

Does electoral violence affect vote choice and willingness to vote? Conjoint analysis of a vignette experiment

Roxana Gutiérrez-Romero

School of Business and Management, Queen Mary University of London

Adrienne LeBas

Department of Government, American University

Abstract

Across many new democracies, voters routinely elect candidates associated with violence. Though electoral violence is common, there is little understanding of how it affects voting behaviour. This article examines how electoral violence affects turnout and vote choice. To this end, a vignette experiment is set in a nationally representative survey in Kenya, where electoral violence has been present since the 1990s. In the experiment, voters choose between two rival politicians. The experiment randomises candidates' attributes, their rumoured use of electoral violence, and their record of reducing poverty. Conjoint analysis is used to isolate the effects of the candidates' randomised attributes on turnout and vote choice. In contrast to the assumptions made in the literature on electoral violence, voters are less likely to vote for candidates rumoured to have used electoral violence, even when the candidate is a coethnic or a copartisan. This sanctioning effect, however, is not consistent across all voters. Victims of electoral violence and the poorest respondents are less likely to sanction candidates rumoured to have used violence, especially when these candidates have a good record of reducing poverty. The results show that voting turnout decreases when participants are asked to choose between candidates who are rumoured to have used electoral violence. These results are robust to including respondent and interviewer characteristics that might have affected participation in the experiment and how respondents voted. These findings explain why candidates using violence can win elections and why electoral violence has been difficult to eradicate in settings characterised by clientelism and instances of political discourse justifying the use of violence.

Keywords: Electoral violence, voting, turnout, ethnicity, Kenya.

Corresponding author: r.gutierrez@qmul.ac.uk

Introduction

Electoral violence affects many new democracies. An estimated 60% of African elections were marred by violence between 1990 and 2008 (Straus & Taylor, 2013). Even when electoral violence has not resulted in regime change, it has caused thousands of casualties, population displacement, and protracted political crises, as in Kenya, Nigeria, and Zimbabwe. Electoral violence has also imposed substantial economic costs by displacing workers, disrupting transport networks, exports and imports, and increasing food prices (Dercon & Gutiérrez-Romero, 2012). Despite these costs, voters in many new democracies routinely elect candidates who are associated with election violence or who have criminal pasts (Banerjee et al., 2014).

In Kenya, for instance, President Uhuru Kenyatta and Deputy President William Ruto were elected in 2013 with the overwhelming support of their coethnics, despite being under indictment by the International Criminal Court (ICC). Both candidates were accused of orchestrating the violence around the 2007 election, which claimed twelve hundred lives and displaced a quarter of a million people. India, the world's biggest democracy, is another case where voters often elect politicians facing criminal charges (Vaishnav, 2017). More broadly, using data on 786 elections in 155 countries, Collier & Hoeffler (2009) show that politicians who use illegal electoral practices –including electoral violence – remain in office three times longer than their clean counterparts. Why do politicians with violent or criminal reputations get elected?

The literature on new democracies has argued that voters might be willing to vote for violent candidates with whom they share affective ties, such as ethnicity or partisan affiliation, expecting that these candidates might bring benefits if elected (Chandra, 2004). The theoretical literature argues that politicians who use pre-electoral violence often aim at weakening opponents' support (Collier & Vicente, 2012). Violence is assumed by this literature not to trigger sanctioning by the violent politicians' core constituents, even if these voters may disapprove of the use of violence in the abstract (de Figueredo & Weingast, 1999; Chaturvedi, 2005; Collier & Vicente, 2012). The assumption of 'electorally costless' violence when it comes to core supporters is also reflected in empirical studies. For instance, Hafner-Burton, Hyde & Jablonski (2014) argue that incumbents' use of pre-electoral violence is effective in reducing opposition support but ignore any effect on the incumbent support. In fact, Wilkinson & Haid (2009) and Lynch (2014) suggest violence may boost ethnic solidarity and increase the support of core constituencies for politicians who use this strategy. Despite the significance of these issues, there is limited evidence on the extent to which voters sanction candidates that use electoral violence.

This article contributes to the empirical literature on electoral malpractice by testing the assumptions commonly made by theories on the use of violence. First, is violence 'electorally costless' for politicians when it comes to their core supporters –coethnic or copartisan– or does rumoured involvement in violence reduce support even among these voters? Second, are poor or

previously victimised voters more or less likely to support candidates rumoured to have used electoral violence? These vulnerable groups may be more susceptible to clientelism and to politicians' justifications of the use of violence, so we might expect them to be less likely to punish politicians for violence.

Since electoral violence occurs amid other irregularities, such as rigging and vote-buying, it is difficult to isolate its effects on voting behaviour using observational data. To overcome this challenge, we use a vignette experiment set in a nationally representative survey in Kenya. Specifically, the experiment presents each respondent with two rival candidates and randomly varies each candidate's rumoured use of electoral violence and their record of reducing poverty, which we view as signalling some of the benefits the candidate might deliver if elected. This randomisation allows us to assess these cues as well as whether respondents' affective ties, such as ethnicity or partisan affiliation, affect vote choice. Finally, respondents also have the choice of refusing to vote, allowing us to test whether rumours of electoral violence affect turnout and vote choice.

The experiment offers three key findings. First, using conjoint analysis, we find that core and non-core constituents are less likely to vote for candidates rumoured to have used electoral violence. This sanctioning is not sufficient to guarantee electoral defeat, however, since the violent candidate might still get sufficient support to win the (imaginary) election if his coethnic and copartisan support base is sufficiently large. One reason for this second finding is that victims of electoral violence and poor voters compose a considerable percentage of the population, and we find that these voters are less likely to sanction candidates rumoured to have used electoral violence. This is especially the case when candidates are reported to have a good record of reducing poverty. Vulnerable populations, it seems, prioritise other factors when evaluating violent politicians. These performance-based trade-offs have not previously been discussed in the literature on vote choice and violence. Third, we find that respondents are more likely to decline to cast a ballot when they face a choice between candidates rumoured to have used electoral violence. This suggests that perceived violations of democratic norms dampen willingness to vote in the absence of potential victimisation.

The article makes significant contributions to the emerging literature on electoral violence. Our findings shed light on why violent politicians may continue to win elections, even as they face the loss core constituents due to their use of violence. In the next sections, we develop our hypotheses, discuss the country setting, the experiment, and results. The last section presents our conclusions.

Mobilising voters

In this section, we develop hypotheses about how voters might react to electoral violence in settings characterised by clientelism and by the use of violence and intimidation by electoral candidates.

The carrot: Clientelism

In Kenya, as in many other African countries, there is clear evidence of identity-based voting, insofar as individuals tend to favour coethnics and ethnic groups tend to coordinate their votes around one political party. The preference for coethnics is partly due to the low availability of information about candidate quality, but it also reflects voters' expectations of more benefits from coethnic candidates if elected (Posner, 2005). Voters are often mobilised through the promised distribution of resources to their ethnic group, rather than broader reaching policies (Wantchekon, 2003). Incumbent politicians also tend to reward their coethnics and their areas of origin with more investments and services (Franck & Rainer, 2012). Ethnicity also helps parties to build alliances with other groups when reliance on a single ethnic group would yield insufficient votes to win an election. These mobilisation tactics illustrate why voters' ethnic identity and partisan affiliation intertwine and why much of African politics remains clientelist (Isaksson & Bigsten, 2017).

One concern with this clientelism is that voters might prioritise ethnic loyalties and fail to sanction their politicians for poor performance or violations of democratic norms. This concern is seen across studies that suggest ethnic diversity explains Africa's dismal economic performance and political instability (Easterly & Levine, 1997). Nonetheless, in countries where ethnicity strongly predicts voting, voters downplay concerns that ethnicity is driving their voting choices (Carlson, 2016). Instead, voters claim to focus on candidates' qualities such as performance and honesty (Bratton & Kimenyi, 2008). Indeed, some studies have found that African voters do hold their politicians accountable, such as in Kenya where MPs perceived to have abused local anti-poverty projects were voted out in elections (Gutiérrez-Romero, 2013). Bratton, Bhavnani & Chen (2012) provide broader cross-country evidence, suggesting that African voters care about ethnicity but also weigh their economic interests when making vote choices.

Experimental studies in Africa provide further evidence that voters favour coethnic politicians but are also driven by performance cues. For example, Conroy-Krutz (2013) finds that 80% of respondents in Uganda voted for their coethnic when faced with a choice between two candidates in which no information other than ethnicity was revealed. Reliance on ethnicity decreased, however, when voters had to choose between coethnics who had negative characteristics and non-coethnics with more positive ones. As in Ichino & Nathan (2013), Conroy-Krutz finds that members of smaller ethnic groups are more likely to support a non-coethnic candidate, suggesting that voters behave strategically to maximize their expected clientelistic benefits.

The stick: Electoral violence

Across many new democracies, electoral violence is used alongside clientelism to influence vote choice. Electoral violence can be understood as any coercive force, including assassination, physical harm, forced displacement, intimidation, and disruption of the ballot or the count (Bekoe, 2012). According to the theoretical literature, politicians often use pre-electoral violence to reduce the

turnout of rivals' supporters, presumably without significantly affecting the voting support and turnout of their core voters (de Figueredo & Weingast, 1999; Chaturvedi, 2005; Collier & Vicente, 2012).

Empirical studies have found support for some of these assumptions, but other aspects of vote choice have not been fully examined. Globally, electoral violence is more likely to occur before the election when incumbents face a credible threat of losing (Hafner-Burton, Hyde & Jablonski, 2014); in Africa, most electoral violence occurs prior to the election and is associated with the incumbent party (Straus & Taylor, 2013). Perhaps unsurprisingly, weaker partisans and individuals who lack an affective tie to election candidates report higher fears of victimisation and lower (self-reported) turnout than do coethnic supporters or non-swing voters (Gutiérrez-Romero, 2014; Wallsworth, 2016). Field experiments and observational studies in Africa also corroborate that electoral violence dampens turnout (Collier & Vicente, 2014).¹ These results coincide with studies inside and outside of Africa that find that turnout is affected by voters' fear or anticipation of disruption at the polls and by their disillusionment with the legitimacy of the electoral process (Birch, 2010; von Borzyskowski & Kuhn 2020; Young 2020).

Experimental studies have also investigated how candidates' alleged criminality affects vote choice. For instance, Banerjee et al. (2014) find that Indian voters, regardless of their ethnicity and political knowledge, prefer honest non-coethnic candidates free of criminal records over corrupt coethnic candidates. Similarly, survey experiments in Kenya and Pakistan find that voters, regardless of their ethnicity or whether they had been victims of electoral violence, reduce their support for politicians known to use electoral violence (LeBas, 2010; Blair et al., 2013; Rosenzweig, 2015). This sanctioning behaviour is also consistent with voters' self-reported rejection of the use of electoral violence when asked directly, as with other violations of democratic norms such as vote-buying and corruption. For instance, according to the third round of the Afrobarometer, about 80% of those surveyed in 18 countries state that the use of political violence is 'never justified'. Nonetheless, these beliefs seem inconsistent with the electoral success of politicians accused of using violence, as the 2013 Kenyan elections demonstrate. To reconcile these findings, the literature has argued that voters may decline to sanction coethnic or copartisan candidates for criminality because they value the access to clientelism that comes with election of coethnics. That is, voters may be willing to 'trade-off' or be accepting of negative signals if expected benefits of a candidate's election are sufficiently large (Winters & Weitz-Shapiro, 2013; Vaishnav, 2017).

Some kinds of voters are more likely to engage in this kind of trade-off. First, voters with historical grievances or past experience of violence are more likely to find politicians' justifications for the use of violence plausible. For instance, in Kenya, as in much of Sub-Saharan Africa, electoral

¹ Voting participation in Africa, as in other regions, is also driven by voters' partisan affiliation, education, age, residency and wealth (Kuenzi & Lambright, 2011).

violence has often been intertwined with ethnic tensions over land, which politicians often exploit (Boone, 2011). Politicians also argue that violence is sometimes necessary to safeguard communities against rivals (Lynch, 2014). These considerations help explain why Kenyans who were victims of violence in the 2007 election period were more likely to perceive violence as acceptable in support of a just cause than non-victims of violence (Gutiérrez-Romero 2014). Drawing on a survey-embedded experiment in the Rift Valley, the epicentre of violence in Kenya, Horowitz & Klaus (2018) find that appeals to ethnic grievances over land were generally unsuccessful – *except* among those respondents who were land insecure and indigenous to the province. Though Horowitz and Klaus’s hypothetical candidates do not state that they would use violence to resolve ethnic land grievances, their results are suggestive of differential response to ethnic appeals in these settings.

Second, we expect that poor voters are less likely to sanction violent candidates from whom they expect clientelistic benefits. Poor voters are the core target audience for clientelistic mobilization (Weitz-Shapiro, 2012; Jensen & Justesen, 2014). This may be because they need these benefits more than wealthier voters, or they may be more risk-averse, preferring local or private benefits today over uncertain future rewards (Stokes, 2009). For instance, Nathan (2016) finds that poorer urban voters in Ghana prefer targeted patronage benefits over universal public policies, in contrast to wealthier voters.² Although the poor might also disapprove of the use of violence, they might face countervailing pressure to vote for a politician offering potential benefits. These voters would be more likely to discount rumoured violence, we believe, if they receive countervailing information about good performance in office or when considering a vote for a coethnic candidate.

Based on this discussion, we test the four hypotheses listed below about how candidates’ attributes might affect voting behaviour by conducting a conjoint experiment. Compared to traditional surveys, conjoint experiments have the advantage of reducing social desirability response bias by presenting candidates with multiple attributes that may guide their voting choices, allowing us to directly observe the effects of individual attributes and the trade-offs that voters make (Hainmueller, Hopkins & Yamamoto, 2014).

In our experiment, each respondent is presented with two rival candidates mentioning only the candidates’ names and partisan affiliation. Respondents are then asked to cast votes for their preferred candidate in two sequential ballots, which allow us to test the following hypotheses.

Hypothesis 1: In the absence of any information about the candidates other than their name (conveying ethnicity) and partisan affiliation, voters are more likely to vote for a coethnic or copartisan candidate over other candidates.

² Consistent with our expectations, Kramon (2013) presents experimental evidence that poor Kenyan voters are more responsive to coethnic politicians’ vote-buying than non-poor and non-coethnic voters.

The experiment then provides respondents with additional randomised information about the same pair of candidates regarding their use of electoral violence and past performance when in office. Respondents are asked to cast a second ballot for their preferred candidate. Unlike previous experiments on electoral violence, we simultaneously provide negative (rumoured use of electoral violence) and positive (good record of reducing poverty) candidates' attributes to more closely mimic the trade-offs that voters might make during elections. As discussed in the next section, Kenyan voters experience various forms of electoral violence, ranging from politicians using incendiary rhetoric against some groups to hiring gangs to intimidate opponents' supporters. Since there is a high degree of impunity for these kinds of electoral crimes in Kenya, our script focuses on the effect of candidates being rumoured to have ordered killings and hired gangs during elections. We then test the following hypotheses.

Hypothesis 2: After learning about the candidate's rumoured use of electoral violence, voters still are more likely to vote for a violent coethnic or copartisan candidate over other candidates.

Hypothesis 3: Victims of electoral violence and poorer voters are less likely to sanction candidates rumoured to have used electoral violence since they may be more susceptible to politicians' mobilisation tactics, albeit due to different theoretical logics.

In Kenya, the use of violence has been justified as group defence and as a means of redressing ethnic grievances or past injustices. Thus, we expect that victims of past electoral violence would be less likely to sanction a candidate who is rumoured to have used electoral violence but who also has a good record of using development funds well in the past. We also expect poorer voters to make such trade-offs due to potential clientelist benefits of electing the candidate.

Hypothesis 4: The experiment's voting turnout decreases when candidates are rumoured to have used electoral violence.

We do not expect that fear of victimisation will affect turnout since all our respondents are aware that they are participating in a game. Still, our experiment's turnout could be affected if voters sanction candidates for violations of democratic norms such as the use of electoral violence.

Setting

Like other countries in the developing world, Kenya has experienced several episodes of electoral violence and is characterized by ethnic tensions that partly result from government bias in the allocation of jobs and public goods (Adar & Munyae, 2001). From independence in 1963 until its electoral defeat in 2002, the Kenya African National Union (KANU) ruled the country. In the early 1990s, due to pressure from domestic actors and foreign donors, Kenya returned to multiparty electoral competition. However, the 1992 and 1997 elections resulted in forced displacement and

state-sponsored violence undertaken by the hiring of gangs and ethnic militia, particularly in the Rift Valley. Political discourse emerged justifying violence as a legitimate form of ethnic self-defence (Human Rights Watch, 1993).

In 2002, KANU named Uhuru Kenyatta as its candidate, but he was defeated by a fellow Kikuyu, Mwai Kibaki, who was backed by a broad multi-ethnic coalition of opposition parties. This peaceful electoral turnover did not end the use of electoral violence by candidates. Before the 2007 elections, the coalition that elected President Kibaki unravelled, and Kibaki ran against his former ally, prominent Luo politician Raila Odinga. Support for the two candidates was split along ethnic lines: Kikuyus supported their coethnic Kibaki, as did most northerners, while Luos, Kalenjins, and western Kenyans lined up behind Odinga. In a nationally representative survey conducted before the polls, Dercon & Gutiérrez-Romero (2012) found that a third of respondents experienced electoral intimidation directly and witnessed politically-connected gangs intimidating other voters.

After Kibaki was announced as the winner, Odinga claimed the election was rigged. Political protest over the election results gave way to large-scale violence. By the time Kibaki and Odinga signed a power-sharing agreement in February 2008, an estimated twelve hundred people had been killed, and a quarter of million people had been displaced. The International Criminal Court (ICC) issued several indictments based on evidence that implicated prominent political figures in the organization of both pre- and post-electoral violence. Among those indicted were Uhuru Kenyatta and William Ruto, who belonged to rival parties at the time.³ They later formed a partisan alliance in the run-up to the 2013 general elections. Opposition challenger Raila Odinga ran against Kenyatta for the presidency, again drawing on the support of his Luo coethnics and other western Kenyans.

The pending ICC cases shaped much of the debate of the 2013 campaign, and Kenyatta and Ruto argued their indictments were unjust punishment for ‘protecting’ their communities from rivals (Lynch, 2014). Rather than costing Kenyatta and Ruto votes, the ICC indictments seem to have solidified the support of each candidate’s coethnic voters. Kenyatta won the presidential election with 50.5% of the vote. The Kenyatta-Ruto alliance won 48% of the National Assembly seats against Odinga’s alliance’s 40%. Even though the campaign was characterized by violence, resulting in the loss of as many as 500 lives (Human Rights Watch, 2013), 86% of registered voters cast their votes, the highest turnout on record in Kenya.

Despite the economic and human costs of electoral violence, Kenyans have complicated views about its use and how to assign blame. In the survey conducted for this article, described in the next section, a quarter of respondents believe politicians must intimidate their opponents and hire gangs to win elections. However, when asked about the credibility of the allegations of violence

³ Ruto was Odinga’s ally in the 2007 election, explaining why most of Ruto’s Kalenjins coethnics voted for Odinga. When Ruto defected to Kenyatta’s ticket in 2013, he took the Kalenjini vote with him.

against Kenyatta, Ruto, and Odinga, respondents discount allegations of violence if it involves their coethnic or copartisan politicians, but are more likely to see rival politicians as responsible for the 2017–2018 violence (Figure A.1 in Online appendix A). These perceptions are not surprising since most respondents (55%) also believe that rivals spread rumours about candidates' use of violence to divide the electorate.

The experiment

The experiment was embedded in a nationally representative survey in Kenya in December 2013, nine months after the general election. The survey consists of 1,210 face-to-face interviews with voting-age adults who were randomly selected from 80 out of 290 constituencies across the country. Online appendix B discusses the sampling in detail, and respondents' main characteristics are reported in Table A.1.

At the beginning of the experiment, all respondents were asked to imagine there was a vacancy for a Member of Parliament in their constituency with only two candidates contesting the seat. To equalise the candidates' qualifications and promises, respondents were told:

‘Both candidates have equal experience in politics, and both have been elected MPs before but in other constituencies similar to this one during the 2008–2013 administration. Both candidates are promising to improve the economy of your community. They are:

Candidate 1 is John Onyango, and he is running for the ODM party.

Candidate 2 is John Kamau, and he is running for the TNA party.’

To reduce the possibility of socially desirable responses we did not mention explicitly the candidates' ethnicity. Onyango and Kamau are common surnames associated with the Luo and Kikuyu ethnic groups, respectively, and would be recognised as such by most Kenyans.⁴

Once the two candidates were introduced, respondents were given a voting ballot, which included the two contending candidates' names and the name and official logo of their respective parties (Figure B.1). Respondents were asked to cast their votes for their preferred candidate by placing their marked folded ballots in a transparent bag containing other ballots. They were reassured that no-one would reveal their name and vote choice.⁵ The vote choice in this ballot helps us assess whether respondents prefer a candidate with whom they have an affective tie in the absence of any

⁴ The ODM and TNA parties ran presidential candidates in 2013 who were Luo and Kikuyu, hence our cues are consistent. Since ethnic and partisan identities overlap, there is no theoretical payoff to priming them separately.

⁵ The ballots were individually coded so they could later be matched to their respective questionnaires. However, the voting was secret in the sense that no other participant or interviewer saw how a respondent voted.

other information. Responses are unlikely to be driven by social desirability bias, so we expected a strong correlation between respondents' ethnicity and their vote choice in this ballot.

After voters cast this first ballot, we provided further information about the same pair of candidates, and respondents were then asked to cast a second ballot for their preferred candidate. The randomisation of candidates' attributes was done so that voters might encounter candidates who are not ideal on every possible aspect, as occurs in real elections. Candidates might have some appealing attributes but also some undesirable ones. Thus, when voters are asked to choose between candidates, they need to prioritise some attributes over others.

Specifically, we simultaneously provided respondents two randomised attributes on each candidate: performance record (intended to signal the potential rewards from electing the candidate) and use of electoral violence.⁶ Each of these attributes has three levels, which are summarised in Table 1. The experiment did not include all the possible combinations of the levels that each of the candidates' attributes can have, as we excluded candidates who are not theoretically interesting, such as those possessing only positive characteristics. Table II presents the nine candidate pairings that serve as 'treatments' to which respondents are randomly assigned. Each participant was randomly assigned to only one of these nine 'treatments'. As shown in Figure A.2, respondent characteristics are well-balanced across treatments. The full experiment script and ethical considerations involved can be found in Online appendix B.

Table I

In terms of the candidate's violence attribute, respondents were either provided with no information concerning the candidate's use of electoral violence, were told that the candidate 'has never used electoral violence',⁷ or were told that the candidate was rumoured to have hired gangs and ordered a murder during the 2007 elections but had not been arrested. This last cue places the candidate's rumoured use of violence in a strategic pre-electoral context. Most Kenyans would recognise these hired gangs as a common form of electoral intimidation, since gangs have been used as part of state-sponsored electoral violence and have also been employed by the opposition (Human Rights Watch, 1993; Dercon & Gutiérrez-Romero, 2012). We chose not to elaborate on the reasons for the rumoured use of electoral violence –whether used as self-defence or in retaliation– since voters hear such rumours without knowing the motive or the degree of involvement of particular politicians.

⁶ We randomised by asking each respondent to select a number between one and nine. As described in Online appendix B, the pilot of the experiment suggested this was the easiest way to randomise.

⁷ We categorically stated that the candidate 'has never used violence,' which we believe most voters will perceive as a positive cue. Had we presented all the positive and the negative cues on violence as merely rumours, we would have risked rendering these cues as noninformative.

Nevertheless, our experiment supplies a context to aid respondents in assessing and interpreting the rumours of electoral violence by providing information on the candidate's ethnicity, partisan alliance, and performance record. Our analysis also considers respondent characteristics, including whether they have been victims of electoral violence and their area of residence.

In terms of the candidate's performance attribute, respondents were provided with no information concerning the candidate's performance, were told that the candidate's performance record was unknown, or were told that the candidate had a good performance record of reducing poverty. To signal good performance, we provided information about the candidate's past use of a programme known as the Constituency Development Fund (CDF). Through this programme, regarded as Kenya's main poverty reduction strategy, MPs receive discretionary funds to sponsor projects in their constituencies.⁸ The CDF has a reputation for being subject to abuse, which has been documented in audits. Thus, respondents receiving a positive performance cue were told:

‘Back when [the candidate] was an elected MP, he was among the few MPs that were randomly and independently audited... The audit found that he spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.’

As Table II shows, for treatments in which one candidate was assigned a positive performance cue, respondents were told that for the other contending candidate, ‘we do not have any information about how this candidate used the CDF ... because he, like the majority of other MPs, was not audited.’ This scenario is realistic, as not all MPs have been audited and ensures that a neutral cue is not interpreted as a negative one.

Table II

Analysis and discussion

We start by analysing the votes cast in the first ballot, where voters were only provided with candidates' names and partisan affiliations. As Table A.2 shows, there is a strong correlation between respondents' ethnicity and their vote choice, supporting our first hypothesis. Specifically, of those who cast a valid ballot and had the opportunity to vote for a coethnic candidate (the Kikuyu and Luo), 87% voted for their coethnic candidate. For the rest of respondents who did not have a coethnic candidate in the imaginary election, over 70% voted for their ‘ally candidate’, that is, the candidate

⁸ Gutiérrez-Romero (2013) finds that 40% of Kenyans who voted to re-elect their MP did so primarily because the MP managed the CDF well.

affiliated with the political party that the respondents' ethnic group supported in the 2013 presidential election.⁹ This is consistent with our starting expectations.

Conjoint analysis of the two ballots

After the first vote was cast, each respondent was randomly assigned to one out of the nine scripts that provided further information about the candidates' performance record and use of electoral violence. Respondents were then asked to cast a second ballot. We analyse together the vote choice made in these two ballots using the conjoint approach proposed by Hainmueller, Hopkins & Yamamoto (2014). This method allows us to estimate the net causal effect of each of the candidate's attributes on vote choice, known as the *average marginal component effect* (AMCE), and any interaction effect between linked attributes.

These casual effects are obtained by estimating an Ordinary Least Squares (OLS) regression. As shown in Equation (1), the dependent variable, the vote choice made in the two ballots is regressed on a set of dummy variables indicating the various levels of each of the candidates' attributes (omitting the respective reference categories) and the interaction terms between attributes. The AMCEs for each of the attribute's levels is obtained as the weighted average of the regression coefficient of the attribute's level and its interaction regression coefficients with other linked attributes. In experiments that do not fully develop all the potential combination of the various levels of attributes, like ours, the AMCE excludes interactions between attributes that do not exist or effects that are deemed implausible. In Online appendix C, we explain further how we obtain these AMCEs.

The candidate attributes presented in the first ballot serve as the baseline, where respondents vote without knowing about candidates' performance records or whether the candidates had used electoral violence. Hence, the AMCE serves as a difference-in-means that measures the change in the probability of selecting a candidate with a given attribute value (e.g. those shown in the second ballot) from the baseline probability of selecting that candidate. To account for a possible non-independence of choices, we cluster the standard errors at the respondent level.

$$output_{ijk} = \beta_0 + \beta_1 I_{ijk} + \beta_2 V_{ijk} + \beta_3 P_{ijk} + \beta_4 V_{ijk} * P_{ijk} + \beta_5 X_{ijk} + \varepsilon_{ijk} \quad (1)$$

where *output* represents whether the respondent voted for a particular candidate in each of the two ballots or not. We have up to 4,840 observations since each of the 1,210 respondents *i* was asked to cast two separate ballots (*k* choice tasks) selecting the most preferred of the *j* alternative candidates

⁹ Table A.2 also shows the voting in the first ballot is remarkably similar to how respondents claimed to have voted in previous presidential elections.

presented [1,200 x 2 x 2].¹⁰ I is a dummy variable that indicates whether the candidate is Onyango/ODM party (the reference category) or Kamau/TNA party. V and P represent the candidates' attributes of violence and performance record; each attribute has three levels as shown in Table I. We exclude the first level of these variables, which serve as the reference categories. We also add the interactions between the linked attributes V and P . ε denotes the error.

We estimate Equation (1) with and without a set of controls represented by vector X . This vector is added to estimate the standard errors more efficiently and contains respondents' ethnicity, sex, age, education, province of residence, rural residence, whether the respondent feels close to a party, the respondent's experience of food shortages within the past year, and experience of electoral violence since 1992. We also include the interviewers' mother tongue (linked to ethnicity) and years of experience in the field, since these characteristics may shape respondents' choices (Carlson, 2016). X also includes the district index of ethnolinguistic fractionalisation proposed by Alesina et al. (2003) and the polarisation index proposed by Montalvo & Reynal-Querol (2003), since these indices might influence vote choice (Isaksson & Bigsten, 2017).¹¹ The regression coefficients of Equation (1) are presented in Table A.3.

We graphically present only the AMCE of the candidate's rumoured use of violence and performance record. These AMCEs represent the change in probability of a respondent selecting a candidate when a given attribute value is compared to the attribute considered in the baseline. The dots in Figure 1 denote the AMCEs estimates, the lines illustrate the 95% confidence intervals, and the dots without the lines denote the baseline categories. Figure 1 also shows the interaction between the rumoured use of electoral violence and good performance. This interaction measures the percentage point difference between the AMCE of a candidate with a good performance record who was also rumoured to have used electoral violence and the AMCE of that same candidate when his performance is not mentioned to respondents. Throughout our analysis, we consider the results of all respondents together. Separately, we also show those who had a coethnic candidate in the election (the Kikuyu and Luo) and those without a coethnic candidate.¹²

¹⁰ Some respondents chose not to vote, or invalidated their ballot by voting for both candidates or submitting a blank ballot. We exclude these three cases from our choice analysis, leaving us with 3,724 valid observations. As shown in the section on turnout, our results are not affected by excluding respondents that did not cast a valid vote.

¹¹ Since there are no official ethnic statistics at the district or lower administrative level, we estimate these district-level indices using our survey data on respondents' mother tongue, which largely denotes individuals' ethnicities.

¹² Figure A.3 shows the same conclusions are reached if the sample is divided according to the respondents' ethnicity and the party they claimed to have voted for in the 2013 presidential election.

Electoral violence

Compared to a candidate for whom there is no information on whether he has used electoral violence, a candidate rumoured to have used electoral violence experiences a decline in support by nearly nine percentage points. Figure 1 shows that this sanctioning is nearly identical if analysing all the respondents together, analysing separately those with and without a coethnic candidate, or adding X controls.

Figure 1

Figure 1 also shows candidates with a good performance record lose voting support if they are rumoured to have used electoral violence (about six percentage point decrease). This sanctioning effect is present if analysing all respondents together or if examining those without a coethnic candidate. However, importantly, for the Kikuyu and Luo who can cast votes for coethnics, the good performance record cue seems to offset the effect of rumoured violence, since this interaction effect is statistically insignificant for these respondents.

Since nearly a quarter of our participants believe politicians must use violence to win elections, a relevant question is whether candidates are sanctioned for not using violence. Instead, the results show candidates who have never used electoral violence are about eight percentage points more likely to receive votes than candidates for whom no information is revealed on whether they have used electoral violence. This boost in support is larger than the boost in support for those who are known to have performed well in office.

Bearing in mind that the good performance cue is *always* paired with rumored use of violence, the AMCE effect of having a good performance record is positive but statistically insignificant. In the experiment, we only paired candidates who have never used violence with the cue that offers no information on the candidate's performance. This may explain why this weaker performance cue has a positive and statistically significant effect when analysing both all respondents together and also those without a coethnic candidate.

To better illustrate how the randomised attributes affect vote choice, Figure 2 displays five candidates' profiles shown to our respondents. The y-axis shows the AMCEs of these candidates' attributes over a candidate with baseline attributes. The x-axis shows the candidate's predicted probability of being elected that corresponds to the 1st, 25th, 75th, and 99th percentile of the overall probability distribution of being elected. This figure shows that candidates who have never used violence have a higher probability of being elected, even if their performance is unknown, than candidates who are rumoured to have used electoral violence.

Figure 2

Interaction with respondent characteristics

It is possible that respondents penalised candidates rumoured to have used electoral violence, but only if these candidates were not coethnics (or copartisans). To test the effect of these affective ties more directly, we re-estimate the AMCEs using the specification shown in Equation (2).

$$\begin{aligned} \text{output}_{ijk} = & \delta_0 + \delta_1 I_{ijk} + \delta_2 V_{ijk} + \delta_3 P_{ijk} + \delta_4 R_{ijk} + \delta_5 (V_{ijk} * P_{ijk}) + \delta_6 (V_{ijk} * R_{ijk}) + \delta_7 (P_{ijk} * R_{ijk}) + \delta_8 (V_{ijk} * P_{ijk} * R_{ijk}) \\ & + \delta_9 X_{ijk} + \varepsilon_{ijk} \end{aligned} \quad (2)$$

where R is a dummy variable that takes the value of one either if the respondent has the same ethnicity as the contending candidate or if the candidate is a partisan ‘ally’ of the respondent’s ethnic group. We grouped respondents without a direct coethnic candidate as ‘allies’ of one of the candidates, based on the 2013 presidential election ethnic voting patterns. We determined the preferred presidential candidate of each ethnic group using the 2013 exit poll results by Ferree, Gibson & Long (2014). By this logic, we identified Embu, Kalenjin, Meru, and Pokot respondents as allies of the vignette’s TNA candidate Kamau. We separately coded Kamba, Kisii, Luhya, Maasai, Mijikenda, Somali, Taita, Teso, and Turkana respondents as allies of the ODM candidate Onyango.

Figure 3 shows the candidate’s identity (ethnicity/partisan affiliation) strongly influences vote choice. For example, taking into account the two ballots, the Kikuyu and Luo respondents are 73 percentage points more likely to support their coethnic candidate than a non-coethnic candidate. Similarly, respondents of other ethnicities are 40 percentage points more likely to support their respective ‘ally’ candidate than a ‘non-ally’ candidate.

We find that respondents penalise their respective coethnic / ally candidates by 16 percentage points if rumoured to have used electoral violence, which remains the case whether analysing all respondents together or separately for those having an ‘ally’ candidate. If we analyse the Kikuyu and Luo respondents separately, the effect of their coethnic candidate using violence is negative, but this effect is not statistically significant. Nonetheless, even for these respondents, we find instances where they also penalise their coethnic candidates for using violence. For example, just as with the rest of the respondents, the Kikuyu and Luo penalise their well-performing coethnic candidate by nearly 15 percentage points when he is also rumoured to have used electoral violence. It is unclear why the positive performance cue does not offset the sanctioning effect of the rumoured use of violence, unlike in the previous sub-section. This sanctioning suggests that voters penalise their own coethnic candidates more severely for violations of democratic norms, such as the use of electoral violence, more than they do other candidates.

The sanctioning effect of coethnic candidates using violence is large, but it is not large enough for these candidates to lose the majority of their coethnic vote (i.e. over 50%). Using violence can, however, become ‘electorally costly’ for candidates when we consider the voting reaction of their weaker supporters (those in a partisan alliance). The sanctioning effect of voters on their respective ‘ally’ candidate is large enough for that ally candidate to lose the majority of their votes. Thus, in contrast to the literature’s prediction and our second hypothesis, these results suggest core constituents sanction their own politicians for using violence despite the benefits they might bring.¹³

Figure 3

Victims of electoral violence

We move on to test whether victims of electoral violence penalise violent candidates differently from those not having experienced electoral violence. To assess the effect of the candidates’ attributes, we revert to our original regression specification. We estimate Equation (1) separately for the respondents who claimed to have been victims of electoral violence and those who reported no such experience. The regression coefficients are presented in Table A.4, while Figure 4 shows the effects of interest.

Approximately 24% of respondents claimed to have been affected by electoral violence since 1992, suffering either personal injury, destruction of property, economic loss, or forced displacement. Supporting our third hypothesis, these victims of violence are less responsive to candidates rumoured to have used electoral violence than other respondents. Overall, victims of electoral violence do not penalise candidates rumoured to have used electoral violence, in contrast to non-victims. That is the case whether analysing all respondents together or whether examining the sample separately for the Kikuyu/Luo and the rest of the respondents (Figure 4).

A plausible explanation for these results could be that these respondents perceive candidates’ use of electoral violence as a justifiable means of ‘defense’ against the attacks of others. We find some evidence for this. For example, non-victims of electoral violence increase their support for non-violent candidates by approximately ten percentage points. In contrast, there is no statistically significant effect on victims’ votes for candidates who have never used violence. Similarly, non-victims of electoral violence sanction well-performing candidates if they are rumoured to have used electoral violence, whereas victims of electoral violence do not sanction these well-performing candidates.

¹³ One could argue that some voters might face a ‘psychological cost’ of changing their vote once they have voted for a candidate in the first ballot. Even if such a cost existed, it would reinforce our findings because most of our AMCEs have been statistically significant, showing that voters sanction candidates rumoured to have used electoral violence.

Poverty

To analyse the voting choices of poorer respondents, we categorise respondents as poor if their household has four or fewer assets from a list of 15 durable assets.¹⁴ Based on this measure, 14% of respondents are identified as poor, which correlates well with whether respondents stated in our survey ‘their household to be fairly poor’. We analyse Equation (1) separately for poor and non-poor voters.¹⁵ The regression coefficients are shown in Table A.5 and the main effects in Figure 5.

Poor respondents do not penalise candidates for their rumoured involvement in electoral violence, in contrast to non-poor respondents. That is the case when analysing all the respondents together or whether examining separately the Kikuyu and Luo. These poorer respondents, including the poorer ally voters, also fail to penalise candidates rumoured to have used violence if they have a good performance record, in contrast to non-poor respondents.

Thus, as with victims of electoral violence, poorer respondents seem to be less responsive to electoral violence cues than wealthier respondents. One explanation for this similar behaviour is that poorer respondents are more likely to be victims of electoral violence. Our survey, however, reveals the same percentage of poor and non-poor respondents have been victims of electoral violence (approximately 24%). To test the robustness of our results, we re-ran our analysis, this time excluding victims of electoral violence. The results remain the same. That is, poor respondents who have never been victims of electoral violence fail to sanction candidates rumoured to have used electoral violence, in contrast to non-poor respondents (Figure A.4). Overall, these results support our third hypothesis and our intuition about the effect of poverty on the strength of clientelistic linkage.

Figure 4

Figure 5

Turnout

Taking into account the two ballots, 83% of respondents cast a vote. This turnout is high and remarkably similar to the turnout of the 2013 presidential election. Also, as in real elections, a small minority among those who cast a ballot submitted invalid ballots (with 6% of respondents leaving them blank and 1% selecting both candidates), as shown in Table A.6. Although some respondents did not cast valid ballots, our experiment does not suffer from attrition given that all respondents

¹⁴ The assets were: house, land, cattle, oven, fridge, washing machine, computer, phone, mobile phone, book, radio, television, bicycle, motorcycle, and car.

¹⁵ We also fit a fully interacted model to the pooled data confirming victims of violence and poor voters display different voting behaviour when compared with other respondents.

completed the rest of the questionnaire. Nonetheless, it is possible that some respondents might have sanctioned candidates rumoured to have used violence, choosing not to vote for them or for the alternative candidate, despite facing no risks of victimisation and having zero costs of going to the polls.

To assess whether candidates' attributes affected turnout and test our fourth hypotheses, we estimate regression as in Equation (1). For this analysis, our dependent variable takes the value of one if the respondent refused to cast a valid vote or spoiled a voting ballot and zero otherwise. The vector of controls X is the same as before, except that we add a proxy for 'civic duty to vote' measured by whether the respondent claimed to have voted in the 2013 presidential election.

Figure 6 shows there is no difference in turnout for candidates rumoured to have used electoral violence over those for whom no mention was made about their use of electoral violence. When we compare a candidate rumoured to have used electoral violence to one who is explicitly reported to have never used violence, the turnout of allies is reduced by six percentage points. It is possible that ally voters are more sensitive to cues on violence than other voters, since pre-electoral violence tends to be targeted toward voters like them who have weaker political loyalties.

As mentioned earlier, the theoretical literature on electoral malpractice assumes that candidates using violence do not suffer reduced turnout of their core supporters (Collier & Vicente, 2012). To test this assumption more directly, we use the specification shown in Equation (2), where we interact the candidate's attributes with whether the respondent is a coethnic or an ally. As Figure 7 shows, the turnout of the Kikuyu, Luo and ally voters is unaffected when comparing their candidates rumoured to have used violence to those candidates for whom there is no information on use of violence. As before, the turnout of allies is reduced when comparing candidates rumoured to have used electoral violence to those who are explicitly reported to have never used violence. In this set of comparisons, however, the turnout for the Kikuyu and Luo voters is reduced for their coethnic candidates. Overall, these results suggest that perceived violations of democratic norms, such as rumours of candidate use of violence, may reduce turnout even among the candidates' core supporters.

Figure 6

Figure 7

Although the experiment's turnout was high, our analysis of vote choice could be biased if some respondents with specific characteristics chose not to cast a valid vote. In Online appendix D, as a robustness check, we test for a potential self-selection bias using Heckman regressions (Heckman, 1979). Based on those results, we confirm no self-selection issues affect the results presented, which

might explain why we also failed to find any differentiated turnout behaviour between victims and non-victims of electoral violence and poor and non-poor voters (Figures A.5 and A.6).

Conclusion

This article makes several notable contributions to the study of election violence. Contrary to the assumptions commonly made in the theoretical literature on electoral malpractice, violent politicians may be sanctioned by voters, including their core constituents (de Figueredo & Weingast, 1999; Chaturvedi, 2005; Collier & Vicente, 2012). So if voters are willing to cast votes against violent politicians, why do these candidates still win? Our experiment cannot, of course, rule out the possibility of social desirability bias guiding some respondents' decisions, but we feel that these considerations are lessened given the experimental design. Our findings suggest that the success of violent candidates is likely to be driven by the size of candidates' core supporter base and by the presence of voters who are less likely to sanction violence. Violent candidates were able to win elections in our experiment because they lost the support of only some of their coethnic and copartisan supporters, leaving a sizeable enough support base to ensure victory. Furthermore, their victory was also due to the weak sanctioning of violence by some vulnerable groups (victims of past electoral violence and those living in poverty), especially for candidates who had a good record of reducing poverty.

Previous experimental and observational studies in Kenya and elsewhere in Africa suggest that voters condemn the use of electoral violence regardless of whether it is used in self-defence or in retaliation (Rosenzweig, 2015). Participants in our nationally representative survey also reported that they disapprove of the use of violence. Once these participants were put into an imaginary election context, however, they discounted the rumours of their coethnic or copartisan candidates having used electoral violence, especially if they expected these candidates to deliver more benefits if elected.

Our results contrast with other experiments that suggest voters, regardless of their ethnicity, wealth or victimisation, sanction candidates rumoured to have used electoral violence. Our opposing results are partly because our experiment presented voters with two alternative candidates instead of only one (e.g., LeBas, 2010; Rosenzweig, 2015). Our experiment also counterbalanced the candidates' negative attributes – i.e., rumoured use of electoral violence – with potential benefits that the candidates might bring. This is a novel design that departs from past survey experiments that ask participants to choose between candidates possessing only negative or positive attributes (Conroy-Krutz, 2013; Banerjee et al., 2014). Participants in our experiment were also presented with candidates rumoured to have ordered murders and hired gangs in the 2007 election. This strategic use of electoral violence is common and often justified by political actors in Kenya, where a third of voters experience electoral intimidation via politically connected gangs ahead of elections (Dercon &

Gutiérrez-Romero, 2012). Thus, our experiment provided a realistic setting, comparable to the choices that voters often face in new democracies marred by recurrent violence.

Overall, our findings help to deepen our understanding of why vulnerable voters might prioritise other factors when evaluating politicians rumoured to be associated with violence. A good record of reducing poverty is a key concern among many African voters (Bratton, Bhavnani & Chen, 2012), but we also provide evidence that trade-offs between violence and performance are more likely among some segments of the population. Our findings thus help to explain why politicians risk using violence, despite the loss of some of their core constituents' votes. We feel this sheds light on why electoral violence has been difficult to eradicate in settings characterised by clientelism and by political discourse that justifies the use of violence.

Acknowledgements

We thank the journal's editors and reviewers, Regina Bateson, Catia Batista, Elizabeth Carlson, Donald Green, Ryan Jablonski, Hye-Sung Kim, Patrick Kuhn, Ryan Moore, Adam Pepelasis, Olga Shemyakina, Pedro Vicente, and Matthew Wright for comments. We thank Amanda McDermott for research assistance.

Funding

We acknowledge the financial support of the UK Department for International Development and the Spanish Ministry of Science and Innovation (EC02013-46516-C4-1-R).

Replication data

The replication files and the Online appendices can be found at <http://www.prio.org/jpr/datasets>.

References

- Adar, Korwa G & Isaac M Muniya (2001) Human rights abuse in Kenya under Daniel Arap Moi, 1978-2001. *African Studies Quarterly* 5(1): 1–17.
- Alesina, Alberto; Arnaud Devleeschauwer, William Easterly, Sergio Kurlat & Romain Wacziarg (2003) Fractionalization. *Journal of Economic Growth* 8(2): 155–194.
- Banerjee, Abhijit; Donald P Green, Jeffery McManus & Rohini Pande (2014) Are poor voters indifferent to whether elected leaders are criminal or corrupt? A Vignette experiment in rural India. *Political Communication* 31(3): 391–407.
- Bekoe, Dorina (2012) Introduction: The scope, nature and pattern of electoral violence in Sub-Saharan Africa. In: Dorina Bekoe (ed.) *Voting in Fear: Electoral Violence in Sub-Saharan Africa*. Washington, DC: United States Institute for Peace, 1–13.
- Birch, Sarah (2010) Perceptions of electoral fairness and voter turnout. *Comparative Political Studies* 43(12): 1601–1622.
- Blair, Graeme; C Christine Fair, Neil Malhotra & Jacob N Shapiro (2013) Poverty and support for militant politics: Evidence from Pakistan. *American Journal of Political Science* 57(1): 30–48.
- Boone, Catherine (2011) Politically allocated land rights and the geography of electoral violence. *Comparative Political Studies* 44(10): 1311–1342.
- Bratton, Michael; Ravi Bhavnani & Tse-Hsin Chen (2012) Voting intentions in Africa: Ethnic, economic or partisan? *Commonwealth & Comparative Politics* 50(1): 27–52.
- Bratton, Michael & Mwangi S Kimenyi (2008) Voting in Kenya: Putting ethnicity in perspective. *Journal of Eastern African Studies* 2(2): 272–289.
- Carlson, Elizabeth (2016) Identifying and interpreting the sensitivity of ethnic voting in Africa. *Public Opinion Quarterly* 80(4): 837–857.
- Chandra, Kanchan (2004) *Why Ethnic Parties Succeed: Patronage and Ethnic Head Counts in India*. Cambridge: Cambridge University Press.
- Chaturvedi, Ashish (2005) Rigging elections with violence. *Public Choice* 125(1–2): 189–202.
- Collier, Paul & Anke Hoeffler (2009) Democracy's achilles heel: Or how to win an election without really trying. CSAE working paper 2009-08 (<https://www.csae.ox.ac.uk/papers/democracy-153-s-achilles-heel-or-how-to-win-an-election-without-really-trying>).
- Collier, Paul & Pedro C Vicente (2012) Violence, bribery, and fraud: the political economy of elections in Sub-Saharan Africa. *Public Choice* 153(1–2): 117–147.
- Collier, Paul & Pedro C Vicente (2014) Votes and violence: Evidence from a field experiment in Nigeria. *Economic Journal* 124(574): F327–F355.
- Conroy-Krutz, Jeffrey (2013) Information and ethnic politics in Africa. *British Journal of Political Science* 43(2): 1–29.

- de Figueredo, Rui JP & Barry R Weingast (1999) The rationality of fear: Political opportunism and ethnic conflict. In: Barbara F Walter & Jack L Snyder (eds) *Civil Wars, Insecurity, and Intervention*. New York: Columbia University Press, 261–302.
- Dercon, Stefan & Roxana Gutiérrez-Romero (2012) Triggers and characteristics of the 2007 Kenyan electoral violence. *World Development* 40(4): 731–744.
- Easterly, William & Ross Levine (1997) Africa's growth tragedy: Policies and ethnic divisions. *Quarterly Journal of Economics* 112(4): 1203–1250.
- Ferree, Karen E; Clark C Gibson & James D Long (2014) Voting behavior and electoral irregularities in Kenya's 2013 Election. *Journal of Eastern African Studies* 8(1): 153–172.
- Franck, Raphael & Ilia Rainer (2012) Does the leader's ethnicity matter? Ethnic favoritism, education, and health in Sub-Saharan Africa. *American Political Science Review* 106(2): 294–325.
- Gutiérrez-Romero, Roxana (2013) Decentralisation, accountability and the 2007 MP elections in Kenya. *Journal of Development Studies* 49(1): 72–94.
- Gutiérrez-Romero, Roxana (2014) An inquiry into the use of illegal electoral practices and effects of political violence and vote-buying. *Journal of Conflict Resolution* 58(8): 1500–1527.
- Hafner-Burton, Emilie M; Susan D Hyde & Ryan S Jablonski (2014) When do governments resort to election violence? *British Journal of Political Science* 44(01): 149–179.
- Hainmueller, Jens; Daniel J Hopkins & Teppei Yamamoto (2014) Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis* 22(01): 1–30.
- Heckman, James J (1979) Sample selection bias as a specification error. *Econometrica* 47(1): 153–161.
- Horowitz, Jeremy & Kathleen Klaus (2018) Can politicians exploit ethnic grievances? An experimental study of land appeals in Kenya. *Political Behavior* (July): 1–24.
- Human Rights Watch (1993) Divide and rule: State-sponsored ethnic violence in Kenya. In: *In Secondary Divide and Rule: State-Sponsored Ethnic Violence in Kenya*. New York: Human Rights Watch.
- Human Rights Watch (2013) *High Stakes. Political Violence and the 2013 Elections in Kenya*. New York: Human Rights Watch.
- Ichino, Nahomi & Noah Nathan (2013) Crossing the Line: Local Ethnic Geography and Voting in Ghana. *American Political Science Review* 107(2): 344–61.
- Isaksson, Ann-Sofie & Arne Bigsten (2017) Clientelism and ethnic divisions in African countries. *African Affairs* 116(465): 621–647.
- Jensen, Peter Sandholt & Mogens K Justesen (2014) Poverty and vote buying: Survey-based evidence from Africa. *Electoral Studies* 33(March): 220–232.
- Kramon, Eric (2013) Vote buying and electoral turnout in Kenya. In: Michael Bratton (ed.) *Voting*

- and Democratic Representation in Africa*. Boulder, CO: Lynne Rienner.
- Kuenzi, Michelle & Gina MS Lambright (2011) Who votes in Africa? An examination of electoral participation in 10 African countries. *Party Politics* 17(6): 767–799.
- LeBas, Adrienne (2010) *Ethnicity and the Willingness to Sanction Violent Politicians*. Afrobarometer Working Paper 125.
- Leeper, Thomas J; Sara B Hobolt & James Tilley (forthcoming) Measuring subgroup preferences in conjoint experiments. *Political Analysis* DOI: <https://doi.org/10.1017/pan.2019.30>.
- Lynch, Gabrielle (2014) Electing the ‘alliance of the accused’: The success of the Jubilee Alliance in Kenya’s Rift Valley. *Journal of Eastern African Studies* 8(1): 93–114.
- Montalvo, Jose G & Marta Reynal-Querol (2003) Religious polarization and economic development. *Economics Letters* 80(2): 201–210.
- Nathan, Noah L (2016) Does participation reinforce patronage? Policy preferences, turnout and class in urban Ghana. *British Journal of Political Science* 49(1): 229–255.
- Posner, Daniel N (2005) *Institutions and Ethnic Politics in Africa*: Cambridge University Press.
- Rosenzweig, Steven C (2015) Dangerous disconnect: How politicians’ misperceptions about voters lead to violence in Kenya. Working Paper (http://cega.berkeley.edu/assets/miscellaneous_files/117_-_Rosenzweig_Dangerous_Disconnect_-_May_14_2015_-_ABCA.pdf).
- Stokes, Susan C (2009) *Political Clientelism*. In: Carles Boix & Susan C Stokes (eds). *The Oxford Handbook of Comparative Politics*. Oxford: Oxford University Press, 604–627.
- Straus, Scott & Charles Taylor (2013) Democratization and electoral violence in Sub-Saharan Africa. In: Dorine A Bekoe (ed.) *Voting in Fear: Electoral Violence in Sub-Saharan Africa*. Washington, DC: USIP Press.
- Vaishnav, Milan (2017) *When crime pays: Money and muscle in Indian politics*: Yale University Press.
- von Borzyskowski & Patrick M Kuhn (2020) Dangerously informed: Voter information and pre-electoral violence in Africa. *Journal of Peace Research* 57(1): XX-XX.
- Wallsworth, Gregory (2016) *Popular Participation and Political Violence*. Ph.D. Dissertation, Economics, Michigan State University.
- Wantchekon, Leonard (2003) Clientelism and voting behavior: Evidence from a field experiment in Benin. *World Politics* 55(03): 399–422.
- Weitz-Shapiro, Rebecca (2012) What wins votes: Why some politicians opt out of clientelism. *American Journal of Political Science* 56(3): 568–583.
- Wilkinson, Steven I & Christopher J Haid (2009) Ethnic violence as campaign expenditure: Riots, competition, and vote swings in India. Working paper, University of Chicago.
- Winters, Matthew S & Rebecca Weitz-Shapiro (2013) Lacking information or condoning corruption: when do voters support corrupt politicians? *Comparative Politics* 45(4): 418–436.

Young, Lauren (2020) Who dissents? Self-efficacy and opposition action after state-sponsored election violence. *Journal of Peace Research* 57(1): XX-XX.

GUTIERREZ-ROMERO, Roxana. PhD in Economics (University of Cambridge, 2006); Reader, Queen Mary University of London (2019–). Research interests: development and violence.

LEBAS, Adrienne. PhD in Political Science (Columbia University, 2006); Associate Professor, American University (2009–). Research interests: electoral politics and violence.

Table I. Candidate attributes

Attribute	Level 1	Level 2	Level 3
Identity	Onyango(Luo)/ODM	Kamau(Kikuyu)/TNA	
Violence	No mention	Has never used electoral violence	Rumoured to have used electoral violence
Performance	No mention	Unknown performance, not audited	Good performance in reducing poverty

Table II. Pairs of candidates presented for each ballot

	Candidate 1: Onyango/ODM	Candidate 2: Kamau/TNA
	Ballot 1	
	No mention on use of violence & no mention of performance	No mention on use of violence & no mention of performance
	Ballot 2 (Respondent hears one of the nine randomised scripts)	
1	Has never used electoral violence & no mention of performance	Has never used electoral violence & no mention of performance
2	Has never used electoral violence & no mention of performance	Has used electoral violence & no mention of performance
3	Has never used electoral violence & unknown performance	Has used electoral violence & good performance
4	Has used electoral violence & no mention of performance	Has never used electoral violence & no mention of performance
5	Has used electoral violence & no mention of performance	Has used electoral violence & no mention of performance
6	Has used electoral violence & unknown performance	Has used electoral violence & good performance
7	Has used electoral violence & good performance	Has never used electoral violence & unknown performance
8	Has used electoral violence & good performance	Has used electoral violence & unknown performance
9	Has used electoral violence & good performance	Has used electoral violence & good performance




Figure 1. Effects of candidate's attributes on vote choice

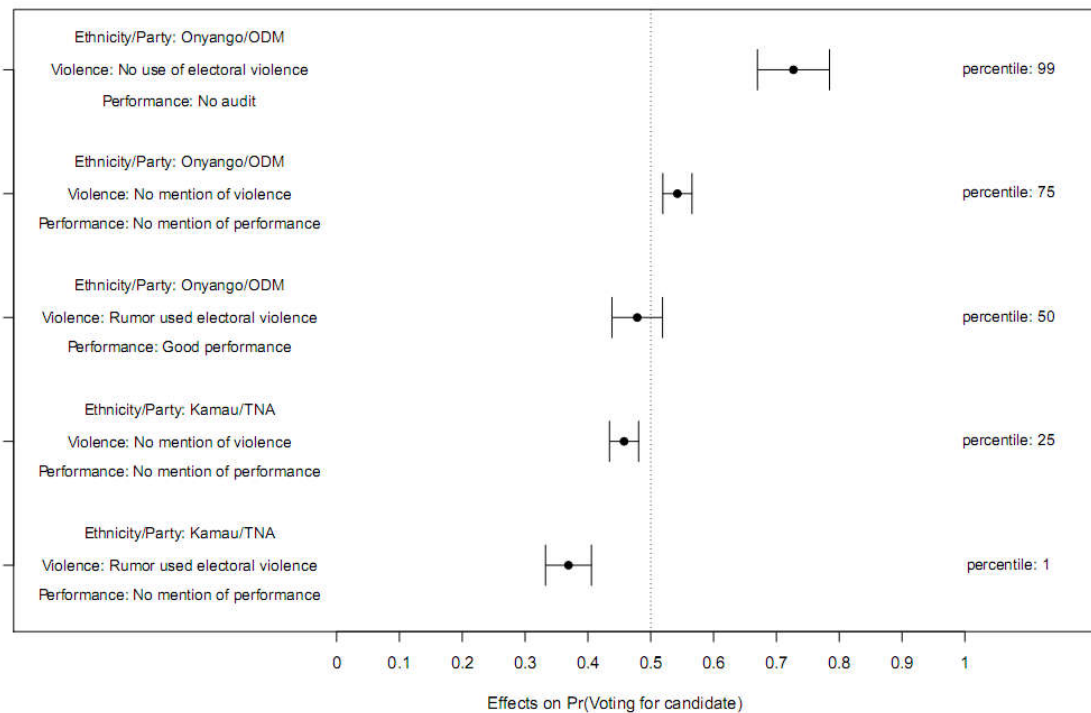


Figure 2. Effects of candidates' profiles on vote choice and probability of being elected


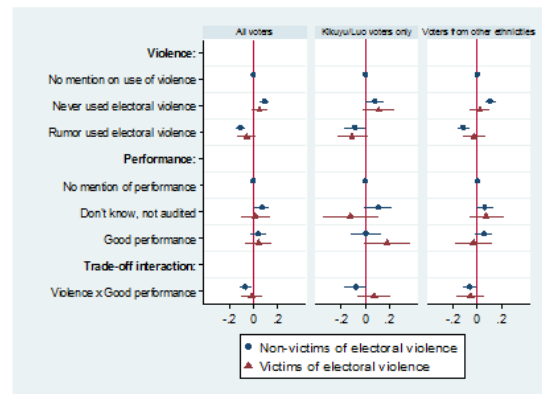
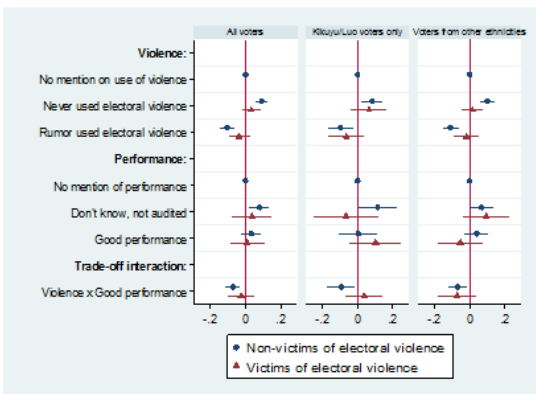


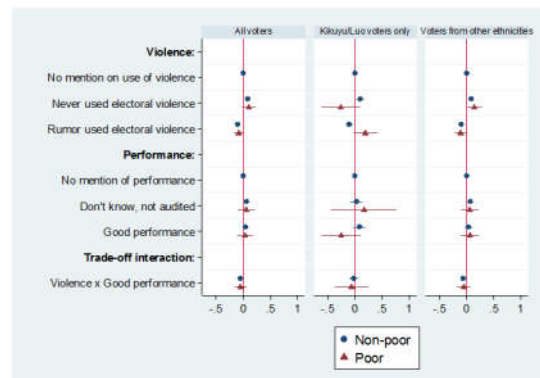
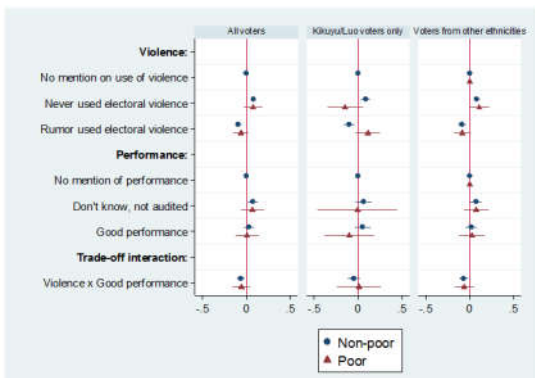
Figure 3. Effects of coethnic/ally candidate's attributes on vote choice



Panel A (without controls)

Panel B (with controls)

Figure 4. Effect of candidate's attributes on the vote choice of victims of electoral violence



Panel A (without controls)

Panel B (with controls)

Figure 5. Effect of candidate's attributes on the vote choice of poor respondents

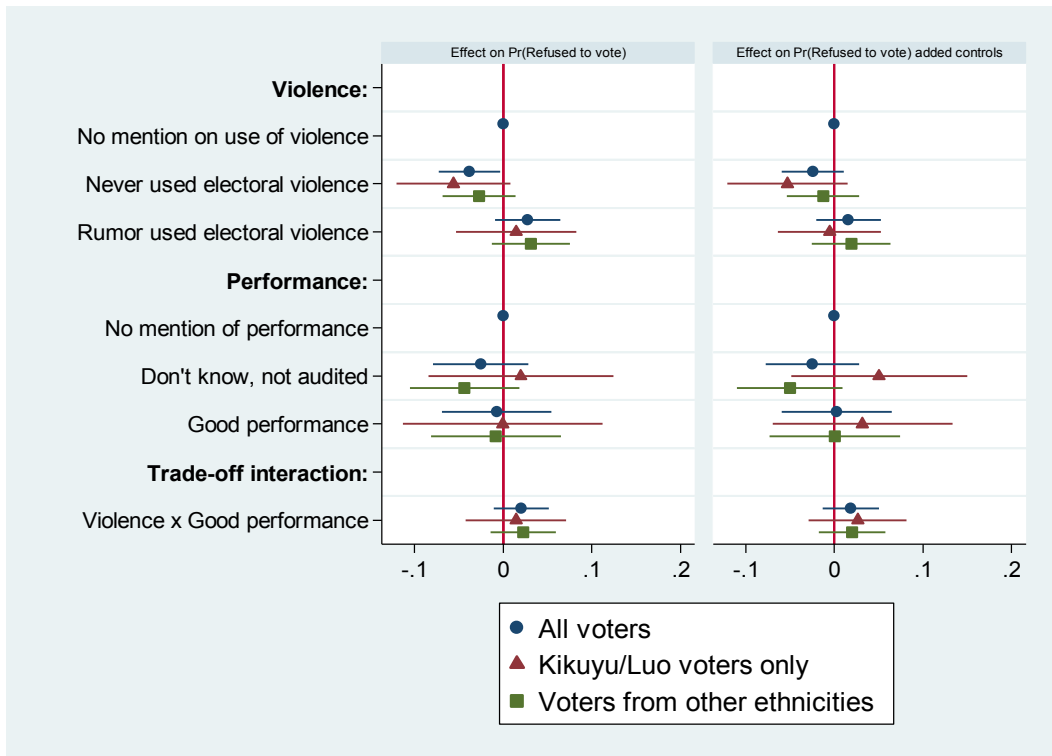


Figure 6. Effects of candidate's attributes on refusing to vote or spoiling ballots

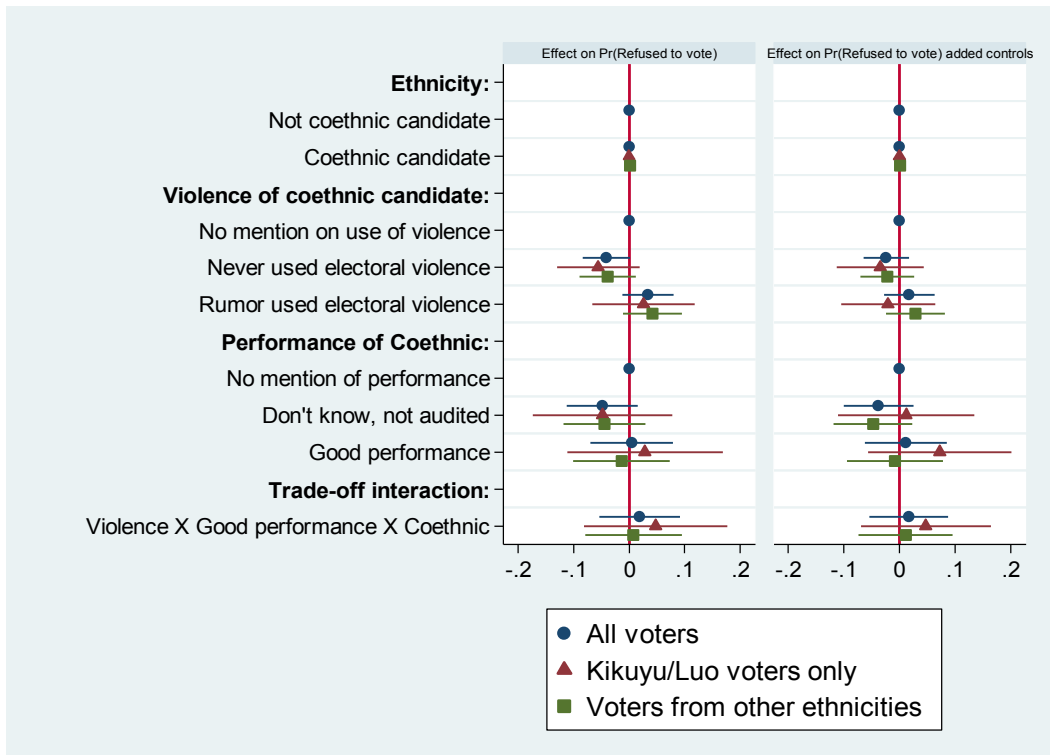


Figure 7. Effects of coethnic/ally candidate's attributes on refusing to vote or spoiling ballots

Online Appendix

Appendix A

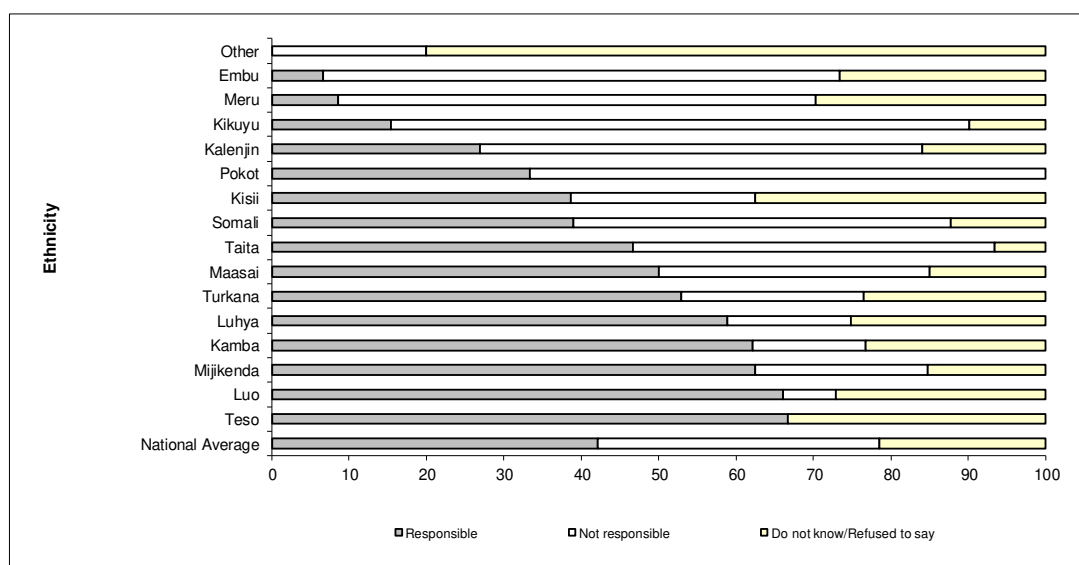


Figure A.1. Believes president Uhuru Kenyatta is responsible for the violence of 2007 disputed election

Table A.1. Characteristics of respondents

	Percent
Aged 18-26	34.2
Female	47.8
Secondary or higher education attainment	76.8
Has been victim of electoral violence at least once since 1991	24.1
Food deprived	48.8
Feels close to a political party	41.1
Lives in a rural area	57.9
Ethnicity	
Kikuyu	16.9
Luo	15.0
Luhya	14.7
Kamba	10.2
Meru	7.9
Kisii	7.5
Kalenjin	10.1
Maasai	1.7
Mijikenda	6.1
Taita	2.7
Somali	3.4
Pokot	0.3
Turkana	1.6
Teso	0.3
Embu	1.3
Ethnicity not stated	0.4
Province	
Nairobi	7.2
Central	8.6
Eastern	16.2
Rift Valley	24.0
Nyanza	15.1
Western	12.7
North Eastern	3.2
Coast	13.0
Observations	1210




Figure A.2 Respondent characteristics balancing tests by candidates' attributes

Table A.2 Voting choice in the experiment's first ballot and respondents' self-reported voting in previous presidential elections

	First Ballot in our Experiment ^a		2007 Actual Presidential Election		2013 Actual Presidential Election	
	John Kamau (Kikuyu origin)	John Onyango (Luo Origin)	Mwai Kibaki (Kikuyu origin)	Raila Odinga (Luo origin)	Uhuru Kenyatta (Kikuyu origin)	Raila Odinga (Luo origin)
Kikuyu	85.2	14.8	90.7	6.2	87.7	8.0
Luo	11.6	88.4	15.2	81.1	6.8	91.2
Luhya	26.1	73.9	37.3	42.9	28.8	60.2
Kamba	29.1	70.9	19.0	42.1	14.9	79.3
Meru	80.7	19.3	92.6	6.2	89.5	4.7
Kisii	30.3	69.7	41.7	53.3	35.1	52.7
Kalenjin	69.1	30.9	32.5	58.8	73.0	13.5
Maasai	73.7	26.3	30.0	70.0	47.1	41.2
Mijikenda	18.6	81.4	25.5	67.3	9.6	80.8
Taita	40.0	60.0	27.3	68.2	30.0	65.0
Somali	46.7	53.3	21.2	63.6	35.7	64.3
Turkana	50.0	50.0	64.3	35.7	40.0	53.3
Embu	85.7	14.3	92.3	7.7	84.6	7.7
Others	28.6	71.4	40.0	60.0	28.6	42.9
All Groups	44.7	55.3	46.2	43.3	45.3	47.5

^a Percentage of valid votes in the first ballot.

Table A.3 Effects of candidate's attributes on voting choice as reported in Equation (1)

	All respondents		Kikuyu & Luo		Other ethnic groups	
	(1)	(2)	(3)	(4)	(5)	(6)
Candidate's identity (baseline, Luo/ODM)						
Kikuyu/PNU	-0.085*	-0.110*	0.074	0.046	-0.153*	-0.170*
	(0.028)	(0.030)	(0.052)	(0.060)	(0.033)	(0.035)
Candidate's electoral violence attribute (baseline, no mention on use of electoral violence)						
Has never used electoral violence	0.079*	0.085*	0.074*	0.079*	0.081*	0.089*
	(0.015)	(0.017)	(0.026)	(0.033)	(0.018)	(0.019)
Rumoured to have used electoral violence	-0.088*	-0.094*	-0.085*	-0.085*	-0.089*	-0.099*
	(0.017)	(0.018)	(0.030)	(0.035)	(0.020)	(0.021)
Candidate's performance attribute (baseline, no mention of performance)						
Not audited yet, thus performance is unknown	0.041	0.055	0.016	0.033	0.048	0.059
	(0.043)	(0.046)	(0.078)	(0.086)	(0.050)	(0.054)
Good performance record of reducing poverty	0.025	0.041	0.040	0.061	0.019	0.036
	(0.026)	(0.028)	(0.045)	(0.053)	(0.031)	(0.033)
Interaction between candidate's performance and electoral violence attributes						
Never used violence x Not audited yet thus performance is unknown	0.065	0.017	0.085	0.006	0.057	0.022
	(0.060)	(0.066)	(0.110)	(0.128)	(0.070)	(0.076)
Other controls						
No		Yes	No	Yes	No	Yes
Constant	0.542*	0.567*	0.463*	0.470*	0.576*	0.607*
	(0.014)	(0.022)	(0.026)	(0.039)	(0.017)	(0.024)
Observations	3,724	3,230	1,134	910	2,570	2,312
R-squared	0.023	0.026	0.020	0.012	0.040	0.044

Robust standard errors clustered by respondent. Significance level * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$. The dependent variable equals to 1 if the respondent voted for a particular candidate in each of the two ballots and 0 otherwise. The controls used but not shown in the table are: respondent's ethnicity, sex, age, education, province of residence, residency area (urban/rural), whether the respondent feels close to a party, has experienced of food shortages within the past year, and has experienced of electoral violence since 1992. Also included the ethnolinguistic fragmentation and polarisation indices at the district level and interviewers' mother tongue and their number of years of experience conducting interviews. Appendix C describes how the AMCEs shown in Figure 1 were estimated using the regression coefficients presented in this table. The violence and performance attributes V and P have three levels each. As the experiment does not include any of the interactions $V1*P2$, $V1*P3$ and $V2*P3$, these three interactions are empty and cannot be estimated. Due to multicollinearity the regression coefficients of the interactions $V3*P2$ and $V3*P3$ are not shown in this table, but these interactions were obtained by re-running the same regression just reparametrizing the parameters to rescue the interaction $V3*P3$ which is shown in Figure 1.

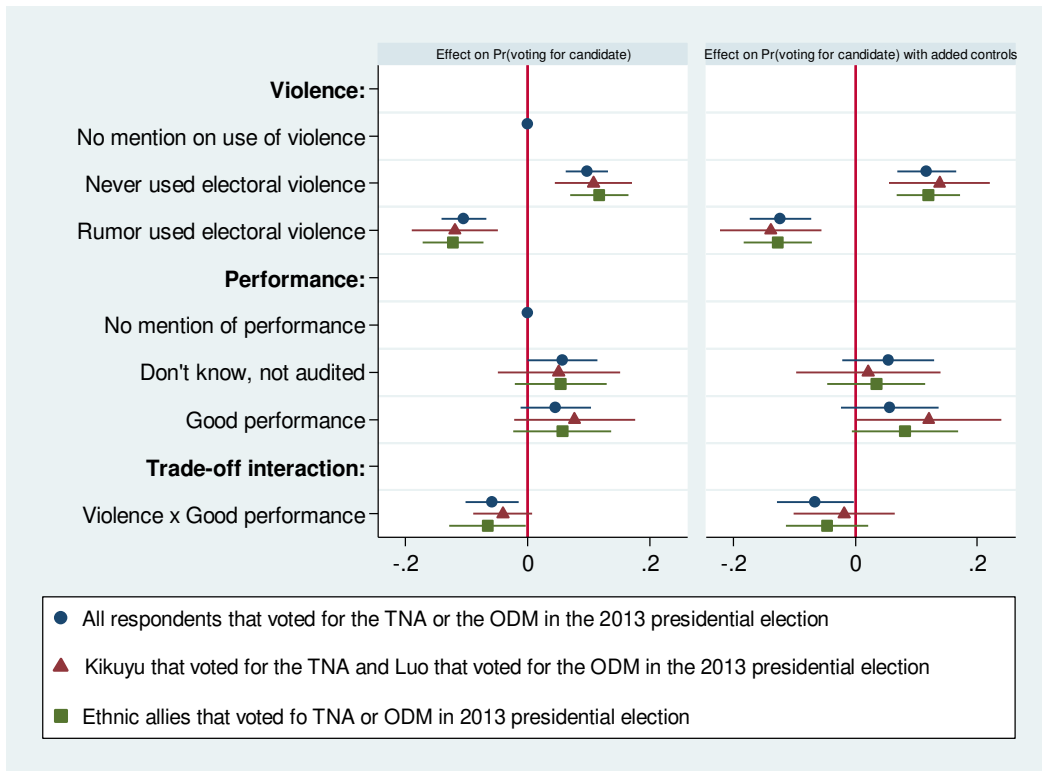


Figure A.3 Effects of candidate's attributes on voting choice by respondent's ethnicity and 2013 partisan support

Table A.4 Effects of candidate's attributes on voting choice by those who have experienced electoral violence at least once since 1992

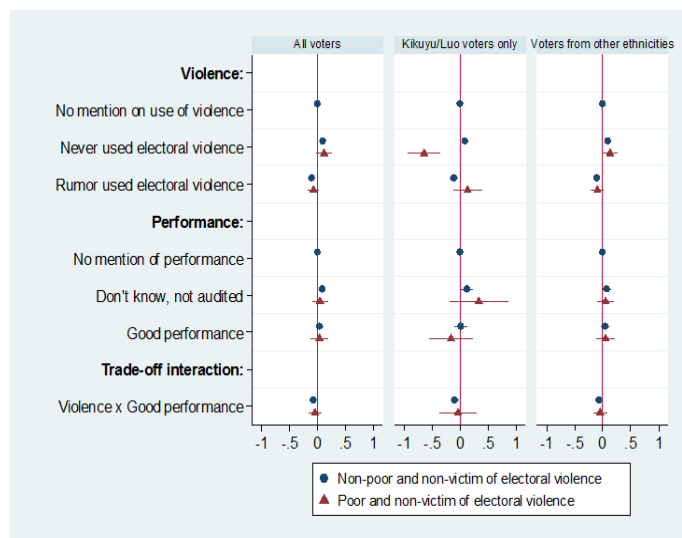
	All respondents				Kikuyu & Luo				Other ethnic groups			
	Non-victim of electoral violence	Non-victim of electoral violence	Victim of electoral violence	Victim of electoral violence	Non-victim of electoral violence	Non-victim of electoral violence	Victim of electoral violence	Victim of electoral violence	Non-victim of electoral violence	Non-victim of electoral violence	Victim of electoral violence	Victim of electoral violence
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Candidate's identity (baseline, Luo/ODM)												
Kikuyu/PNU	-0.047 (0.032)	-0.073* (0.035)	-0.204* (0.058)	-0.230* (0.064)	0.190* (0.061)	0.153* (0.072)	-0.153* (0.092)	-0.158 (0.108)	-0.132* (0.037)	-0.146* (0.040)	-0.240* (0.074)	-0.269* (0.081)
Candidate's electoral violence attribute (baseline, no mention on use of electoral violence)												
Has never used electoral violence	0.093* (0.018)	0.096* (0.019)	0.032 (0.027)	0.053 (0.035)	0.081* (0.030)	0.079* (0.040)	0.065 (0.052)	0.111 (0.068)	0.101* (0.021)	0.107* (0.022)	0.015 (0.031)	0.024 (0.043)
Rumoured to have used electoral violence	-0.105* (0.020)	-0.106* (0.021)	-0.036 (0.030)	-0.057 (0.038)	-0.098* (0.036)	-0.089* (0.044)	-0.065 (0.052)	-0.108 (0.065)	-0.108* (0.023)	-0.117* (0.025)	-0.018 (0.037)	-0.026 (0.048)
Candidate's performance attribute (baseline, no mention of performance)												
Not audited yet, thus performance is unknown	0.038 (0.049)	0.050 (0.053)	0.035 (0.092)	0.053 (0.101)	0.059 (0.095)	0.073 (0.104)	-0.122 (0.146)	-0.079 (0.181)	0.027 (0.056)	0.038 (0.061)	0.145 (0.109)	0.121 (0.121)
Good performance record of reducing poverty	0.032 (0.030)	0.043 (0.033)	0.010 (0.049)	0.043 (0.059)	0.003 (0.055)	0.007 (0.066)	0.102 (0.076)	0.183* (0.100)	0.041 (0.035)	0.057 (0.038)	-0.053 (0.066)	-0.030 (0.077)
Interaction between candidate's performance and electoral violence attributes												
Never used violence x Not audited yet thus performance is unknown	0.083 (0.068)	0.046 (0.075)	0.001 (0.126)	-0.070 (0.143)	0.112 (0.126)	0.069 (0.148)	0.112 (0.213)	-0.090 (0.273)	0.085 (0.079)	0.052 (0.086)	-0.100 (0.151)	-0.095 (0.173)
Other controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Constant	0.524* (0.016)	0.544* (0.033)	0.602* (0.029)	0.623* (0.037)	0.405* (0.031)	0.439* (0.052)	0.576* (0.046)	0.733* (0.109)	0.566* (0.019)	0.615* (0.046)	0.620* (0.037)	0.625* (0.044)
Observations	2,820	2,468	904	762	758	614	376	296	2,042	1,846	528	466
R-squared	0.025	0.025	0.044	0.055	0.062	0.042	0.029	0.039	0.039	0.042	0.068	0.082

Robust standard errors clustered by respondent. Significance level * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$. The dependent variable equals to 1 if the respondent voted for a particular candidate in each of the two ballots and 0 otherwise. The controls used but not shown in the table are: respondent's ethnicity, sex, age, education, province of residence, residency area (urban/rural), whether respondent feels close to a party, has experienced of food shortages within the past year, and has experienced of electoral violence since 1992. Also included the ethnolinguistic fragmentation and polarisation indices at the district level and interviewers' mother tongue and their number of years of experience conducting interviews. Appendix C describes how the AMCEs shown in Figure 4 were estimated using the regression coefficients presented in this table. The violence and performance attributes V and P have three levels each. As the experiment does not include any of the interactions $V1*P2$, $V1*P3$ and $V2*P3$, these three interactions are empty and cannot be estimated. Due to multicollinearity the regression coefficients of the interactions $V3*P2$ and $V3*P3$ are not shown in this table, but these interactions were obtained by re-running the same regression just reparametrizing the parameters to rescue the interaction $V3*P3$ which is shown in Figure 4.

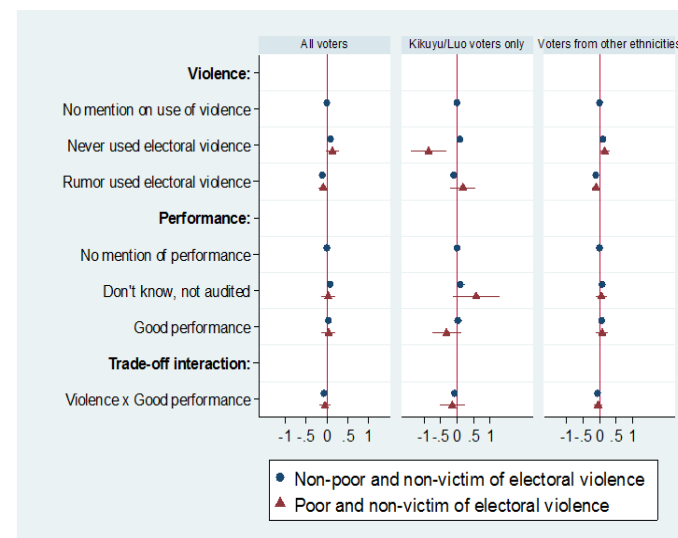
Table A.5 Effects of candidate's attributes on voting choice by poor and non-poor respondents

	All respondents				Kikuyu & Luo				Other ethnic groups			
	Non-poor	Non-poor	Poor	Poor	Non-poor	Non-poor	Poor	Poor	Non-poor	Non-poor	Poor	Poor
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Candidate's identity (baseline, Luo/ODM)												
Kikuyu/PNU	-0.071*	-0.100*	-0.170*	-0.172*	0.089*	0.046	-0.102	0.065	-0.148*	-0.161*	-0.177*	-0.215*
	(0.031)	(0.033)	(0.073)	(0.084)	(0.054)	(0.062)	(0.195)	(0.255)	(0.037)	(0.039)	(0.080)	(0.089)
Candidate's electoral violence attribute (baseline, no mention on use of electoral violence)												
Has never used electoral violence	0.079*	0.085*	0.077	0.101	0.085*	0.095*	-0.150	-0.261	0.077*	0.084*	0.107*	0.141*
	(0.015)	(0.017)	(0.055)	(0.064)	(0.026)	(0.034)	(0.104)	(0.182)	(0.018)	(0.020)	(0.060)	(0.071)
Rumoured to have used electoral violence	-0.093*	-0.098*	-0.062	-0.080	-0.100*	-0.105*	0.112	0.192	-0.090*	-0.098*	-0.087*	-0.115*
	(0.018)	(0.020)	(0.044)	(0.051)	(0.031)	(0.037)	(0.072)	(0.117)	(0.021)	(0.024)	(0.049)	(0.058)
Candidate's performance attribute (baseline, no mention of performance)												
Not audited yet, thus performance is unknown	0.026	0.035	0.105	0.160	0.026	0.033	-0.112	0.037	0.023	0.031	0.137	0.186
	(0.047)	(0.051)	(0.100)	(0.118)	(0.082)	(0.090)	(0.273)	(0.421)	(0.057)	(0.061)	(0.108)	(0.125)
Good performance record of reducing poverty	0.029	0.045	0.005	0.030	0.051	0.084	-0.102	-0.255	0.018	0.033	0.023	0.060
	(0.028)	(0.031)	(0.068)	(0.077)	(0.046)	(0.056)	(0.144)	(0.188)	(0.034)	(0.037)	(0.075)	(0.086)
Interaction between candidate's performance and electoral violence attributes												
Never used violence x Not audited yet thus performance is unknown	0.090	0.050	-0.073	-0.214	0.069	-0.004	0.211	0.258	0.099	0.072	-0.131	-0.263
	(0.064)	(0.071)	(0.164)	(0.208)	(0.113)	(0.132)	(0.457)	(0.766)	(0.077)	(0.084)	(0.177)	(0.227)
Other controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Constant	0.536*	0.555*	0.585*	0.598*	0.455*	0.493*	0.551*	0.873	0.574*	0.595*	0.588*	0.650*
	(0.015)	(0.025)	(0.037)	(0.078)	(0.027)	(0.045)	(0.098)	(0.860)	(0.018)	(0.027)	(0.040)	(0.078)
Observations	3,196	2,784	528	446	1,044	836	90	74	2,136	1,940	434	372
R-squared	0.023	0.026	0.036	0.037	0.026	0.016	0.023	0.034	0.040	0.043	0.043	0.059

Robust standard errors clustered by respondent. Significance level * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$. The dependent variable equals to 1 if the respondent voted for a particular candidate in each of the two ballots and 0 otherwise. The controls used but not shown in the table are: respondent's ethnicity, sex, age, education, province of residence, residency area (urban/rural), whether respondent feels close to a party, has experienced of food shortages within the past year, and has experienced of electoral violence since 1992. Also included the ethnolinguistic fragmentation and polarisation indices at the district level and interviewers' mother tongue and their number of years of experience conducting interviews. Appendix C describes how the AMCEs shown in Figure 5 were estimated using the regression coefficients presented in this table. The violence and performance attributes V and P have three levels each. As the experiment does not include any of the interactions $V1*P2$, $V1*P3$ and $V2*P3$, these three interactions are empty and cannot be estimated. Due to multicollinearity the regression coefficients of the interactions $V3*P2$ and $V3*P3$ are not shown in this table, but these interactions were obtained by re-running the same regression just reparametrizing the parameters to rescue the interaction $V3*P3$ which is shown in Figure 5.



Panel A (without controls)



Panel B (with controls)

Figure A.4 Effect of candidate's attributes on the voting choice of poor and non-poor respondents that have never been victims of electoral violence

Table A.6 Valid votes and refusals in the first and second ballot

	First ballot	Second ballot (respondent hears one of nine randomised scripts)								
		1	2	3	4	5	6	7	8	9
%Voted for Onyango (Luo candidate)	44.8	54.6	52.73	61.74	25.27	33.12	34.11	29.23	41.3	32.04
%Voted for Kamau (Kikuyu candidate)	36.2	34.4	20.91	24.35	45.6	38.31	37.21	54.62	31.52	24.27
%Voted for both candidates in ballot	0.9	1.3	1.82	1.74	0.55	0.65	1.55	1.54	1.09	1.94
%Left ballot in blank	4.0	3.3	5.45	4.35	5.49	14.29	8.53	3.85	10.87	7.77
%Refused to vote, said was indifferent between both candidates	6.1	3.9	14.55	4.35	10.44	9.09	10.85	5.38	10.87	9.71
%Refused to vote without giving any reason	8.0	2.6	4.55	3.48	12.64	4.55	7.75	5.38	4.35	24.27
Number of observations	1,210	154	110	115	182	154	129	130	92	103

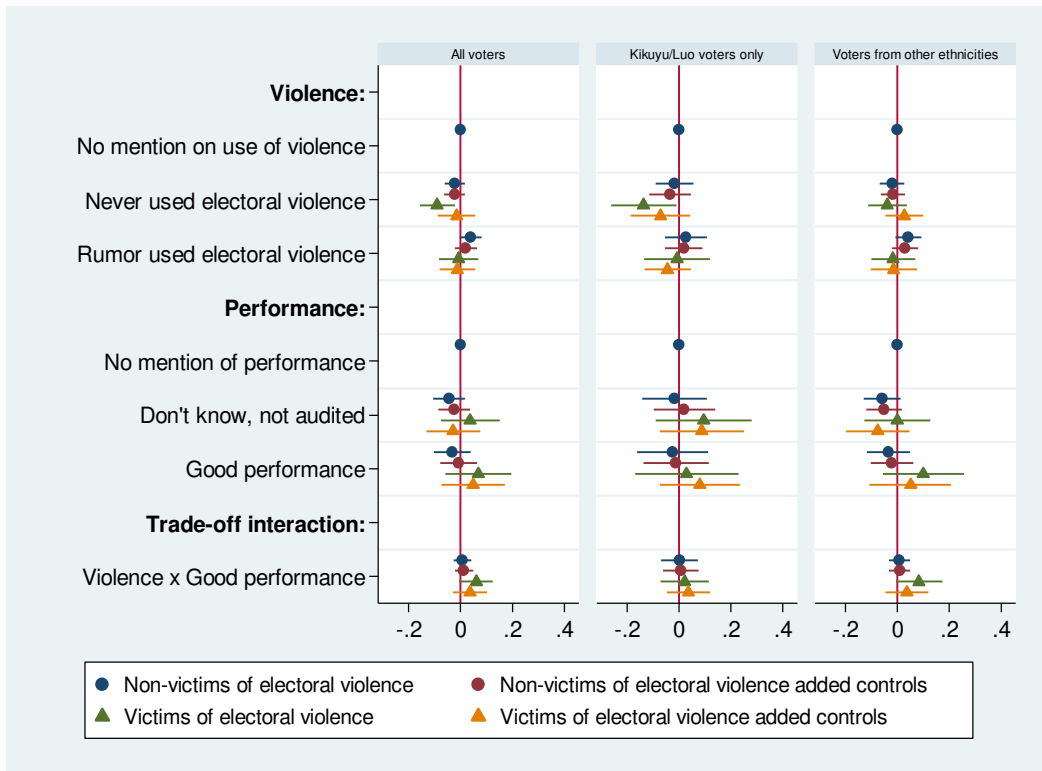


Figure A.5 Effects of candidate's attributes on refusing to vote or spoiling ballots by respondents that have been victims and non-victims of electoral violence

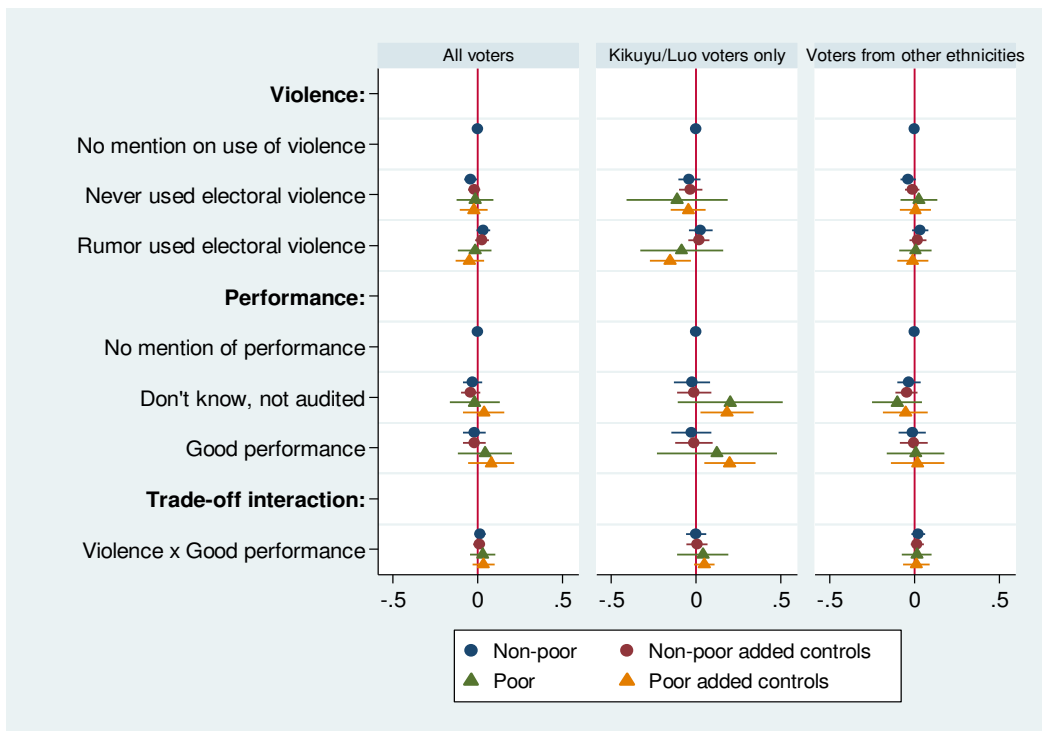


Figure A.6 Effects of candidate's attributes on refusing to vote or spoiling ballots by poor and non-poor respondents

Appendix B

Sampling and ethical considerations

The experiment is set in a nationally representative survey we conducted in Kenya, in December 2013. The national representation of the sample allows to test our research hypotheses across a range of different sub-populations of interest, which would have been impossible to do had we focused on a single area or sub-population as is commonly done in similar conjoint experiments.

The households and interviewed respondents were selected using the same sampling method that the Afrobarometer follows. Specifically, we drew our random sample following five stages. First, the sample was stratified by province and by urban and rural areas. Then, as the second stage, we randomly selected primary sampling units (PSU's) using the Kenyan Census. In the third stage, within each PSU, the field manager randomly selected a sampling start point. Then as the fourth stage, within the selected sampling point the interviewers randomly selected households. To this end, in each sampling point, the team of interviewers were told to walk in opposite directions from each other. In the last fifth stage within the randomly selected household, the interviewer randomly selected an individual respondent. To ensure a gender balance in the sample, interviewers alternated in each household between interviewing a man and interviewing a woman. Also, in line with standard practices, about eight interviews were clustered within each selected PSU to keep costs within budget. The sampling method we followed is the same as the one used by the Afrobarometer, described in further detail at:

<http://www.afrobarometer.org/surveys-and-methods/sampling-principles>

The questionnaire was carefully designed to satisfy the ethical, data confidentiality and data protection requirements set by our sponsor the UK Department for International Development (DFID). Whenever possible, the questions on victimisation were designed to also be compared to published research on victimisation in Kenya.

None of the survey participants, interviewers and dedicated personnel were at risk of harm during the interview or as a direct consequence of fieldwork. Before the interview took place, the randomly selected respondents were told that every person in the country had an equal chance of being included in our survey, and there was no penalty for refusing to participate. All survey participants were informed that they were going to participate in an election game, and all provided their full consent. The survey respondents were reassured their identity would not be revealed, and all the information provided would be anonymous and confidential.

As part of the experiment embedded in the survey, respondents were asked to pick a number between one and nine. This number was used by the enumerator to pick the script about the candidates' contending in an imaginary election. We chose to randomise the candidates' attributes in this way for two main reasons. First, this randomisation turned out to be the easiest


to implement in the field with our 80 enumerators working across the country. It also reduced the risk of enumerators losing other equipment that we could have used to randomise treatments (such as dice or deck of cards). Second, our extensive pilot work in Nairobi (interviewing more than 150 people) showed that respondents felt more reassured to say aloud a number of their own choice, instead of other methods we tested, such as picking a random card from a deck of cards. In this pilot, we had a couple of instances where respondents became suspicious when asked to randomly choose a card. Some respondents questioned whether all the cards had the same number, or whether the enumerator had sufficiently shuffled the deck of cards to be reassured that the game was fair. To minimise these rare instances of respondents questioning how the randomisation had been done, we decided to randomise by allowing respondents to pick a number of their own choice between one and nine. In our pilot in Nairobi we did not encounter any major biases by randomising in this way. Figure A.2 shows the randomisation worked reasonably well in the actual experiment as the characteristics of respondents across the treatments do not suffer from major biases.



As shown below, our voting choice involves two hypothetical male candidates contending in an imaginary parliamentary election. Like many other similar experiments in Sub-Saharan Africa, we ask voters which ‘candidate’ they prefer rather than which ‘party’ they prefer. This choice mimics more closely the choices that voters make in elections. The choice between candidates is also of interest in African elections, since candidates’ ethnicity is often more relevant for voters than their other attributes such as party affiliation. This is particularly the case for Kenya, where political parties have been formed every election as mere vehicles for (ethnic) groups to form alliances.¹ We focus on male candidates as the majority of MP candidates and elected MPs in Kenya are male. Thus, our respondents would be used to seeing male candidates contending in their constituency. Had we instead compared a male to a female MP, we would have introduced an additional treatment, where it would have been even harder to assess whether respondents view both candidates as equally competent. To ensure that candidates were viewed as equally competent and equally experienced, the script states that both candidates had been elected MPs in 2008–2013 in other constituencies similar to the respondent’s home. If we had instead stated that the candidates had no previous experience that information would have resulted in an additional treatment, which our experiment was not

¹ In the experiment we did not ask respondents to choose between hypothetical candidates at presidential level since this type of choice risks conflating the intended candidates’ characteristics with those of the politicians actually holding these offices. Nonetheless, given how similar Kenyans vote in real parliamentary and presidential elections (Gutiérrez-Romero, 2013) the results from our experiment might be more generalisable.

designed to test. Our focus is to test whether well-performing candidates, rumoured to have used electoral violence, still attract voting support.

As part of the experiment, respondents were asked to place their folded voting ballot in a see-through bag, alongside other folded ballots. Ballots were individually coded such that they could be later matched to their respective questionnaires. These numbers were clearly visible in the ballot, as shown in Figure B.1. However, the voting was secret in the sense that no other participant, enumerator, or supervisor ever saw how respondents voted. Also respondents were reassured that we would respect their confidentiality and no-one would be able to identify their name with their vote. As part of our ethics policy, we instructed the enumerators and field supervisors to give the bags containing the ballots to our data-entry team. This data-entry team was the one responsible for matching the ballots to their respective questionnaire. Further, to ensure the confidentiality of the survey participants, the names of the respondents were not recorded in the dataset containing the responses nor their contact details.

VOTE FOR ONE CANDIDATE ONLY 

Candidate: John Onyango ORANGE DEMOCRATIC MOVEMENT (ODM)	Tick here if vote for this candidate
	
Candidate: John Kamau The National Alliance (TNA)	Tick here if vote for this candidate
	

Once you finish voting, please fold the voting ballot and place it in the bag provided by the interviewer. Thank you.

Q1 [or Q2 depending if first or second ballot]	Survey ID number
--	------------------

Figure B.1 Voting ballot used in the first and the second vote

At the end of the experiment, participants were debriefed, where interviewers reassured participants that they had taken part in a mere game, where their voting choices would not be revealed along with any information that could identify them. None of our participants expressed any concerns about the election game, nor about the questions asked during the survey. Below, we describe the script used for the experimental game.

The script of the survey experiment

We are now going to play an elections game.

Imagine for a moment that there is a vacancy in your constituency for a Member of Parliament. So there will be a by-election held in a few weeks time. Only two candidates will be competing for the MP seat in your constituency.

Imagine that both candidates have equal experience in politics, both have been elected MPs before, but in other constituencies similar to this one during the 2008-2013 administration.

They did not contend in the 2013 elections in the constituencies they used to be MPs due to the administrative changes of boundaries that some areas had.

Both candidates are promising to improve the economy of your community. They are:

Candidate 1 is John Onyango, and he is running for the ODM party.

Candidate 2 is John Kamau, and he is running for the TNA party.

1.A Remembering that they will be the only candidates contending in this imaginary election, I am going to ask you to secretly cast your vote in this ballot for the person you would vote for between these two candidates.		
Your vote will be placed in this bag mixed with other ballots. Your vote is private, we will respect your confidentiality and no-one will identify your name with your vote.		
A	Refused to vote, respondent said was indifferent between candidates >Go to <i>Q1.B</i>	5
B	Refused to vote without giving any reason >Go to <i>Q1.B</i>	6
OPTIONS C, D, E, F WILL BE FILLED BY ENUMERATOR'S SUPERVISOR		
C	Respondent voted for Candidate 1, John Onyango in the ballot.	1
D	Respondent voted for Candidate 2, John Kamau in the ballot.	2
E	Respondent voted for both candidates in the ballot .	3
F	Respondent did not vote for either of candidates, left voting ballot in blank.	4

1.B Now I'm going to tell you a little bit more about these candidates. But first I would like you to pick any number between one and nine. [Interviewer: Circle which number the respondent chose. Read the corresponding statement of that number]

Number 1	1
Number 2	2
Number 3	3
Number 4	4
Number 5	5
Number 6	6
Number 7	7
Number 8	8
Number 9	9
Refused to select a number> Thank respondent, continue interview with question 3.	999

1	Neither John Onyango nor John Kamau has ever used violence to try winning elections.
2	<p>John Onyango has never used violence to try winning elections.</p> <p>John Kamau is rumoured to have ordered one murder and hired gangs during the 2007 elections. He has not been arrested for these alleged crimes.</p>
3	<p>John Kamau is rumoured to have ordered one murder and hired gangs during the 2007 elections. He has not been arrested for these alleged crimes. Back when John Kamau was an elected MP, he was among the few MPs that were randomly and independently audited about how MPs used the CDF fund, the fund that MPs are given to reduce poverty in their constituencies. The audit found that Kamau spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.</p> <p>John Onyango has never used violence to try winning elections. We do not have any information about how Onyango used the CDF back when he was an elected MP because he, like the majority of other MPs, was not audited.</p>
4	<p>John Onyango is rumoured to have ordered one murder and hired gangs in the 2007 elections. He has not been arrested for these alleged crimes.</p> <p>John Kamau has never used violence to try winning elections.</p>
5	Both John Onyango and John Kamau are rumoured to have ordered one murder and hired gangs during the 2007 elections. They have not been arrested for these alleged crimes.
6	<p>John Kamau is rumored to have ordered one murder and hired gangs during the 2007 elections. He has not been arrested for these alleged crimes. Back when John Kamau was an elected MP, he was among the few MPs that were randomly and independently audited about how MPs used the CDF fund, the fund that MPs are given to reduce poverty in their constituencies. The audit found that Kamau spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.</p> <p>John Onyango is rumoured to have ordered one murder and hired gangs in the 2007 elections. He has not been arrested for these alleged crimes. We do not have any information about how Onyango used the CDF, back when he was an elected MP because he, like the majority of other MPs, was not audited.</p>
7	<p>John Onyango is rumoured to have ordered one murder and hired gangs during the 2007 elections. He has not been arrested for these alleged crimes. Back when John Onyango was an elected MP, he was among the few MPs that were randomly and independently audited about how MPs used the CDF fund, the fund that MPs are given to reduce poverty in their constituencies. The audit found that Onyango spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.</p> <p>John Kamau has never used violence to try winning elections. We do not have any information about how Kamau used the CDF back when he was an elected MP because he, as the majority of other MPs, was not audited.</p>

8	<p>John Onyango is rumoured to have ordered one murder and hired gangs during the 2007 elections. He has not been arrested for these alleged crimes. Back when John Onyango was an elected MP, he was among the few MPs that were randomly and independently audited about how MPs used the CDF fund, the fund that MPs are given to reduce poverty in their constituencies. The audit found that Onyango spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.</p> <p>John Kamau is also rumoured to have ordered one murder and hired gangs in the 2007 elections. He has not been arrested either for these alleged crimes. We do not have any information about how Kamau used the CDF back when he was an elected MP because he, like the majority of other MPs, was not audited.</p>
---	---

9	<p>John Onyango and John Kamau are both rumoured to have ordered one murder and to have both hired gangs in the 2007 elections. They have not been arrested for these alleged crimes. Back when Onyango and Kamau were elected MPs, both of them were among the few MPs that got randomly and independently audited about how they managed their CDF funds, the fund that MPs are given to reduce poverty in their constituencies. The audits found they both spent the CDF fund for its intended purpose, to tackle poverty, and used it for good quality projects.</p>
---	--

2. With this new information, I am going to ask you to secretly cast your vote in this ballot again for your preferred candidate between these two. Again, your vote will be placed in this bag and be mixed with other ballots. Your vote is private, we will respect your confidentiality and no-one will identify your name with your vote. As before you can tick in this ballot for your preferred candidate between these two.

A	Refused to vote, respondent said was indifferent between candidates > GO TO Q3	5
B	Refused to vote without giving any reason > GO TO Q3	6
OPTIONS C, D, E, F WILL BE FILLED BY ENUMERATOR'S SUPERVISOR		
C	Respondent voted for Candidate 1, John Onyango in the ballot.	1
D	Respondent voted for Candidate 2, John Kamau in the ballot.	2
E	Respondent voted for both candidates in the ballot .	3
F	Respondent did not vote for either of candidates, left voting ballot in blank.	4

We have finished playing the election game. Thank you!

Appendix C

Estimating the effects of candidates' attributes

Voting experiments seek to examine whether showing voters one set of candidate's attributes as opposed to another would change the respondents' voting choice. By randomly varying aspects of the candidates' attributes, such as their performance record and rumoured use of electoral violence, researchers can attribute any differences in voting choice to the experimental manipulation. The key challenge in experiments that manipulate various treatment dimensions (e.g. various candidates' attributes) is to identify the causal effects of the various treatment components. Traditional experimental designs typically can only identify the causal effect of the experimental manipulation as a whole but are unable to decompose the treatment effect of each of the candidate's attribute shown to voters. Specifically, traditional experimental designs estimate the *average treatment effect* (ATE) but not the composite treatment effects associated with each candidate's attribute. As shown in Equation (C.1), ATE measures the expected difference in responses Y_i (in our case, the voting choice in the first and the second ballot) for two different candidates' profiles sets t_1 and t_2 .

$$\pi(t_2, t_1) = E[Y_i(t_2) - Y_i(t_1)], \quad (\text{C.1})$$

Experiments offering complex treatments, like ours, face the challenge of deriving meaningful interpretations from ATE. Since voters have been asked to compare pairs of candidates that have complex profiles, consisting of a large number of attributes with multiple levels, it is difficult to directly interpret the overall effect estimated. In practical terms, what does it mean to compare two profiles of candidates, t_1 and t_2 , whose attributes differ on multiple dimensions?

To be able to estimate the causal effect of many treatment components simultaneously, Hainmueller, Hopkins & Yamamoto (2014) propose using conjoint analysis and identify the individual treatment component, which they refer as the *average marginal component effect* (AMCE). The AMCE represents the change in the population probability that a profile would be chosen if the value of its l th component were changed from t_1 and t_2 , averaged over all the possible values of the other components given the joint distribution of the profile attributes. That is, AMCE represents the marginal effect of attribute l averaged over the joint distribution of the remaining attributes, as shown in Equation (C.2).²

² That is the AMCE measures the expected change in the probability of a candidate being elected, when a given candidate's attribute (e.g. a candidate rumoured to have used violence) is compared to the rest of research-chosen baseline candidate's attributes.

$$\hat{\pi}(t_2, t_1, p(\mathbf{t})) = \sum_{(t, \mathbf{t}) \in \tilde{T}} \{E[Y_{ijk} | T_{ijkl} = t_2, T_{ijk[-l]} = t, \mathbf{T}_{i[-j]k} = \mathbf{t}] - E[Y_{ijk} | T_{ijkl} = t_1, T_{ijk[-l]} = t, \mathbf{T}_{i[-j]k} = \mathbf{t}]\} \times p(T_{ijk[-l]} = t, \mathbf{T}_{i[-j]k} = \mathbf{t} | (T_{ijk[-l]}, \mathbf{T}_{i[-j]k}) \in \tilde{T}),$$

(C.2)

where $\hat{\pi}$ denotes AMCE as a function of the distribution of the treatment components, $p(\mathbf{t})$. Y_{ijk} denotes a vector of potential outcomes for respondent i in her k th choice task that would be observed when the respondent receives the sequence of profile attributes represented by t . T_{ijk} represents the treatment given to respondent i as the j th profile in her k th choice task, whose l th component T_{ijkl} corresponds to the l th attribute of the profile.

Hainmueller, Hopkins & Yamamoto (2014) make three simplifying to estimate the AMCE, assumptions that can be either guaranteed to hold by the design of the experiment or partially tested. These assumptions are that the treatments have been randomised, there are no carryover effects, and there are no profile-order effects. The assumption of no carryover effects suggests that potential outcomes always take the same value as long as voters receive the same candidate's profiles and the same choice tasks. The assumption of no profile-order effects implies that whenever estimating causal effects, researchers can ignore the order in which the profiles were presented to respondents and pool the information across profiles. Indeed, it is common in conjoint experiments to present respondents with several pairs of candidates, where respondents are asked which candidate they would prefer in each of the multiple pairs of candidates shown. These responses are then pooled to estimate the respective AMCE, choosing a relevant reference category.

We acknowledge that it is perhaps possible that in experiments like ours where some candidates have some socially undesirable attributes (such as rumoured use of electoral violence), respondents could face a 'psychological cost' of changing their voting choice once committing to a specific candidate in the first round of voting. To reduce such potential psychological cost, by design, our experiment first presented to all respondents the same pair of candidates mentioning only their name and political affiliation. Only after respondents cast their first ballot for their preferred candidate, we revealed further information about the candidate's rumoured use of electoral violence and performance record.

Estimating the AMCEs

Another advantage of the conjoint analysis proposed by Hainmueller, Hopkins & Yamamoto (2014) is that it is possible to estimate AMCE using a linear regression. As shown in Equation (C.3), the dependent variable, the voting choice made in each of the two ballots is regressed on a set of dummy variables indicating the various levels of each of the candidates' attributes and the interaction terms between attributes (omitting the respective researcher-chosen reference categories to avoid multicollinearity). By including in this linear regression the interactions

with any other linked attributes, it is also possible to estimate the causal interaction effects between attributes (the so-called *average component interaction effect*).

$$Y_{ijk} = \beta_0 + \beta_1 I_{ijk} + \beta_2 V_{ijk} + \beta_3 P_{ijk} + \beta_4 V_{ijk} * P_{ijk} + \beta_5 X_{ijk} + \varepsilon_{ijk} \quad (C.3)$$

where Y_{ijk} represents whether the respondent voted for a particular candidate in each of the two ballots or not. I, V, P are a set of dummy variables representing the candidate's attributes shown in the first and second ballot.

In Equation (C.3) the reference categories are the baseline attributes shown in the first ballot, where there was no mention of candidates' rumoured use of electoral violence or their performance record. Thus, the estimated AMCEs represent the change in probability of a respondent selecting a candidate when a given attribute value is compared to the attribute considered in the reference category. In other words, the AMCEs measures to what extent showing voters one set of candidate's attributes in the second ballot as opposed to the one shown in the first ballot changed the respondents' voting choice.

The AMCEs for each of the attribute's levels is obtained as the weighted average of the regression coefficient of the attribute's level and its interaction regression coefficients with other linked attributes. In experiments like ours that do not fully develop all the potential combinations of the various levels of attributes, the estimation of the AMCE must exclude interactions between attributes that do not exist in the experiment, and effects that are deemed implausible.

Specifically, in our experiment, the candidates' attributes on violence V and performance P have three levels each. Table II illustrates all the combination of attributes shown to respondents in the first and second ballot. In Table C.1 below, we summarise the combination of attributes, showing the restrictions imposed on the experiment, in other words, the combinations of candidates' attributes that we excluded. Specifically, we excluded three possible combinations between the attributes of violence and performance ($V1*P2$, $V1*P3$, $V2*P3$) as our research was not interested in assessing exclusively the effects of presenting candidates only with positive characteristics. Instead, we are interested in closely resembling voter decision making in settings where competing candidates have a combination of good (performance) and some undesirable attributes (potential use of electoral violence). The combination of candidates' attributes presented therefore allow us to observe the trade-offs that voters must make when choosing between candidates that are theoretically relevant.

Table C.1 The combination of candidates' attributes on violence and performance considered in the experiment.

	Level P1: No mention of performance	Level P2: Performance unknown, has not been audited	Level P3: Good performance
Level V1: No mention of whether has used electoral violence or not	included		
Level V2: Has not used electoral violence	included	included	
Level V3: Rumoured to have used electoral violence	included	included	included

The respective AMCE shown in this article were obtained using the regression coefficients from the regression expressed in Equation (C.3) and estimated as described in the Equations (C.4)–(C.7) shown below. To facilitate showing the restrictions applied and weights used, we use a simplified notation. In Equations below, we merely omit showing the respective regression coefficients associated with each attribute of violence ($V1$, $V2$, $V3$), and performance ($P1$, $P2$, $P3$). This simplified notation might also help reading the accompanying replication code.

-The effect of a candidate with the second level of the attribute of violence, $V2$, (the candidate has never used electoral violence) compared to a candidate with the first level attribute of violence $V1$, the baseline category (no mention on whether the candidate has used or not electoral violence) is defined in Equation (C.4) shown below. This AMCE estimates the expected change in probability of voting for a candidate with the attribute $V2$ when compared to a candidate with the attribute $V1$.

$$V2 + [V2 + V2*P2]*0 + [V2 + V2*P3]*0 \quad (C.4)$$

The last two elements take the weight of zero since their effect with respect to the baseline category $V1$ is deemed implausible, since the experiment does not include candidates with the linked characteristics $V1*P2$ nor $V1*P3$.

- The effect of a candidate with the third level of the attribute of violence, $V3$, (the candidate is rumoured to have used violence) compared to a candidate with the first level attribute of violence $V1$, the baseline category (no mention on whether the candidate has used or not electoral violence) is defined as:

$$V3 + [V3 + V3*P2]*0 + [V3 + V3*P3]*0 \quad (C.5)$$

The last two elements take the weight of zero since their effect with respect to the first level $V1$ is rendered implausible, since the experiment did not include candidates with the linked characteristics $V1*P2$ nor $V1*P3$.

-The effect of a candidate with the second level of the attribute of performance, $P2$, (the candidate's performance is unknown as has not been audited) compared to a candidate with the

first level attribute of performance $P1$, the baseline category (no mention on whether the candidate has good or unknown performance) is defined as:

$$P2*0 + [P2 + V2*P2]*\frac{1}{2} + [P2 + V3*P2]*\frac{1}{2} \quad (C.6)$$

The first element $P2$ takes the weight of zero since the experiment did not include candidates with the linked characteristics $V1*P2$. The other two elements receive a weight of $\frac{1}{2}$ which equals the corresponding cell probabilities.

-The effect of a candidate with the third level of the attribute of performance, $P3$, (the candidate has good performance record) compared to a candidate with the first level attribute of performance $P1$, the baseline category (no mention on whether the candidate has good or unknown performance) is defined as:

$$P3*0 + [P3+ V2*P3]*0 + [P3 + V3*P3] \quad (C.7)$$

The first element $P3$ takes the weight of zero since the experiment did not include candidates with the linked characteristics $V1*P3$. The second element has the weight of zero since the experiment did not include candidates with the baseline of the linked characteristics $V2*P1$.

We also ran a separate regression specification (reported as Equation (2) in the article) where we added the interaction with a dummy variable, R , indicating whether the respondent has the same ethnicity/political alliance as the candidate. In this second specification, the AMCEs were obtained as described in Equations (C.8)–(C.11) shown below.

-The effect of a candidate with the second level of the attribute of violence, $V2$, compared to a candidate with the first level attribute of violence $V1$, the baseline category is defined as:

$$V2 + V2*R + [V2 + V2*R + V2*P2+V2*P2*R]*0 + [V2 + V2*R + V2*P3+V2*P3*R]*0 \quad (C.8)$$

-The effect of a candidate with the third level of the attribute of violence, $V3$, compared to a candidate with the first level attribute of violence $V1$, the baseline category is defined as:

$$V3 + V3* R + [V3+V3* R + V3*P2+V3*P2*R]*0 + [V3+V3* R+V3*P3+V3*P3*R]*0 \quad (C.9)$$

-The effect of a candidate with the second level of the attribute of performance, $P2$, compared to a candidate with the first level attribute of performance $P1$, the baseline category is defined as:

$$[P2+P2*R]*0 + [P2+P2*R+V2*P2+V2*P2*R]* \frac{1}{2} + [P2+P2*R+V3*P2+V3*P2*R]* \frac{1}{2}$$

(C.10)

-The effect of a candidate with the third level of the attribute of performance, $P3$, compared to a candidate with the first level attribute of performance $P1$, the baseline category is defined as:

$$[P3 + P3*R]*0 + [P3 + P3*R+V2*P3+V2*P3*R]*0 + [P3 + P3*R+V3*P3+V3*P3*R]$$

(C.11)

Estimating the differences of AMCE between sub-groups

In conjoint analysis, the AMCEs are also commonly estimated across sub-samples of the population by splitting the sample into two or more.³ Leeper, Hobolt & Tilley (2019) explain that whenever comparing the effect size estimated in these sub-samples, any difference in AMCE simply reflects that there is a difference in effect size, which is distinct from a potential difference in absolute level of preferences between groups. That is, conjoint analysis of subgroups should be interpreted merely as the difference in casual effects size. Although differences in AMCEs provide important insights into the descriptive variation in preference within-groups, and the differences in the causal effect size of different attributes, the AMCEs do not provide information about the absolute levels of favourability toward profiles with each feature.

Leeper, Hobolt & Tilley (2019) also explain that AMCEs can be sensitive to the reference categories chosen and that it is misleading to compare AMCEs between subgroups whenever each subgroup uses a different reference category. For this reason, across our conjoint study and subgroup analysis, we use the same reference category, that is the candidates' profile shown in the first ballot. Moreover, in the article, we also explained the relevance of our chosen baseline category. That is, the ethnic and political affiliation profile of the two candidates closely resembles the type of candidates that voters face in real elections. As our descriptive analysis shows, (e.g. Table A.2), the voting choice observed in the first ballot also closely captures the voting split typically displayed in real Kenyan elections.

³ For instance, in our case, we have split the sample into those who have been victims of violence and those who have not, as well as wealthier and poorer respondents. For each of these sub-groups we have estimated the AMCEs.

Appendix D

Taking into account both ballots, 83% of respondents cast their ballots. Like in real elections, among those who did cast a ballot a small minority submitted invalid ballots (by leaving them blank 6% or selecting both candidates 1%), and another 17% of respondents did not cast a ballot at all.⁴

Although some respondents chose not to cast (valid) ballots, our experiment does not suffer from attrition given that all respondents remained cooperative and completed the rest of the questionnaire. More accurately, the experiment suffers from missing survey responses as some respondents chose not to cast valid ballots.

In the conjoint choice analysis presented thus far, we have focused only on the respondents that cast valid ballots. This analysis would be unbiased, provided that the missing (valid) ballots are missing at random and there are no effects on the parameter estimates from ignoring these missing responses. However, we could face sample selection bias if the respondents who chose not to cast a valid ballot were systematically different to those who did cast a valid ballot. If that is the case, we might face bias due to non-random sample selection, as some respondents might have elected themselves out of casting a valid ballot perhaps due to their characteristics or the attributes of the candidates presented.

In this appendix, we test whether responses are missing at random by estimating a probability model in which the dependent variable takes the value of one for respondents that chose not to cast a valid ballot and zero otherwise. In this regression we include as covariates the candidates' attributes in each of the two ballots (our treatments) as well as the respondents' characteristics. Note that in this regression analysis we are not estimating the average marginal component effects of the candidate's attributes. Simply, we are testing whether the candidates' and respondent's characteristics are associated with the probability of respondents refusing to cast a valid vote.

In Figure D.1 we present the effect of the respondents' and candidates' characteristics on the probability of respondents refusing to cast a valid ballot. The dots in this figure represent the size of the effect and the lines the 95% confidence intervals. Note that the attributes of the candidates (our treatments) do not affect the probability of respondents refusing to cast a valid ballot. But, we find some small biases statistically significant at 10% level in terms of the respondent's age, education level, and whether the respondent is poor or not (the respondent's

⁴ As Table A.6 shows, about 8% of respondents told the interviewer that were indifferent between the two candidates thus did not wish to submit a ballot, and a further 8% of respondents gave no explanation as why they refused to submit a ballot.

overall wealth index, whether the respondent has experienced food shortages within the past year, and whether the respondent is categorised as poor given that her/his household has four or fewer assets from a list of 15 durable assets).

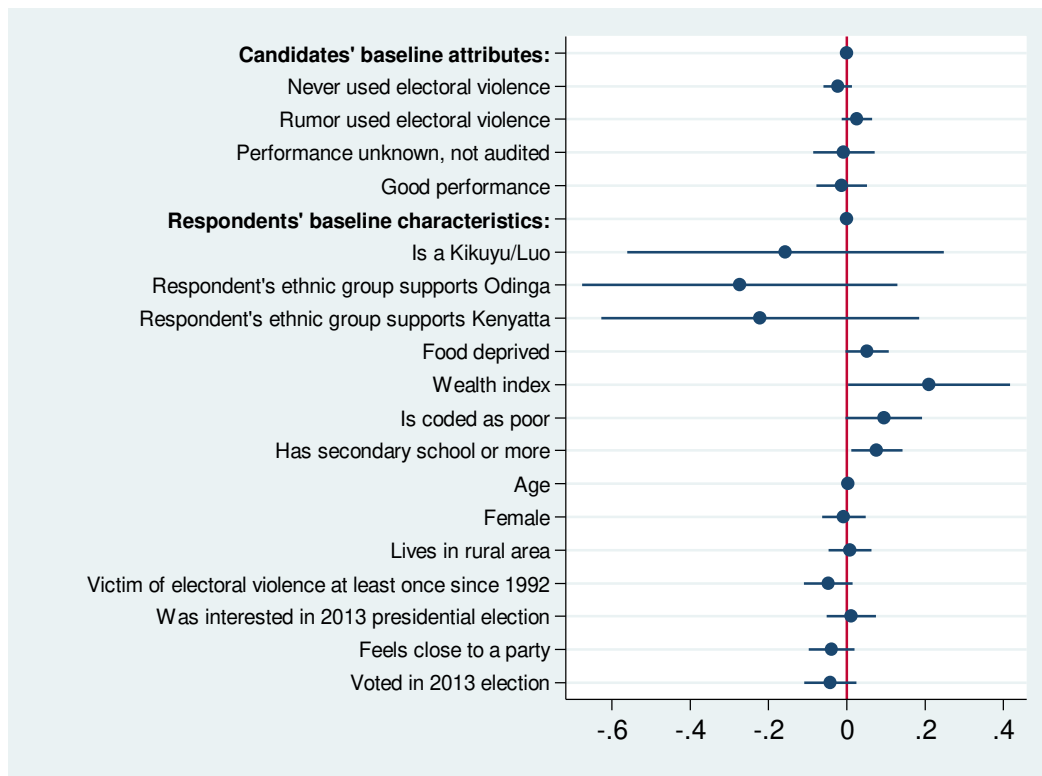


Figure D.1 Effects of candidate's attributes and respondents' characteristics on the probability of the respondent refusing to cast a valid ballot

The possibility of self-selection bias

Figure D.1 suggests that missing responses are not entirely random. In particular, poorer respondents were more likely to have refused to cast a valid ballot. Thus, we use the Heckman two-step estimator model to *detect* and *correct* if necessary for a potential self-selection bias due to missing responses (Heckman, 1979).⁵ This method is estimated in two stages. The first stage estimates the observed and unobserved factors that affected non-response. In the second

⁵ Another method that is commonly used to deal with missing responses is to re-weight the outcome regression results (in our case voting choice) to take into account the probability associated with missing responses. This inverse probability method has the major disadvantage of being valid only if the missing responses are due entirely on observable characteristics. If the missing responses (or attrition) are based on both observable and unobservable characteristics, as is often the case, then the two-stage Heckman method employed here would be more suitable.

stage, the Heckman estimator controls for the effect of non-random selection in the outcome of interest (respondent's voting choice).

Specifically, in the first step, we estimate the probability of a respondent refusing to cast a valid vote, as shown in Equation (D.1). This regression, known as the Heckman selection, includes as covariates the candidate's attributes on identity, violence and performance (I , V , and P), the respondents', interviewers' and area characteristics, denoted by vector X . This Heckman selection regression must also contain at least one variable, so-called- external instrument Z , to safeguard against a potential collinearity between the selection regression and the voting choice regression to be used in the second stage. The instrument Z should influence whether respondents cast a valid vote, but should not affect how respondents voted (so it can be excluded from the voting choice model regression).

As external instrument Z we include the interviewers' identification numbers, which were randomly assigned by our field team. We use these identification numbers as interviewers' fixed effects, which help us to assess whether a respondent's abstention from casting a valid vote is associated with any other interviewers' characteristics, other than the interviewer's years of experience and mother tongue (ethnicity) already included in X , that might have increased respondent's willingness to participate in the voting game yet unlikely to have affected how the respondent voted.

$$refused_{ijk} = \gamma_0 + \gamma_1 I_{ijk} + \gamma_2 V_{ijk} + \gamma_3 P_{ijk} + \gamma_4 V_{ijk} * P_{ijk} + \gamma_5 X_{ijk} + \gamma_6 Z_{ijk} + u_{ijk} \quad (D.1)$$

Refused indicates if the respondent i cast a valid vote for one of the two candidates j (in either of the two ballots k) or not. The vector X includes respondents' ethnicity, sex, age, education, province of residence, rural residence, whether the respondent feels close to a party, whether the respondent claimed to have voted in the 2013 presidential election, has experienced food shortages within the past year, and has experienced electoral violence since 1992. We also included the district index of ethnolinguistic fractionalisation and polarisation indices and the interviewers' mother tongue and years of experience in the field. The residual is denoted by u . This residual is used to construct a new variable called the inverse mills ratio to be used in the second stage of the Heckman method.⁶

In the second stage of the Heckman method, we estimate the probability of whether respondents voted for one of the candidates, as shown in Equation (D.2). In this regression we analyse the voting choices made in the two ballots together. Specifically, the covariates

⁶ This variable is constructed as the ratio of the standard normal probability density function over the cumulative density function of the residual for each respondent.

included in this second stage regression are again the candidate's attributes on identity, violence and performance (I , V , and P) and the external instrument Z . This second-stage of the Heckman method also adds the inverse mills ratio $\lambda(Z_i\gamma)$ as an additional explanatory variable. Thus, this Heckman regression corrects for the potential bias in self-selection by incorporating both the observed and unobserved characteristics that affected non-response in the selection equation. We run this second stage regression with and without adding the vector X to assess to what extent respondents' and interviewers' observed characteristics affected voting choice.

$$choice_{ijk} = \phi_0 + \phi_1 I_{ijk} + \phi_2 V_{ijk} + \phi_3 P_{ijk} + \phi_4 V_{ijk} * P_{ijk} + \phi_5 X_{ijk} + \rho \sigma_u \lambda(Z_i\gamma) + v_{ijk} \quad (D.2)$$

σ_u represents the standard deviation of the residual u . ρ represents the correlation between the unobserved determinants of respondents refusing to cast a valid vote u and v the unobserved determinants of voting for a candidate.

Whenever the regression coefficient ρ is statistically significant, it suggests there is evidence of a significant bias from missing responses in the first selection model.⁷ This result would imply that those who refused to cast valid votes are significantly different from those who cast valid ballots. In this case, the Heckman's two-step estimator would provide unbiased and consistent estimates of the factors affecting the probability of choosing a candidate, by taking into account both the observed and unobserved characteristics that drove non-response.

To simplify the presentation of the Heckman two-stage regressions, in Table D.1 we present exclusively the coefficient ρ and whether it is statistically significant or not. As that table shows, there is no evidence of self-selection when analysing all the sample, or when analysing the Kikuyu or Luo, or those belonging to other ethnicities. There is no evidence of self-selection for the subgroups of victims and non-victims of electoral violence. We only find evidence of self-selection when analysing the Kikuyu and Luo respondents that are classified

⁷ A limitation of the Heckman two-stage method is that if the first selection regression is not well specified, then the variables in this regression will not predict well enough the bias in non-response and the coefficient ρ will have weaker power to detect self-selection bias. Nonetheless, even when the selection equation is not perfectly specified, an alternative way to assess for a potential bias in self-selection is by comparing the predictors of the second-stage Heckman and the regression coefficients of the outcome model when excluding missing responses. We find no major differences when comparing the regression coefficients of the voting choice results presented earlier in the article and the Heckman two-stage estimator. Thus, we conclude that there are no major sample selection biases affecting our results.

as poor. For this sub-group of respondents, as shown in the Figure D.2 once correcting for self-selection bias, our results do not appreciably impact the voting choices presented when excluding the voters that failed to cast a valid ballot. That is, the bias is not sufficient to alter the sign or significance of the AMCEs of each of the candidate’s attributes presented earlier. Once controlling for a likely bias in self-selection of responses, poorer Kikuyu-Luo respondents still fail to sanction candidates rumoured to have used electoral violence, in contrast to the non-poor. Thus, the findings reported earlier on voting choice are a robust description of the drivers of individual vote choice in the experiment.

Table D.1 Likelihood-ratio test of Heckman self-selection

	All respondents				Kikuyu & Luo				Other ethnic groups			
	Basic controls		Extra controls		Basic controls		Extra controls		Basic controls		Extra controls	
	χ^2	p-value	χ^2	p-value	χ^2	p-value	χ^2	p-value	χ^2	p-value	χ^2	p-value
All respondents	0.23	0.63	1.39	0.24	0.77	0.38	0.21	0.64	0.02	0.88	0.18	0.67
Non-victim of electoral violence	1.92	0.17	3.84	0.05	0.65	0.42	0.00	0.95	0.01	0.90	0.02	0.88
Victim of electoral violence	1.40	0.24	1.36	0.24	1.76	0.18	0.09	0.76	0.50	0.48	0.04	0.85
Non-poor	0.00	1.00	0.14	0.71	1.32	0.25	0.19	0.66	0.00	0.95	0.18	0.67
Poor	0.75	0.39	0.06	0.81	340.59	0.00	4339.40	0.00	0.01	0.94	0.01	0.93

The likelihood-ratio test χ^2 tests whether $\rho=0$, where ρ is the correlation between the unobserved determinants of respondents refusing to cast a valid ballot in the experiment u and v the unobserved determinants of voting for a candidate. If this test is statistically significant, it indicates evidence of sample self-selection. The basic controls in the second stage Heckman regression are the candidates’ attributes, and the interviewers’ identification numbers. The extra controls added in the second stage-regression are the respondents’ ethnicity, sex, age, education, province of residence, residency area (urban/rural), whether the respondent feels close to a party, has experienced food shortages within the past year, and has experienced electoral violence since 1992. We also include the ethnolinguistic fragmentation and polarisation indices at the district level and interviewers’ mother tongue, their number of years of experience conducting interviews.

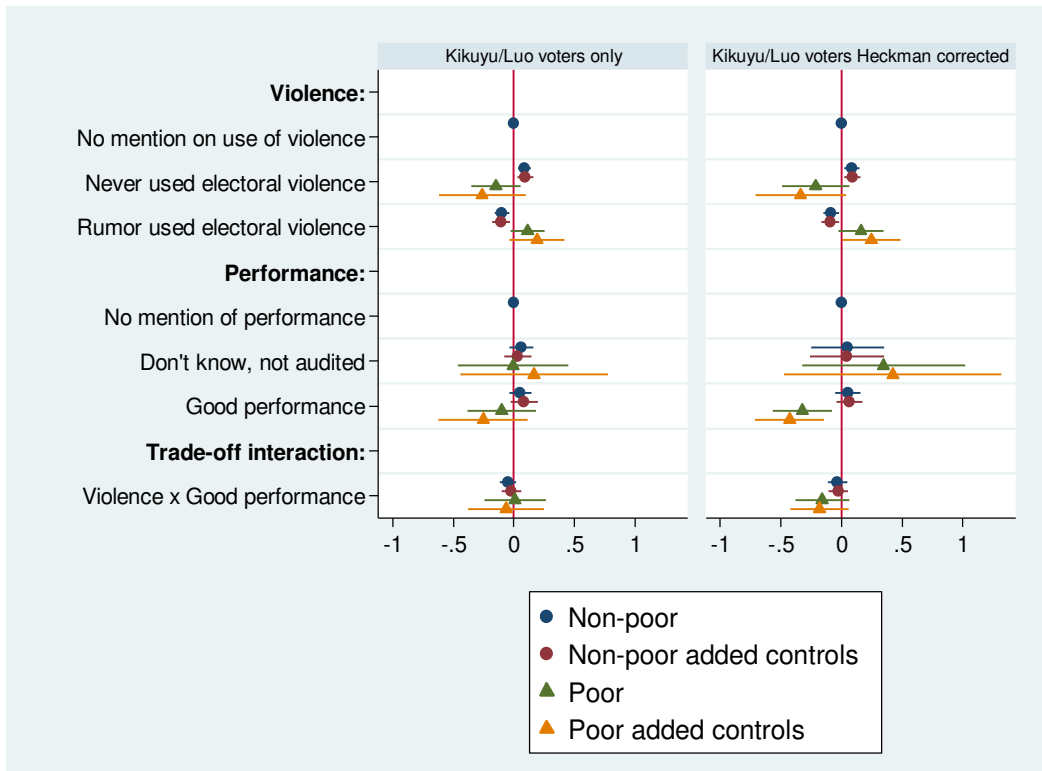


Figure D.2 Effects of candidate's attributes on voting choice by poor and non-poor respondents once controlling for self-selection