^{**}(\theta^l, \theta^h)$ in each*

period, while low–quality officials take the entire endowment.

Proof We proceed by backward induction. Recall that any symmetric re–election strategy is an equilibrium strategy for citizens. Consider now the message stage. Since candidates always prefer to be re–elected, any message strategy implying a future per–period income smaller than the highest income than any official can credibly commit to will be dominated by one that promises such an income. Conversely, any message strategy that implies more than such an income will not be credible and will lose the election. Now, let w^j and l^j be the continuation values for a j-quality ($j = H$ or L) candidate in a credible communication equilibrium conditional on winning or losing the election in the current period. Since our equilibria are stationary, these are the same in each period. Given our assumptions on c, it must be that $w^j - l^j > c > 0$. Also, because we are looking at separating equilibria, it must be that $w^H > w^L$ and $l^H = l^L \equiv l$.[32] We consider first the case where the official chooses an income for citizens that guarantees re–election, conditional on the future income implied by her promises being no worse than those of the challenger. In equilibrium the official must be high–quality, so the challenger will put probability one on the official being high–quality. It follows that, effectively, we have a simple subgame where either a candidate sends a message m implying the maximal credible income, or she does not because any other message would be dominated by message m for the same cost. We have two possible challenger types, and assume that the incumbent keeps her promises, as a high type will do. We then have a game of incomplete information, which we can represent in the payoff table below:

	m	$\neg m$
m	$w^H - c; l - c$	$w^H - c; l$
$\neg m$	$l; \omega - c$	$w^H; l$

[32]Recall that an official's quality is unchanged only if she is re–elected. We use w^j and l to describe these continuation values because we are proceeding by backwards induction, but in equilibrium, given Proposition 2, we have $l = \delta V(x^{**})$, $w^L = u\left(n\theta^L\right) + \delta V(x^{**})$ and $w^H = \frac{1}{1-\delta} u(n\theta^H - x^{**}(n-1)) = u\left(n\theta^H\right) + \delta V(x^{**})$.

Here, ω is the continuation value for the challenger upon winning; this notation emphasises that the incumbent knows with probability 1/2 that she is playing with a $\omega{=}w^H$ challenger and with probability 1/2 she is playing with a $\omega{=}w^L$ challenger. As is clear from the table, the incumbent wins unless she does not make promises and the challenger does. It is also clear that no pure–strategy equilibria exist. However, we do have a mixed–strategy equilibrium where a low–quality challenger does not send a message, a high–quality challenger sends a message with probability $\frac{2c}{w^H-l}$, and the incumbent sends a message with probability $1-\frac{c}{w^H-l}$.[33] Suppose now that the incumbent is a low type who deviated and chose an income compatible with her promises. Then, a low–quality challenger sends no message and a high–quality challenger sends a message with probability $\frac{2c}{w^H-l}$ because they still believe they are dealing with a high–quality official. Given this, the low–quality incumbent will send a message for sure and be re–elected.[34] So, as c converges to zero, we converge to an equilibrium where only the incumbent will send the message m.

Consider now the case where the incumbent chooses an income that does not guarantee re–election. In equilibrium such an official must be low–quality, so the challenger will put probability one on this. Again, effectively, we have a simple subgame where either a candidate sends a message m that implies the maximal credible income, or she does not because any other message would be dominated by message m for the same cost. We have two possible challenger types and assume the incumbent is indeed a low type. Then this game can be summarised by the table below:

	m	$\neg m$
m	$l-c;\omega-c$	$w^L-c;l$
$\neg m$	$l;\omega-c$	$l;\omega$

As is clear from the table, the incumbent loses unless she makes a promise and the challenger

[33] We also have a putative equilibrium where the low–quality challenger sends a message with probability $\frac{2c}{w^H-l}-1$, the high–quality challenger does so with probability 1 and the incumbent sends a message with probability $1-\frac{c}{w^L-l}$. But this requires $2c > w^H-l$, so we disregard it.

[34] In the table above, w^L replaces w^H, and since $w^H > w^L$, the incumbent would send the message for sure.

does not. Again, no pure–strategy equilibria exist, but we have a mixed–strategy equilibrium where a high–quality challenger sends a message with probability 1, a low–quality challenger sends a message with probability $1 - \frac{2c}{w^L - l}$, and the incumbent sends a message with probability $\frac{c}{w^L - l}$.[35] Suppose now that the incumbent is a high type who deviated and did not choose an income compatible with her promises. Then, the high–quality challenger sends message m for sure and the low–quality challenger sends a message with probability $1 - \frac{2c}{w^L - l}$, because they still believe they are dealing with a low–quality official. Given this, the low–quality incumbent will not send a message and will be replaced. So, as c converges to zero, we converge to an equilibrium where only the challenger will send the message m.

Then, proceeding as in Proposition 2 completes the proof. In particular, the message strategies described above imply that a incumbent who provides the maximal income will always be re–elected and one that does not will always be replaced, so the maximal credible income will indeed be $x^{**}(\theta^l, \theta^h)$. Therefore, only high–quality officials will provide this income and send messages, while low–quality officials will take the whole endowment and not bother campaigning. ■

It is immediate to see that Hypotheses 1 and 3 continue to hold under the conditions in Proposition 4. Still, the fact that pooling equilibria with credible promises do not exist, and the fact that x^{**} is higher than x^* under the assumptions made to support Hypothesis 1, suggest that messages will increase the expected level of citizen income (Hypothesis 2).

As discussed, the straw poll allows us to disentangle the impact of communication on voting from that of the incumbent's performance. If straw–poll voting is sincere, and if changes from the straw–poll result to the election result differ systematically between the treatments with and without communication, this will be evidence of the direct impact of the chosen communication strategies. Sincere voting in the straw poll means that voters, absent any new information (such as communication from candidates), vote in it in the same way they intend to vote in the election. In

[35] Again, we also have a putative equilibrium where the low–quality challenger does not send a message, the high quality challenger sends a message with probability $2 - \frac{2c}{w^L - l}$ and the incumbent sends a message with probability $\frac{c}{w^H - l}$. This requires $2c > w^L - l$ and we disregard it.

this setting, where all non–candidate citizens have the same preferences and constitute a majority, it is an equilibrium to vote sincerely.

We now consider message content. Candidates can make "campaign promises", but they also can remind voters about their past behaviour or their opponent's; communicate their or their opponent's quality; or even just make small talk. However, our model can only describe messages that may or may not imply returns x^{**} for the citizens, because game–theoretic models of communication cannot generally pin down the precise content of messages. To bridge that gap, note:

Remark 2

- *Incumbents have no information about challenger quality or track record, whereas challengers can discuss their opponent's track record and from that perhaps infer their quality.*

- *As shown in the proof of Proposition 4, incumbents are much more likely than challengers to send messages when they have kept their promises, while the reverse occurs when they did not keep their promises.*

- *For the incumbent, the best way to convey that her promises have not been broken is to emphasise her past performance when that has been satisfactory, and the best way to make up for broken promises is to encourage citizens to look forward to what she might be able to do for them in the future. Thus, we should expect more retrospective campaigning by an incumbent who has complied with her previous promises, and more prospective campaigning when she has not. In both cases, campaigning will be positive because there is no information about the challenger that the incumbent can credibly convey.*

- *Conversely, it is natural for challengers to focus on incumbent past performance when it has been unsatisfactory, and on future performance when the incumbent has performed well. Thus, we should expect more retrospective and negative campaigning by challengers when the incumbent has not complied with her previous promises, and more prospective and positive campaigning when the incumbent did comply.*

In an equilibrium with credible communication, one should therefore expect incumbent performance with respect to her promises to be positively related to the incomes given to citizens, other things being equal. Another simple consequence of this logic is that in the high–variance treatment (where citizens are more likely to expect higher incomes), that threshold at which citizens are satisfied will increase further. These observations lead to Hypotheses 5 and Hypothesis 6.

We can now summarise our theoretical predictions in Table 10.

Table 10: Summary of theoretical predictions

Cell	Comparison	Predicted effect	Basis	Hypotheses
EH	EL	High–quality officials replaced less often	Prop 3	3, 4
		Low–quality officials replaced more often	Prop 3	3, 4
		Higher citizen incomes	Prop 3, Remark 1	1
MH	EL	Higher citizen incomes	Prop 4, Remark 1	2
MH	ML	For given income, challengers more negative and retrospective messages	Prop 4, Remark 1, Remark 2	5
		For given income, incumbents fewer positive and retrospective messages	Prop 4, Remark 1, Remark 2	6
MH and ML	Higher incomes (within –cell)	Challengers more negative and retrospective messages	Prop 4, Remark 1, Remark 2	5
		Incumbents fewer positive and retrospective messages	Prop 4, Remark 1, Remark 2	6

Notes: E = Election–only, M = Messages, H = High–variance, L = Low–variance

A.2 Additional analysis: message characteristics and polling results

Figure 6 displays point estimates (as open or closed circles) and 95–percent confidence intervals (as vertical line segments) of the straw poll's effect on message characteristics, based on the probits in Table 8. The results visible in the table are apparent in this figure. Additionally, we see that the

Figure 6: Incumbent and challenger message characteristics contingent on straw–poll outcomes (point estimates and 95–percent confidence intervals), from Table 8 regressions

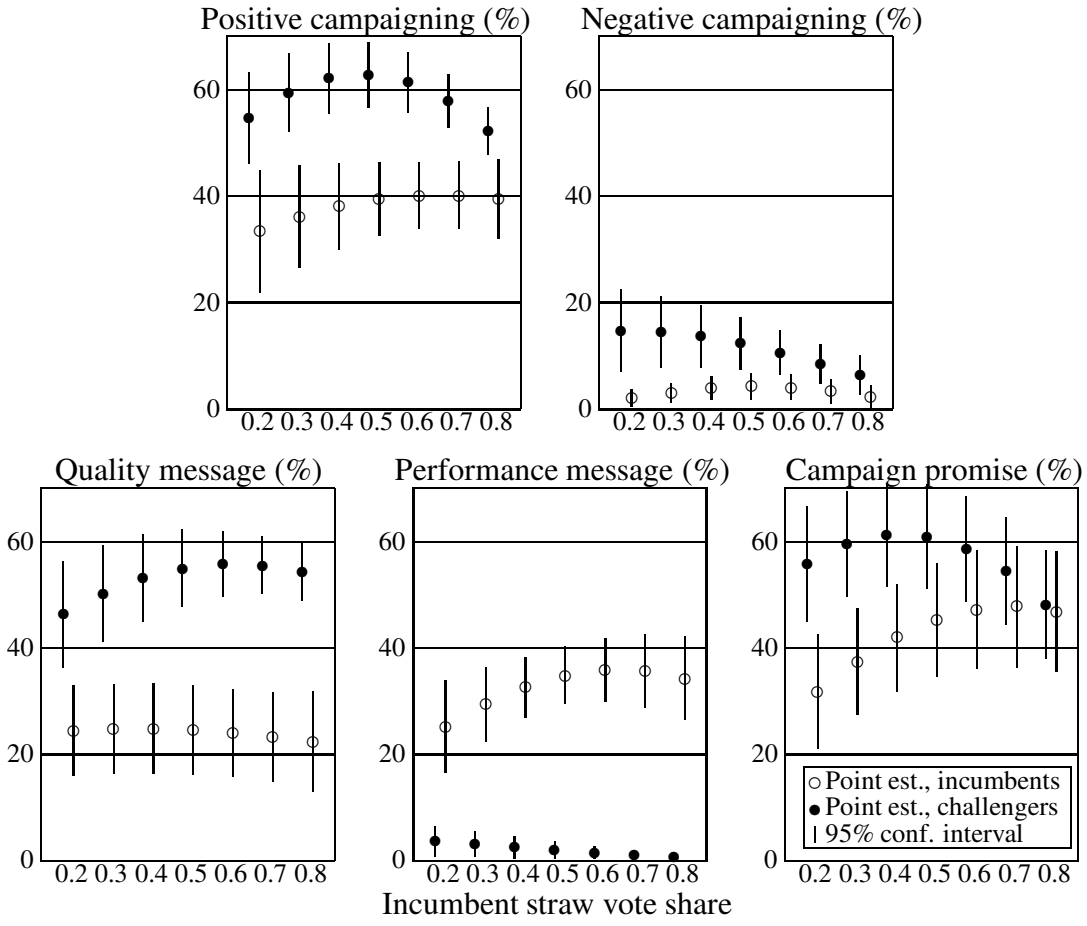

prevalence of a message type is usually highest when the election is expected to be close. The main exceptions are negative campaigning and performance messages by the challenger, both of which are most common when the straw poll favours the challenger, likely reflecting poor performance by the incumbent.

Figure 7 displays marginal effects of messages on the election, conditional on particular straw–

poll results, based on the probits in Table 9. Positive campaigning by either candidate typically

Figure 7: Marginal effects (MEs) of incumbent and challenger messages on election outcomes (point estimates and 95–percent confidence intervals), from Table 9 regressions

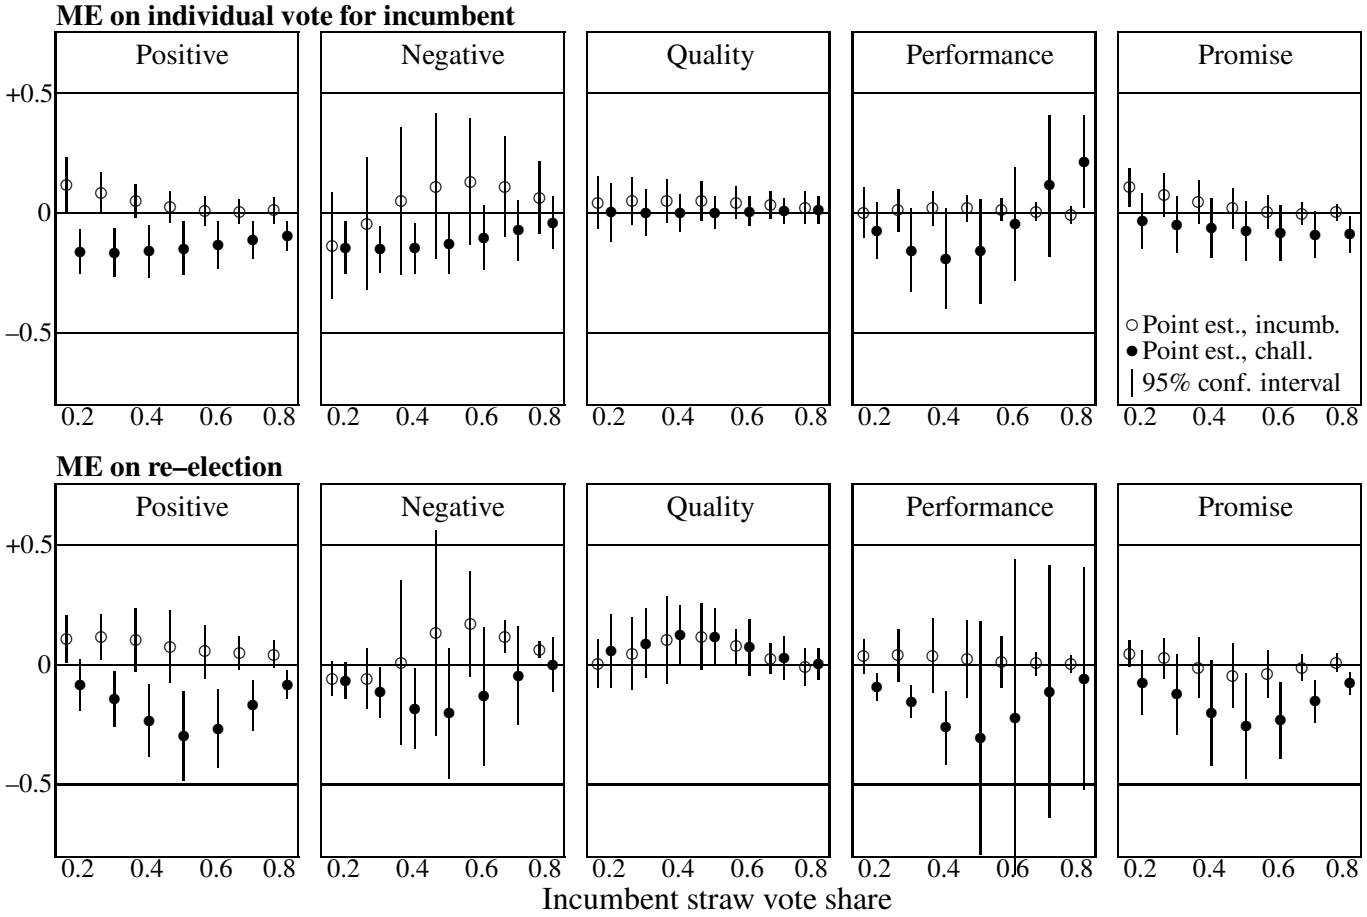

has a positive effect on that candidate's electoral success, though sometimes insignificant for incumbents. Negative campaigning usually benefits the candidate doing it, but only significantly increases the chance of winning when the candidate was likely to win anyway; the same is true for performance messages by the challenger. Finally, campaign promises usually benefit the candidate making the promise in the vote count; those by the incumbent do not significantly affect the outcome, while those by the challenger improve her chances in competitive elections and in those where the incumbent was favoured.

B Sample instructions

Instructions (Part 1)

You are about to participate in a decision making experiment. Please read these instructions carefully. If you have a question at any time, please feel free to ask the experimenter. Please do not talk with the other participants during the experiment, and put away your mobile phones and other devices at this time.

This experiment has two parts. These instructions are for Part 1. After Part 1 is finished, you will receive instructions on your screens for Part 2, which will consist of a new task and a survey.

In Part 1, your decisions will take place within a series of fictional "economies". Each economy lasts for a number of periods, which will be called "years". ***An economy contains the same set of participants for all years of its existence.*** In each year, you can have one of two roles: ***citizen*** or ***official***. There is exactly one official in each economy; everyone else is a citizen. In the economy's first year, the official is chosen randomly by the computer. From the second year on, the official is chosen by an election. When your economy starts, your computer screen will show its population and whether you begin as the official or a citizen. If you are the official, you will also learn your ***quality***. The official's quality is chosen randomly by the computer – equally likely to be either [H cells: ***3*** or ***8***] [L cells: ***5*** or ***6***] – and ***remains the same*** for every year as long as the official stays in office.

At the beginning of a year, your economy receives an ***endowment***: a sum of money to be allocated amongst its population. The amount of the endowment is given by the formula

(Your economy's endowment) = (Your official's quality) × (Population of your economy)

The allocation is made by the ***official***, who chooses how much to take as his/her salary, with the remainder divided evenly amongst the citizens. The official's salary can be any whole number of dollars, between zero and the entire endowment inclusive. After this decision is made, everyone in the economy is informed of their own income, but ***only the official*** knows his/her quality, his/her salary choice, and the economy's endowment.

Elections: After citizens learn their incomes, there is an election between two candidates. One of the citizens is randomly chosen by the computer to be nominated as the ***challenger***, and the official is the ***incumbent***. Each person in the economy (all citizens and the official) ***casts a vote*** for one of the two candidates. ***If the challenger receives more votes than the incumbent, then the challenger wins the election. Otherwise, the incumbent wins the election.***

The challenger is informed of his/her quality – also equally likely to be either [H cells: ***3*** or ***8***] [L cells: ***5*** or ***6***] – as soon as he/she is nominated. No-one else (including the incumbent) is informed of the challenger's quality. If the challenger ***wins*** the election, his/her quality ***stays the same*** for every year as long as he/she stays in office. If the challenger ***loses***, the quality ***is randomly drawn again*** if nominated in a future year.

Straw polls: Before the election, there is a ***straw poll***. Each person in the economy casts a vote for either the incumbent or the challenger. Unlike the election, ***the straw poll is not binding*** – it is just to give information about how the election is likely to go. You should vote in the straw poll the way you are ***intending*** to vote in the election, but you are allowed to change your mind and vote differently in the

election if you like.

[M cells: ***Campaigning:*** After the straw poll results are announced, but before the election, both incumbent and challenger have the opportunity to make ***campaign announcements***. Announcements can be up to 140 characters long, including spaces. The cost of making an announcement is zero. You may choose to make no announcement. If you do make an announcement, it must satisfy these rules:
(1) Announcements must be in ***English***.
(2) You must ***not*** include any ***identifying information*** about yourself (name, age, sex, nationality, …).
(3) You must ***not*** make ***promises or threats*** involving behaviour once this session is finished.
Any announcement satisfying these rules is acceptable. In particular, you may mention information you have that others do not (like your quality, or your salary choice if you are the incumbent), but you are not required to do so.

After both candidates have chosen their announcements (or not to make one), they are seen by everyone in the economy before voting begins. Voters should remember that ***information given in the announcements need not be correct***, and that ***the announcements are not binding*** – they do not restrict the election winner's future decisions.]

Continuing or ending the economy: The number of years an economy exists is determined randomly. At the end of each year, there is a ***20%*** chance that the economy ends, and an ***80%*** chance that it continues for at least another year. If the economy ends, then a new one begins, with a ***new*** – ***randomly chosen*** – set of participants. If the economy continues to the next year, it keeps the same group of participants, with the election winner becoming the official, and the election loser becoming a regular citizen.

Since the number of years an economy lasts is determined randomly, it is possible that some economies will be long and others short. Also, the total number of economies you will take part in is uncertain. It is likely that there will be 5-7 economies in total, but there might be more or fewer.

Sequence of events: The sequence of events in a year is as follows.
(1) Your computer screen displays your economy's population and whether you are the official. The official in your economy chooses his/her salary.
(2) Citizens are informed of their income for the year.
(3) One of the citizens in your economy is designated as the challenger and informed of his/her quality.
(4) There is a straw poll. Everyone votes for either the incumbent or the challenger.
(5) The results of the straw poll are announced: the number of votes for each candidate, and the winner. [M cells: The incumbent and challenger each choose their campaign announcements, or not to make an announcement.]
(6) [M cells: Your computer screen displays both candidates' announcements.] Everyone votes for the incumbent or the challenger.
(7) The computer screen displays the number of votes for each candidate, and the election winner.
(8) The year ends. The computer randomly determines whether the economy ends or continues.

Payments: At the end of each economy, your income from the ***final*** year will be taken as your earnings from that economy. ***So, your total earnings will include the sum of your incomes from the final year of each economy***. Any money you earn in Part 2 of the experiment will be added to your earnings from Part 1. Finally, you will receive a small additional random payment (equally likely to be 2, 3, 4, 5 or 6 cents), and your total will be rounded to the nearest $0.05. Payments are made privately and in cash at the end of the session.