

ORIGINAL ARTICLE

# Does vote buying undermine confidence in ballot secrecy? Theory and experimental evidence

Sergio J. Ascencio<sup>1</sup>  and Han Il Chang<sup>2</sup> 

<sup>1</sup>University of Essex, Colchester, United Kingdom and <sup>2</sup>Political Science and International Relations, Kookmin University, Seongbuk-gu, Korea

**Corresponding author:** Han Il Chang; Email: [hanilchang@kookmin.ac.kr](mailto:hanilchang@kookmin.ac.kr)

(Received 12 August 2023; revised 14 March 2024; accepted 25 March 2024)

## Abstract

Why does vote buying persist under the secret ballot? We argue initiating vote-buying transactions allows politicians to undermine voter confidence in the secret ballot, and thus to induce voter compliance. Our analysis consists of three parts. First, we present evidence from a survey experiment in Mexico that shows receiving material goods from a candidate diminishes voter confidence in ballot integrity. Next, we introduce an informational theory of vote buying that explains this phenomenon. Specifically, we develop a model of vote buying as a signaling game, in which a voter who is *ex ante* uncertain about a politician's capacity to monitor voter behavior learns new information from the politician's actions. Finally, we test the key insights from the model in a lab experiment. Our results suggest that, under certain conditions, offering material goods to voters is sufficient to erode their confidence in ballot secrecy, making vote buying effective.

**Keywords:** Comparative politics; political behavior; elections and campaigns; experimental research; voting behavior

Elections are a cornerstone of representative democracy. They allow citizens to make their voices heard, to influence policymaking, and to hold incumbents accountable for their performance. A precondition for elections to fulfill these vital functions is that citizens are free to signify their political preferences (Dahl, 1971). This ideal is violated under several circumstances, such as when politicians are able to monitor voter behavior at the polls. If politicians can observe how voters cast their ballots, they can induce them to vote against their own preferences through the use of coercion, selective incentives, or a combination of both, as in the case of vote buying. Traditionally, the secret ballot is thought to assure that citizens will enjoy this guarantee. Yet, extensive scholarship shows the use of the secret ballot has failed to curb vote buying.

Why does vote buying persist under the secret ballot? This question has generated a vast literature that studies how politicians overcome the commitment problem inherent in vote buying (e.g., Cox and Kousser, 1981; Stokes, 2005; Nichter, 2008; Finan and Schechter, 2012; Rueda, 2017). Implicit in most of these works is one of two assumptions. One is that voters actually believe the ballot is secret, and the other is that this belief is independent of parties' actions. This overlooks the facts that, even in consolidated democracies, a significant portion of the electorate believes parties can monitor their behavior at the polls (Gerber *et al.*, 2012), and that voters' perceptions of electoral integrity depend not only on electoral administration institutions (Rosas, 2010) but also on idiosyncratic factors (Challú *et al.*, 2020).

This paper challenges these assumptions to advance two ideas. First, we argue the occurrence of vote buying itself can undermine voters' perceptions of ballot secrecy. That is, once a broker

offers a bribe to a voter in exchange for his vote, the voter might become less confident that his vote is secret. Second, we argue low trust in ballot secrecy can induce voter compliance with vote-buying exchanges. Therefore, from the perspective of the voter, it is irrelevant whether the ballot is secret or not; as long as the voter *believes* parties can monitor his behavior with high-enough probability, he should be more likely to comply with the broker's request. Taken together, these arguments imply that, by initiating vote-buying exchanges, politicians can erode voters' confidence in the secret ballot, increasing the likelihood that voters will comply with their demands.

We motivate our analysis by showing that there is, in fact, a relationship between the prevalence of vote buying and voters' ballot-secrecy perceptions. Using survey data from the Mexico Panel Study (2012) and the Afrobarometer (Wave 5), we show that voters who have been offered electoral handouts are more likely to believe that politicians can monitor their vote choices. Since this pattern could be driven by reverse causality or omitted variables, we conduct a survey experiment in Mexico, a country where electoral competition has a strong clientelistic component (Aziz *et al.*, 2022). In the experiment, respondents are shown information from a hypothetical election and then asked how confident they are in the secrecy of their vote. One group of respondents is presented with a scenario that involves them being offered a bribe from a broker in exchange for their vote, while the other is presented with a scenario in which the broker does not make such an offer. We find respondents in the "bribe treatment" are significantly more likely to doubt ballot secrecy.

Next, we present an informational theory of vote buying that explains this empirical pattern. We develop a formal model of the interaction between a clientelistic candidate and a voter. In the model, the candidate chooses whether to offer a bribe to the voter in exchange for his vote, and then the voter decides whether to turn out to vote and what candidate to support. The model's main feature is that the voter is *ex ante* uncertain about the candidate's ability to monitor his behavior. We characterize a separating equilibrium in which the candidate offers a bribe to the voter only when the candidate can monitor the voter's behavior. In this equilibrium, receiving a bribe makes the voter more likely to (1) believe the candidate can monitor his vote choice, and (2) cast his ballot in favor of the clientelistic candidate.

Finally, we present results from a lab experiment designed to test the theory's main implications. We find behavior in the laboratory closely matches the behavior in the separating equilibrium: receiving a bribe from a candidate makes voters more likely to both doubt the secrecy of the ballot and vote for that particular candidate. Most importantly, the data indicate that the informational mechanism we propose (i.e., voters adjusting their beliefs about ballot secrecy) has a sizable impact on making vote buying effective, and that this effect is independent from alternative channels studied in the literature.

Our work is part of an extensive literature that studies the prevalence of vote buying in secret-ballot elections. Several authors suggest parties use their extensive knowledge of voters' preferences, socio-economic status, and social networks, to target those who are more likely to comply (Brusco *et al.*, 2004; Stokes, 2005; Bratton, 2008; Cruz, 2019; Mares and Young, 2019), while others argue they exploit this same knowledge to engage in other types of exchanges, such as turnout and abstention buying (Cox and Kousser, 1981; Nichter, 2008; Morgan and Vardy, 2012). Other works identify social norms, such as loyalty and reciprocity, as the mechanism driving the compliance of bribed voters (Finan and Schechter, 2012; Lawson and Greene, 2014). Finally, some propose different outcome-contingent mechanisms in which brokers' access to disaggregated electoral results enable them to enforce vote-buying transactions (Smith and Bueno de Mesquita, 2011; Gingerich and Medina, 2013; Rueda, 2017).

We contribute to this literature by studying an unexplored mechanism that explains why, under certain conditions, vote buying can be effective despite the secret ballot. In our account, receiving a bribe decreases a voter's confidence in the secret ballot. Thus, we join a group of theories that highlight the *informational role* of vote buying (Muñoz, 2014; Kramon, 2016). In contrast to these works, which argue electoral handouts shape voters' perceptions of candidate

viability, we argue they can also change voters' beliefs about a party's capacity to monitor their behavior at the polls.

This paper is most closely related to work by Ferree and Long (2016), who argue politicians strategically create doubts about ballot secrecy in order to engage in vote buying.<sup>1</sup> Using survey data from Ghana, these authors show that voters' perceptions of ballot secrecy correlate with reports of vote buying and with indicators of parties' campaign intensity. Our work differs from theirs in two respects. First, our main evidence comes from experimental data, which allows us to show the association between vote buying and ballot-secrecy perceptions is causal. Second, we formalize a theoretical mechanism showing that initiating vote-buying exchanges is *sufficient* to decrease voters' confidence in ballot secrecy.

Finally, our work also contributes to the literature on voter perceptions of election integrity. Several works describe voter trust in elections as being shaped by institutional factors (e.g., Rosas, 2010), while others find that supporters of winning candidates tend to show higher trust than those who voted for losing candidates (e.g., Cantú and García-Ponce, 2015). We contribute to this literature by showing that vote buying can decrease voter trust in ballot integrity. While previous research finds similar associations between vote buying and voters' electoral trust (e.g., Bratton, 2008; Oliveros, 2019), to the best of our knowledge we are the first to both provide evidence that this relationship is causal and formalize a theory that explains this pattern.

## 1. A first cut: observational evidence

To motivate our analysis, we document a strong positive correlation between citizens' beliefs that politicians can monitor their vote and the likelihood of having received a private transfer (e.g., gift, bribe) during election time using two widely-studied datasets: the Mexico 2012 Panel Study and the Afrobarometer (Wave 5). We summarize our key findings below and present descriptive statistics and regression results in Appendices A and B.

First, data from the Mexico 2012 Panel Study show that, at the 5 percent significance level, respondents who answered "Yes" when asked if anyone had done a favor for them or offered a gift in exchange for their votes during the electoral campaigns were more likely to disagree with the statement "my vote is always secret, unless I tell someone." Figure 1 (top panel) shows the predicted proportion of responses to this statement as a function of the answer to the vote-buying question. Changing the answer from "No" to "Yes" decreases the proportion of respondents who "Agree completely" with the statement, from 0.74 to 0.64. Meanwhile, for the same change of answers, the proportions of respondents who "Agree somewhat", "Disagree somewhat", or "Disagree completely" increase from 0.21 to 0.28, from 0.03 to 0.05, and from 0.02 to 0.03, respectively.

Our analysis of Wave 5 of the Afrobarometer uncovers an analogous pattern (statistically significant at the 1 percent level). The bottom panel of Figure 1 shows predicted perceptions of ballot secrecy by the frequency with which respondents said to have received electoral handouts from politicians. As the frequency changes from "Never" to "Often", the proportion of respondents who said it is "Not at all likely" that powerful people could find out how they voted decreases from 0.66 to 0.53, whereas the proportions of those who said it is "Not very likely", "Somewhat likely" or "Very likely" increase from 0.17 to 0.21, from 0.11 to 0.17, and from 0.06 to 0.09, respectively.

## 2. A survey experiment in Mexico

While the patterns reported above are consistent with our claim that vote buying undermines voter confidence in ballot secrecy, interpreting this evidence as causal is not automatic. Two

<sup>1</sup>Cruz (2019) presents anecdotal evidence of politicians using several tactics to erode voter trust in ballot secrecy.

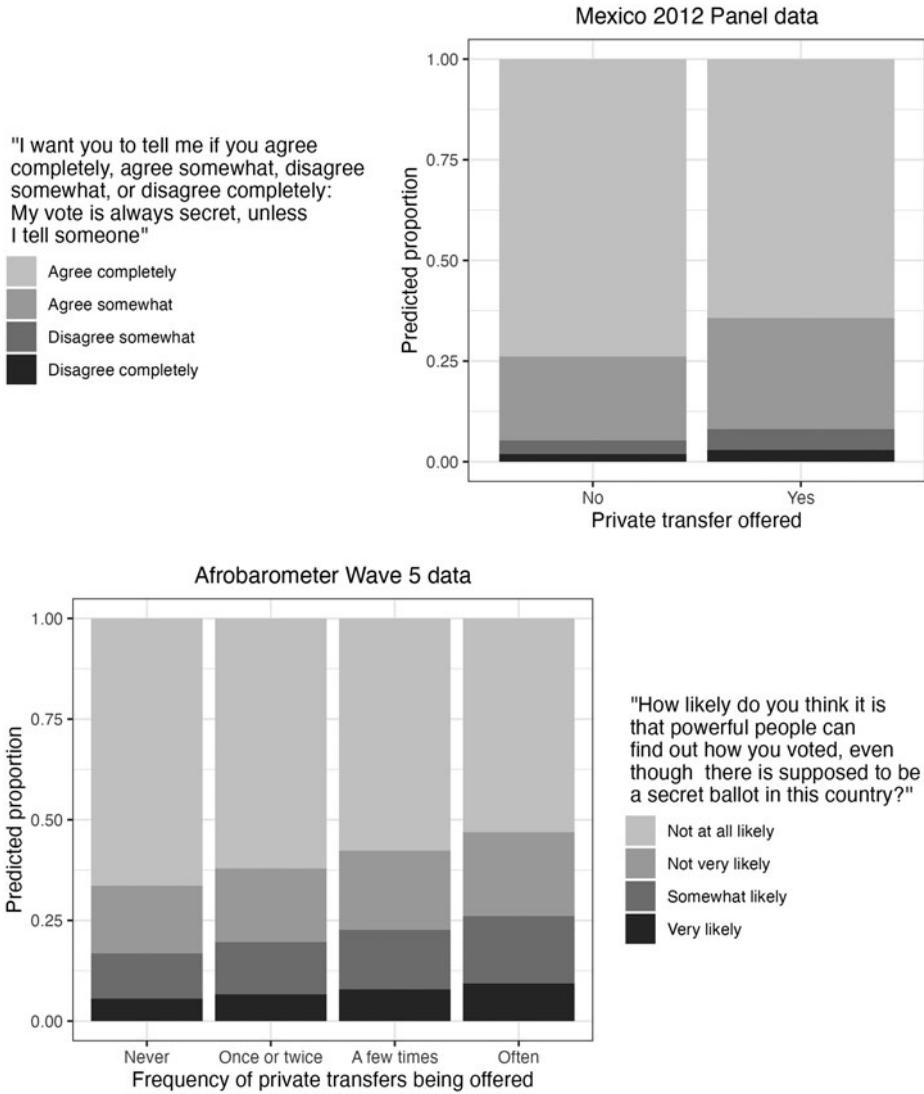


Figure 1. Perceptions of ballot secrecy by experience with vote buying. Note: Proportions are computed using coefficients of models reported in Appendices A and B.

threats to this interpretation stand out. First, this association could be driven by omitted confounders, e.g., weak rule of law could both erode voters' perceptions of electoral integrity and facilitate parties' use of illicit electoral tactics.<sup>2</sup> The second threat is reverse causality, i.e., parties may target voters who doubt their vote is secret. We overcome these challenges by conducting a survey experiment.

2.1 Design

We recruit a national sample of Mexican citizens via Qualtrics. The average age, household monthly income, and education level of the 1043 participants are 39 years old, about 12500

<sup>2</sup>Appendix B.4 shows experiencing vote buying is indeed strongly correlated with institutional-trust measures.

Downloaded from https://www.cambridge.org/core. IP address: 86.30.115.167, on 12 Feb 2026 at 18:49:11, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/psrm.2024.54

pesos, and the second year of high school, respectively; and 50.8 percent of the participants are male. For additional information, including descriptive statistics, see Appendix C.

Upon answering a common questionnaire in Spanish, (reported in Appendix C.1), participants were randomly assigned to one of two—baseline and manipulation—vignettes that described a hypothetical election. The only difference was that in the manipulation vignette respondents were told that a broker working for a candidate they do not support offered electoral handouts to them,<sup>3</sup> whereas in the baseline vignette the broker did not make such an offer. The English translation of the vignettes is as follows:

- [Baseline] Suppose there will be an election for the Chamber of Deputies and you support one of the top two candidates in your constituency. **You are not offered any material goods from a broker working for your candidate's rival.**
- [Manipulation] Suppose there will be an election for the Chamber of Deputies and you support one of the top two candidates in your constituency. **You are offered some material goods from a broker working for your candidate's rival.**

Participants were then asked: “If the broker wanted to find out who you voted for, how likely is it that he can actually find out?” Responses were recorded using a 4-point scale, from “Not likely at all” (0) to “Very likely” (3).

We highlight that the baseline condition explicitly says “You are not offered any material goods from a broker working for your candidate's rival.” Although this sentence could prime respondents to think of vote buying, adding it was necessary to isolate the effect of interest. Specifically, had the baseline vignette not included this sentence, it would have differed from the manipulation vignette along two dimensions: (a) the vote-buying prime, and (b) the private-transfer information (i.e., participants being told that they *did* receive an electoral handout). This would have made it impossible to isolate the effect of the latter, which is our variable of theoretical interest, from the effect of the former. Thus, adding the sentence in question helped us ensure that the results are exclusively driven by the private-transfer information. Moreover, we do not think that mentioning a broker in the baseline vignette is something that Mexican respondents would have perceived as surprising or unusual given the prevalence of vote buying in the country.<sup>4</sup>

## 2.2 Results

We begin by verifying that the random assignment of the experimental vignettes was well implemented. Appendix C.4 reports results of balance tests, which show participants in the baseline and manipulation conditions are statistically indistinguishable along several dimensions.<sup>5</sup>

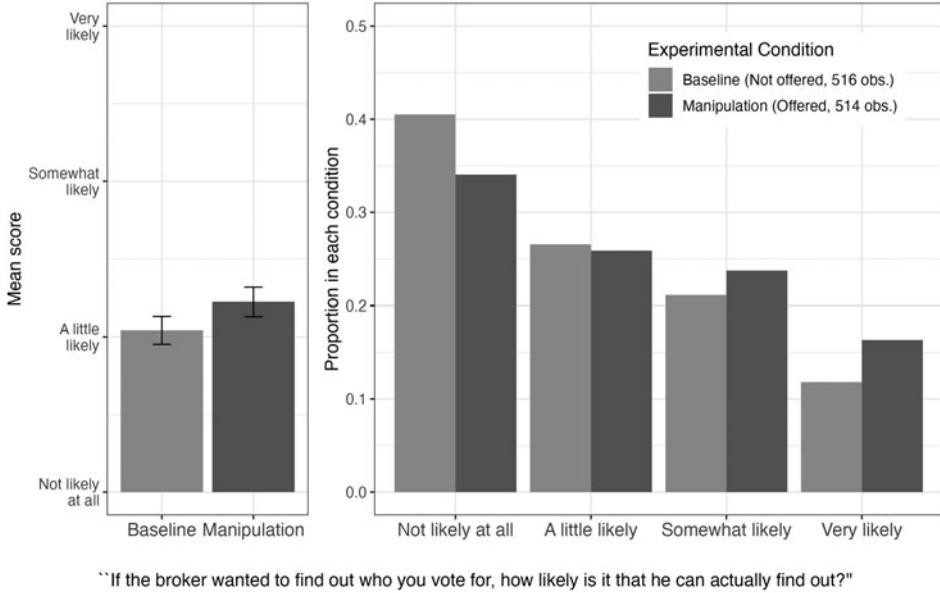
First, we present our main results graphically in Figure 2. The left panel shows that the mean score of the outcome of interest is greater in the manipulation than in the baseline condition (1.224 versus 1.043,  $p$ -value = 0.007). The right panel shows in greater detail how the manipulation vignette increases the share of affirmative answers to the question about the broker's monitoring capacity. In the baseline condition, the shares of participants who chose “Not likely at all”, “A little likely”, “Somewhat likely”, and “Very likely” are 0.405, 0.266, 0.211, and 0.118, respectively. In the manipulation condition, the proportion of participants who chose “Not likely at all” drops by 0.065, whereas for the option “A little likely” this reduction is 0.007. Accordingly, the shares of participants who chose “Somewhat likely” and “Very likely” increased by 0.026 and 0.045, respectively.

Additionally, to assess the statistical significance of the manipulation vignette, we reports results of a series of ordered logistic regression models that gradually include controls from

<sup>3</sup>We use this framing to match what the literature calls *vote buying* rather than *turnout buying* (see Nichter, 2008).

<sup>4</sup>In a nationally representative survey fielded by the national electoral commission in 2020, 79 percent of respondents said they think vote buying is a common practice in Mexico (Aziz et al., 2022).

<sup>5</sup>Appendices C.2-C.3 show details on the covariates used in these tests.



**Figure 2.** Mean score and distribution of beliefs that votes can be monitored by condition, Note: Black vertical lines in left panel indicate 95 percent confidence intervals.

the wide set of variables used in the balance tests. The results, reported in Table 1, show that the manipulation vignette has a positive effect on voter perceptions that their votes can be monitored. Across all models, the coefficient of *Manipulation vignette* is positive and significant at the 1 percent level.

### 3. Vote buying as a signaling game

This section advances an informational theory of vote buying that explains the empirical pattern documented above. Specifically, we develop a model of vote buying as a signaling game in which voters are *ex ante* uncertain about the effectiveness of a candidate’s clientelistic machine. Following previous work, we say a machine is effective or “strong” if it has the capacity to carry out two tasks that are essential to enforce vote buying: (1) monitor voter behavior at the polls, and (2) punish non-compliance.

#### 3.1 Setup

Two candidates, *A* and *B*, run in an election. We model the interaction between *A*, who is a clientelistic candidate (“she”), and an arbitrary voter (“he”), denoted by *V*.<sup>6</sup> Nature begins by drawing *A*’s type as “strong” with probability  $q \in (0, 1)$  and “weak” with probability  $1 - q$ . As mentioned, the difference between these types is that the strong type has the capacity to monitor voter behavior at the polls and punish non-compliance, whereas the weak type lacks these capabilities. Thus, one can interpret these draws as reflecting the effectiveness of the candidate’s brokers or the extensiveness of her clientelistic network in a given community. Candidate *A* observes her own type but *V* only observes the common prior distribution.

The interaction proceeds as follows. First, *A* takes action  $a_A \in \{0, 1\}$  deciding between transferring fixed amount  $t > 0$  to the voter in exchange for his support ( $a_A = 1$ ) or keeping this

<sup>6</sup>This means *B* is not a strategic actor in the model.

**Table 1.** Regression results of the survey experiment data from Mexico

Dependent variable:	Belief that vote can be monitored (0/1/2/3)			
	(1)	(2)	(3)	(4)
Manipulation vignette	0.307** (0.115)	0.352** (0.119)	0.360** (0.120)	0.342** (0.122)
Age		-0.010 (0.006)	-0.011 (0.006)	-0.009 (0.006)
Male		0.030 (0.121)	0.015 (0.122)	0.036 (0.127)
Education		-0.004 (0.093)	0.001 (0.093)	0.005 (0.096)
Income		-0.013 (0.029)	-0.016 (0.029)	-0.012 (0.030)
Any religion		0.101 (0.151)	0.098 (0.152)	0.169 (0.157)
Employment status		0.239 (0.188)	0.231 (0.188)	0.192 (0.193)
Positive reciprocity			-0.027 (0.033)	-0.026 (0.033)
Negative reciprocity			-0.010 (0.080)	-0.013 (0.081)
Confidence on the impact of my vote				0.043 (0.078)
Trust in electoral commission				-0.124 (0.073)
Turnout in 2018 election				-0.322 (0.217)
Offered transfer in 2018 election				0.422** (0.159)
State FE	Included	Included	Included	Included
Observations	1030	974	963	941

Note: Ordered logistic regression models (non-exponentiated coefficients); cutoffs are suppressed in the report; \*  $p < 0.05$ , \*\*  $p < 0.01$ .

amount for herself ( $a_A = 0$ ). That is,  $A$  chooses whether or not to buy  $V$ 's vote at price  $t$ . Next,  $V$  chooses  $a_V \in \{A, B, \emptyset\}$ , indicating whether he abstains ( $a_V = \emptyset$ ), votes for the clientelistic candidate ( $a_V = A$ ), or votes for her opponent ( $a_V = B$ ). Turning out to vote is costly, which means  $V$  pays cost  $c > 0$  unless he abstains.

Payoffs are as follows. First,  $A$  is compensated according to the outcome of the election, and thus her payoff is partially shaped by  $V$ 's action. There is function  $w: \{A, B, \emptyset\} \rightarrow \mathbb{R}_+$  that maps from the voter's set of actions to  $A$ 's payment. To simplify notation, let  $w_{a_V}$  denote  $A$ 's payment when the voter chooses  $a_V$ . Consistent with different broker-compensation schemes described in the literature (e.g., Brierley and Nathan, 2022), we assume  $w_A > w_\emptyset > w_B$ . In words,  $A$  obtains a higher payment as the number of votes for candidate  $A$ , relative to  $B$ , increases.

Following similar vote-buying models,  $V$  derives expressive utility from voting (e.g., Gans-Morse *et al.*, 2014). He receives payoff  $b_i \in \mathbb{R}$  from casting a vote for  $i \in \{A, B\}$  and a payoff of zero from abstaining. Private transfers aside,  $V$  prefers one of the candidates over the other. To facilitate the discussion, we assume  $V$ 's utility of voting for his preferred candidate is equal to the disutility of voting for his non-preferred candidate, i.e., we assume  $b_A = -b_B$ .<sup>7</sup> Finally, if  $V$  fails to vote for  $A$  after the strong type of  $A$  paid  $t$ , he pays sanction  $s > 0$ , which represents a penalty imposed by the candidate.<sup>8</sup> This penalty can be interpreted in several ways. For instance, it could result from the withdrawal of benefits (e.g., Robinson and Verdier, 2013) or reflect the iterated

<sup>7</sup>This assumption is standard in this literature (e.g., Morgan and Vardy, 2012) and is not necessary for our results.

<sup>8</sup>We assume monitoring and sanctioning are costless. This is inconsequential and can be incorporated in the setup. In fact, an interpretation of our setup is that  $A$ 's types face different costs.

Downloaded from https://www.cambridge.org/core. IP address: 86.30.115.167, on 12 Feb 2026 at 18:49:11, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/psrm.2024.54

nature of broker-voter interactions (e.g., Rueda, 2017). While we do not take a stance on the nature of the sanction, we follow these works in assuming that the strong type of *A* can effectively punish voters who renege on their word.

Before presenting our analysis, we highlight that, although we analyze a one-shot exchange between a broker and a voter, several features of the model capture the repeated nature of the interactions between these actors. Most notably, *V*'s prior belief that *A* is the strong type, *q*, can be seen as shaped by past experiences with vote buying. Similarly, the sanction from renegeing, *s*, could capture the loss of future benefits (analogous to punishment strategies in repeated games). Therefore, we believe the model can be interpreted both as a one-shot interaction or as one stage of the ongoing broker-voter relationship.

**3.2. Analysis**

This is a dynamic game of incomplete information, and thus the solution concept we use is Perfect Bayesian equilibrium (henceforth equilibrium). Since our goal is to use the model to analyze how vote buying can undermine voter confidence in ballot secrecy, we focus on an equilibrium in which *vote buying is informative*. Specifically, we search for an equilibrium with two features: (1) there is vote buying, i.e., at least one type of *A* provides the private transfer, and (2) upon receiving the transfer, *V* learns information about *A*'s type. Our main result provides a set of necessary conditions for a *separating equilibrium*, the only equilibrium with these two properties.

We first describe the equilibrium strategies and beliefs, and then discuss our main result. In a separating equilibrium with vote buying, the strong type of *A* offers the transfer to the voter but the weak type does not. The voter's strategy is such that he casts his ballot for *A* after he receives the transfer, and either votes for candidate *B* or abstains otherwise (for details, see Appendix D). Two features of this equilibrium deserve special attention. First, the players' strategies are such that vote buying is effective, i.e., when *V* receives the transfer, he votes for *A*. This occurs because, if vote buying were not effective, the strong type of *A* would be better off not providing the transfer. Second, this equilibrium is fully informative. Upon observing *A*'s action, *V* updates his beliefs about the effectiveness of *A*'s clientelistic machine and, in fact, learns her type.

In the remainder of this section, we characterize the conditions under which this equilibrium emerges. Throughout, we say the voter is a *supporter of  $i \in \{A, B\}$*  if  $b_i > 0$ , and we say he is a *strong supporter of  $i$*  if  $b_i > c$ .<sup>9</sup>

Proposition 1: There exists a separating equilibrium with vote buying only if the following conditions hold:

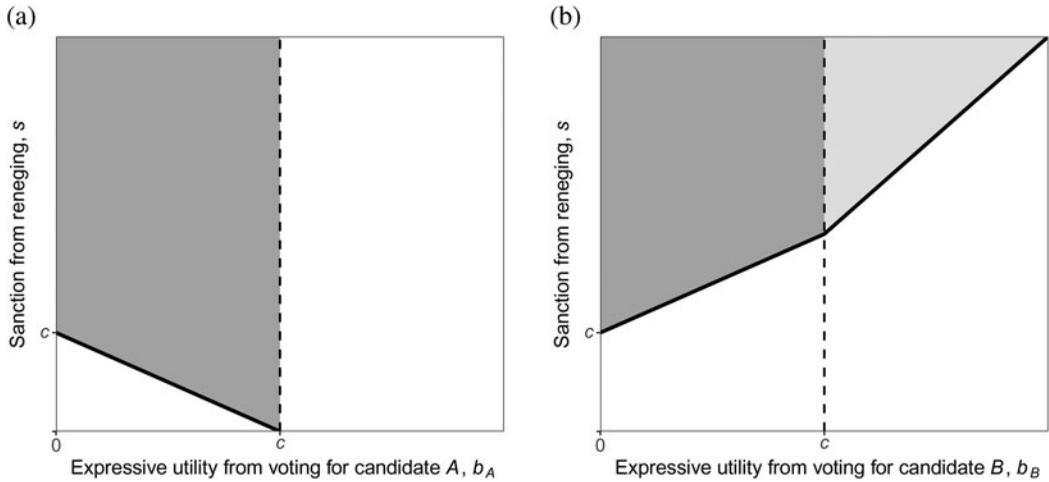
- (1) The voter is not a strong supporter of candidate *A*,
- (2) The sanction from non-compliance is sufficiently high, and
- (3) The private transfer is optimal for *A*.

The proof is in Appendix D. Let us discuss the logic behind this result. Condition (1) identifies the set of potential targets of vote buying. The strong type of *A* only offers the transfer to a voter who would otherwise abstain or vote for *B*. To see why, notice that the behavior of a voter who is a supporter of *i* after he does not receive a transfer is driven by the cost of voting, *c*. The voter strictly prefers to cast his ballot for *i* than to abstain only if he is a *strong supporter of  $i$*  (i.e., only if  $b_i > c$ ). Since a strong supporter of *A* is guaranteed to vote for *A*, whether he receives a transfer or not, *A* is better off not providing the transfer and keeping  $t > 0$  for herself.

Figure 3 illustrates this point. In each panel, the horizontal axis shows *V*'s utility from voting from his preferred candidate; in Panel 3a he is an *A* supporter (i.e.,  $b_A > 0 > b_B$ ) and in Panel 3b he

<sup>9</sup>Since we assumed  $b_A = -b_B$ , *V* can be a supporter of only one candidate.

Downloaded from https://www.cambridge.org/core. IP address: 86.30.115.167, on 12 Feb 2026 at 18:49:11, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/jpsrm.2024.54



**Figure 3.** Parameters that support separating equilibrium described in Proposition 1, (a) Voter is supporter of A ( $b_A > 0$ ), (b) Voter is supporter of B ( $b_B > 0$ ),  
 Note: Shaded areas show parameter values in which the separating equilibrium can emerge.

supports B. The dashed lines show the cost of voting,  $c$ , which drives  $V$ 's behavior when he does not receive the transfer. Candidate A has incentives to provide the transfer to all voters, except those who are strong supporters of A (area to the right of the dashed line in Panel 3a).

Condition (2) indicates that vote buying affects voter behavior. The effectiveness of vote buying depends on the non-compliance sanction,  $s$ . The shaded areas in Figure 3 show the separating equilibrium emerges only if  $V$  is not a strong supporter of A and  $s$  exceeds a minimum threshold, depicted by the solid black line. Among supporters of B (see Panel 3b), this threshold is increasing in  $b_B$ . This reflects the fact that, as  $V$  derives greater utility from voting for B, it is necessary to impose a greater sanction to prevent him from renegeing after receiving  $t$ . In contrast, among A supporters (see Panel 3a), this minimum sanction is decreasing in  $b_A$ , meaning that as the utility of voting for A increases, the smaller the sanction required to induce  $V$  to vote for A. We highlight that the minimum sanction required for this equilibrium to emerge is always lower for A supporters than for B supporters. Thus, unless candidate A can impose a large-enough sanction from renegeing, the model suggests A should target her own supporters.<sup>10</sup>

Finally, condition (3) establishes the optimal size of  $t$ . Candidate A is willing to provide a larger transfer to strong supporters of B than to other voters. This is not to say A chooses an amount to transfer—recall that  $t$  is exogenous. Instead, this means the transfer that is optimal to buy the vote of a strong supporter of B is greater than the one that is optimal for other voters. The two shades in Figure 3 represent the different sizes of the optimal transfer  $t$  in equilibrium. Voters who would abstain in the absence of the transfer receive the same amount, whether they are A or B supporters (darker areas in panels 3a and 3b). Because these voters would not turn out to vote otherwise, from the perspective of A, her transfer changes an abstention into a vote for A.<sup>11</sup> In contrast, when A provides the transfer to a strong supporter of B, who would otherwise vote for B, she is effectively taking a vote away from her rival. Therefore, the optimal transfer for a voter who is a strong supporter of B is greater than the one for other voters (lighter area in Panel 3a).

We conclude by briefly discussing two differences between the separating equilibrium analyzed here and other equilibria with vote buying (see Appendix D). First, the separating

<sup>10</sup>If  $s$  is not sufficiently large (i.e.,  $s \in (0, c)$ ) only A supporters are targeted in this equilibrium. This is consistent with research that argues parties target their own supporters (e.g., Nichter, 2008).

<sup>11</sup>Exchanges with moderate voters resemble what Gans-Morse *et al.* (2014) call turnout buying/double persuasion.

equilibrium is *fully informative*, meaning  $V$  learns  $A$ 's type upon observing her action. In equilibrium, when  $V$  observes  $A$  offering the transfer, he updates his beliefs about the effectiveness of  $A$ 's clientelistic machine; specifically, his belief that  $A$  is the strong type increases from  $q$  to 1. By contrast, if  $A$  does not offer the transfer, the voter's belief that  $A$  is the weak type goes from  $1 - q$  to 1. It should be noted that this dynamic, which highlights the informational role of vote buying, is consistent with the empirical results presented in previous sections of the paper.

Second, the conditions in Proposition 1 do not depend on  $V$ 's prior beliefs about  $A$ 's type,  $q$ , and thus the separating equilibrium can emerge even when  $q$  gets arbitrarily close to zero. Substantively, this means the informational mechanism we propose can operate even in settings in which the *ex ante* likelihood that  $A$  is the strong type, and can thus monitor voter behavior, is very low.

### 3.3. Hypotheses

Based on the separating equilibrium, we derive three hypotheses and test them in a lab, where we artificially create an electoral environment that meets Proposition 1's conditions. Because our hypotheses are stated under the assumption that these conditions are met, they refer to the interaction between a clientelistic candidate and a voter who is not a strong supporter of this candidate.

*Hypothesis 1: The voter is more likely to believe that his vote can be monitored when he receives the private transfer than when he does not.*

To be clear, the results of our survey experiment provide strong support for Hypothesis 1. However, the lab offers an opportunity to test this hypothesis in a different setting, along with the other two hypotheses derived from our theoretical model. The first of these focuses on how receiving the transfer affects the voter's choice.

*Hypothesis 2: The voter is more likely to vote for the clientelistic candidate when he receives the private transfer than when he does not.*

Our final expectation focuses on the candidate's behavior. Typically, given that candidate (or broker) behavior is hard to observe directly, testing hypotheses of this nature would be unfeasible. However, in the lab experiment participants play the role of clientelistic candidates, which allows us to test the following hypothesis:

*Hypothesis 3: The clientelistic candidate is more likely to offer the private transfer when she is the "strong" type (i.e., when she can monitor voter behavior and punish non-compliance) than when she is the "weak" type.*

## 4. Lab experiment

We conduct a lab experiment to evaluate the three hypotheses derived from our theory. As mentioned above, all experimental parameters are set to meet the conditions from Proposition 1's separating equilibrium.<sup>12</sup> We first describe the experimental design and then report the main results.

### 4.1 Design

The experiment consists of several incentivized tasks, which we describe in the order of play at a session. At the start of each session, participants are asked to play the first task, which measures

<sup>12</sup>We highlight that this parameter setup does not prevent the possibility that subjects' behavior in the lab is inconsistent with our theory. For example, subjects could support a candidate who provides an electoral handout out of reciprocity (e.g., Lawson and Greene, 2014), without updating their beliefs about the candidate's monitoring capacity.

their strategic sophistication (Carpenter *et al.*, 2013). The task is a standard  $p$ -beauty contest in which participants, who are randomly grouped, individually choose an integer between 0 and 100. The contest's winner is the participant whose number is closest to  $p$  times the average of all numbers chosen by members of their group (we set  $p$  to  $2/3$ ). The winner is rewarded with 100 experimental tokens (hereafter, tokens). The number chosen by each participant is used as an indicator of their strategic sophistication.

The second task, called the election game, begins without informing participants of whether or not they won the beauty contest. The election game consists of two candidates—one *clientelistic* and the other *programmatic*—and a *voter*. The programmatic candidate does not play a strategic role in the interaction, and is thus played by a computer. The roles of clientelistic candidate and voter are played by participants. Thus, unless there is potential for ambiguity, throughout we refer to the clientelistic candidate as *the candidate*. The strategic environment, described below, is set to meet the conditions in Proposition 1 and is common knowledge to the candidate and the voter.<sup>13</sup>

Initially, the candidate and the voter are endowed with 120 and 40 tokens, respectively. The candidate can be either a weak or a strong type. As before, the difference between these types is that the strong type has the capacity to monitor voter behavior and sanction non-compliance. The candidate is strong with exogenously given probability, which we set to 0.4. At the beginning of the game, the candidate learns his type, but the voter remains uninformed. The interaction proceeds as follows. First, the candidate chooses whether or not to offer a private transfer (40 tokens) to the voter. After observing this action, the voter chooses one option among “voting for *clientelistic* candidate”, “abstention”, and “voting for *programmatic* candidate.” Turning out costs 40 tokens.

The voter's choice affects the probability that the clientelistic candidate wins. This probability is set at 0.6 if he votes for the clientelistic candidate, at 0.4 if he votes for the programmatic candidate, and at 0.5 if he abstains.<sup>14</sup> If the clientelistic candidate wins the election, then he earns 200 tokens, while the voter earns 0 tokens. Otherwise, the clientelistic candidate and the voter receive 0 and 120 tokens, respectively. This payoff structure is such that the voter is what we called a strong supporter of the programmatic candidate.<sup>15</sup> The voter is deducted 40 tokens only if the following conditions are simultaneously met: (1) the candidate offered the 40 tokens to the voter, (2) the voter abstained or voted for the programmatic candidate, and (3) the clientelistic candidate is the strong type. If any of these conditions is not met, the voter is not sanctioned.

After the voter chooses an action, the candidate and the voter receive feedback about both the outcome of the election and their payoffs. Additional feedback, in particular regarding potential punishments for the voter, is provided in a way that reflects how this interaction would operate in the field. The candidate learns the voter's action, and thus whether or not he was punished, only if (i) the candidate offered the 40 tokens to the voter, and (ii) the candidate is the strong type. The voter is never directly informed about the candidate's type, but his payoffs can potentially reveal this information. If the voter received the transfer and then chose an option other than “voting for clientelistic candidate,” the voter pays the sanction of 40 tokens only if the candidate is the strong type. Thus, receiving a punishment unequivocally means that the candidate is strong. In contrast, if the voter receives the transfer and then chooses “voting for clientelistic candidate,” he is never punished, regardless of the candidate's type. In this case, the voter is unable to distinguish if the lack of punishment results from his choice or from the candidate's type.

<sup>13</sup>Participants are informed of the strategic environment before being assigned their roles and playing the game.

<sup>14</sup>The experiment exaggerates voters' tendency to overestimate their votes' impact on election outcomes (Duffy and Tavits, 2008). Including this probabilistic component in the model does not affect the separating equilibrium's substantive features.

<sup>15</sup>Our experimental setup deviates from the theoretical model in that voters do not derive expressive utility from voting. However, we emphasize that the strategic environment used in the experiment is exactly analogous to that introduced in the theory section (and preserves all substantial characteristics of Proposition 1's separating equilibrium); the only difference is the *interpretation* of the voter's payoffs.

Since the candidate's type is determined by chance, the election game is played multiple times. This allows us to collect observations from a diverse range of situations that vary not only by the players' actions but also by whether the candidate has weak or strong monitoring capacity. Before the first round, participants are randomly divided into the roles of either clientelistic candidate or voter. Roles are fixed for the first six rounds and then switched between the sixth and the seventh rounds for the next six rounds. One clientelistic candidate and one voter are randomly and anonymously matched at every round.

Finally, to avoid potential experimenter effects, participants' beliefs about the candidate's monitoring capacity are separately measured in a third task, called the guessing game (or belief-elicitation task), performed after the twelfth round of the election game.<sup>16</sup> The guessing game begins with two participants being randomly selected and asked to, once again, play the election game as the candidate and the voter.<sup>17</sup> The remaining participants are asked to indicate their beliefs about the candidate's monitoring capacity in each of two hypothetical situations—one in which the candidate chooses to offer 40 tokens to the voter and another in which the candidate chooses not to do so (i.e., within-subject design)—using binary options (“No, she/he is unable to monitor” and “Yes, she/he is able to monitor”). Every participant who correctly guesses the candidate's type under the hypothetical scenario that matches the action of the participant playing the role of the candidate receives 320 tokens. To obtain more nuanced measures of the participants' beliefs, respondents are also asked to indicate their confidence in their choices on a scale from 0 to 10, where greater values indicate more confidence. The guessing game is played only once and its result is revealed at the last stage, in which final payments are reported to participants.

Before the payment stage, participants are asked to fill out a short questionnaire including questions regarding age, gender, religion, altruism, and positive reciprocity. We also measure participants' cognitive ability using a three-question cognitive reflection test (CRT) (Frederick, 2005). Appendix E shows the exact wording of the questions. Finally, participants leave the experimental lab with cash that is determined as the sum of earnings from the played tasks; in the case of the second task, the election game, one of the twelve rounds is randomly selected for payment.

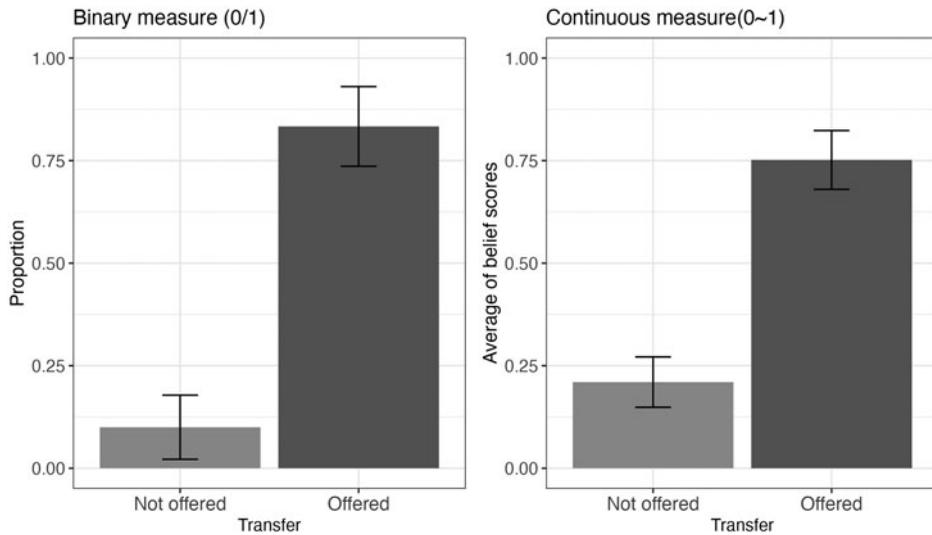
#### 4.2 Main results

The lab experiment was conducted at NYU Abu Dhabi. Instructions on the beauty contest and guessing game were provided on computer screens, while those of the election game were provided on paper (see Appendix E.1). Overall, 70 undergraduate students participated in one of five sessions. Observations for the guessing game, used to elicit beliefs about the candidate's type, were collected from only 60 participants, since two participants are excluded from this task in each session. The average payment was \$ 29.41.

We first report results from the guessing game. As mentioned, we elicited the beliefs of 60 participants regarding the candidate's monitoring capacity, once under the hypothetical scenario that the candidate offered a transfer to the voter and once under the scenario that the candidate did not offer a transfer. In total, then, this game generated 120 ( $= 60 \times 2$ ) observations. Figure 4 (left panel) shows the share of participants who believe the candidate can monitor voter behavior in each scenario. The data strongly support Hypothesis 1: only 10 percent of participants believe the candidate has strong monitoring capacity under the scenario in which the candidate did not offer the transfer, whereas 83 percent believe so under the scenario in which the candidate did offer the transfer. Disaggregating these data shows that 78 percent of participants believed the candidate had strong monitoring capacity only under the scenario in which the candidate offered

<sup>16</sup>If asked about their beliefs during the election game, participants may learn that researchers are interested in how their beliefs affect their behavior and change their behavior in response.

<sup>17</sup>This setup maximizes the number of observations on the beliefs of the remaining participants.



**Figure 4.** Beliefs about candidate's monitoring capacity by hypothetical scenario, Note: The left and right panels show the share of participants who believe the candidate has strong monitoring capacity and the averages of belief scores measured along a continuous scale, respectively. The black lines are 95 percent confidence intervals.

the transfer, which is exactly as expected in Hypothesis 1. The breakdown of the data for the rest of the participants is as follows: 12 percent (5 percent) believed the candidate had strong (weak) monitoring capacity under both hypothetical scenarios, and the remaining 5 percent believed the candidate had strong monitoring capacity only under the scenario in which she *did not* offer the transfer to the voter.

In the guessing game, participants were also asked to indicate their confidence in their choices. We use these reported confidence levels to measure the participants' beliefs along a continuous scale between 0 and 1. This scale assumes that a participant who is completely uncertain about the candidate's monitoring capacity chooses one of the two options with probability 0.5. Under each hypothetical scenario, a participant's score is calculated as  $\frac{100 - 5 \times (10 - \text{Confidence})}{100}$  if the respondent said they believe the candidate is able to monitor, and as  $\frac{5 \times (10 - \text{Confidence})}{100}$  if they believe the candidate is unable to monitor. The average beliefs using this continuous scale, reported in the right panel of Figure 4, are 0.75 (s.d.: 0.28) and 0.21 (s.d.: 0.24) under the scenario that the transfer was and was not offered, respectively.<sup>18</sup>

To provide a more rigorous test of Hypothesis 1, we conduct regression analyses in which our beliefs measures are the dependent variables. Table 2 reports the results of these models, which include different sets of controls. Importantly, we control for the participant's cognitive ability by including *Choice in beauty contest* and *Number of correct answers in CRT*, which were constructed using responses to the beauty contest game and the CRT (for details see Appendices E.3-E.4). Additionally, we include *Role in first six rounds* to control for ordering effects associated with the role participants played first in the election game.

The estimates in Table 2 confirm that participants' assessment of the candidate's monitoring capacity is stronger under the scenario in which the candidate offers the transfer. In Model 3, the

<sup>18</sup>In the election game, the *ex ante* probability that the candidate has strong capacity is 0.4, so participants who were completely unsure, might have guessed accordingly. We create an alternative continuous scale that assumes respondents choose the option "Yes, she/he is able to monitor" with probability 0.4 when they are completely uncertain. Replicating Figure 4 and Table 2 using this measure yields substantively identical results (see Appendices E.2 and E.5).

coefficient of *Hypothetical scenario of transfer provision* indicates that the likelihood a respondent believes the candidate has strong capacity increases from 0.06 to 0.80 when the hypothetical scenario changes from the one where the candidate does not offer the transfer to the one where she does. The results using the continuous measure of beliefs are equivalent. The coefficient of interest in Model 6 reveals a similar gap in continuous beliefs across scenarios (from 0.17 to 0.72).

Next, we turn to the remaining expectations. To reiterate, Hypothesis 2 establishes the voter should be more likely to vote for the clientelistic candidate when he receives the private transfer, while Hypothesis 3 states the clientelistic candidate should be more likely to offer the private transfer when she has strong monitoring capacity. We test these hypotheses by analyzing the behavior of participants in the election game. Recall that each of our 70 participants played six rounds of this game *in each role* (voter or candidate), for a total of 12 rounds. The unit of observation in the analyses that follow is the participant-round. We use random-effects models to address the fact that participants were repeatedly exposed to the same strategic environment across rounds. The models also control for learning effects with a round indicator, *Round*.

**Table 2.** Regression results regarding belief updating

	Dependent variable:					
	Binary belief (0/1)			Continuous belief (0~1)		
	(1)	(2)	(3)	(4)	(5)	(6)
Hypothetical scenario of transfer provision	3.859** (0.568)	4.048** (0.608)	4.158** (0.637)	0.542** (0.047)	0.542** (0.048)	0.542** (0.048)
Role in the first six rounds	0.539 (0.569)	0.575 (0.636)	0.735 (0.737)	0.043 (0.048)	0.046 (0.052)	0.065 (0.059)
Age		0.016 (0.177)	-0.015 (0.187)		0.003 (0.014)	0.001 (0.015)
Female		-0.159 (0.513)	-0.118 (0.551)		-0.028 (0.045)	-0.018 (0.046)
Christian		0.841 (0.648)	1.031 (0.786)		0.059 (0.053)	0.085 (0.064)
Muslim		1.358 (0.922)	1.581 (0.883)		0.085 (0.079)	0.120 (0.076)
Choice in beauty contest			-0.002 (0.011)			0.0001 (0.001)
Number of correct answers in CRT			0.086 (0.366)			0.022 (0.029)
Altruism			-0.016 (0.230)			-0.012 (0.019)
Positive reciprocity			-0.311 (0.195)			-0.022 (0.018)
Observations	120	120	120	120	120	120

Note: Observations are pooled from the *guessing game*: beliefs of 60 participants under two hypothetical scenarios, ( $N = 120 = 60 \times 2$ ); Models 1–3 and 4–6 use logit (non-exponentiated coefficients) and OLS, respectively; standard errors shown in parentheses; cutoff and constant are suppressed in the report; \*  $p < 0.05$ , \*\*  $p < 0.01$ .

Table 3 shows the main results. In Models 1–3, which test Hypothesis 2, the dependent variable is the voter’s action in the round: it takes a value of 1 (–1) if the voter cast their ballot for the clientelistic (programmatic) candidate, and equals zero if the voter abstained.<sup>19</sup> Similarly, we test Hypothesis 3 in Models 4–6, where the dependent variable is an indicator that equals 1 if the candidate offered the transfer, and equals zero otherwise. The main

<sup>19</sup>Results from random-effects multinomial logistic regressions (Appendix E.6) show that offering the transfer increases (decreases) the likelihood of voting for the clientelistic (programmatic) candidate.

Downloaded from https://www.cambridge.org/core. IP address: 86.30.115.167, on 12 Feb 2026 at 18:49:11, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/jpsrm.2024.54

explanatory variables are *Private transfer offered*, which is an indicator of whether the respondent playing the role of the voter was offered the transfer in the round, and *Monitoring capacity*, which indicates if the respondent playing the role of the candidate had strong monitoring capacity in the round.

**Table 3.** Regression results regarding vote choice and transfer provision

	Dependent variable:					
	Vote choice (-1/0/1)			Transfer offer (0/1)		
	(1)	(2)	(3)	(4)	(5)	(6)
Private transfer offered	3.317** (0.343)	3.286** (0.342)	3.329** (0.343)			
Monitoring capacity				2.062** (0.357)	2.063** (0.357)	2.077** (0.358)
Role in the first six rounds	-0.228 (0.582)	-0.289 (0.554)	-0.408 (0.549)	-0.308 (0.699)	-0.453 (0.709)	-0.588 (0.725)
Age		-0.031 (0.165)	0.015 (0.159)		-0.282 (0.214)	-0.266 (0.212)
Female		1.525** (0.562)	1.571** (0.568)		-0.152 (0.700)	0.010 (0.727)
Christian		0.105 (0.620)	-0.313 (0.630)		0.503 (0.786)	0.442 (0.806)
Muslim		0.866 (0.757)	0.797 (0.752)		-0.128 (0.968)	-0.096 (0.984)
Choice in the beauty contest			-0.017 (0.012)			-0.011 (0.017)
Num. of correct answers in the CRT			-0.339 (0.259)			0.249 (0.333)
Altruism			-0.203 (0.208)			0.157 (0.270)
Positive reciprocity			0.249 (0.221)			0.376 (0.285)
Round	0.090 (0.074)	0.087 (0.074)	0.089 (0.074)	-0.485** (0.096)	-0.484** (0.096)	-0.487** (0.096)
Participant random effects (Variance):	4.596 (2.144)	3.847 (1.961)	3.426 (1.851)	6.520 (2.553)	6.265 (2.503)	5.903 (2.430)
Observations	420	420	420	420	420	420

Note: Each observation is a participant-round in the election game; Models 1–3 use the six rounds that each of the 70 participants played as voters ( $N = 420 = 70 \times 6$ ), and Models 4–6 use the six rounds that each of the 70 participants played as clientelistic candidates ( $N = 420 = 70 \times 6$ ); All models use random effect regressions (ologit for Models 1–3 and logit for 4–6); standard errors in parentheses; \* $p < 0.05$ , \*\*  $p < 0.01$ .

These results provide strong support for Hypotheses 2 and 3. The coefficients of *Private transfer offered* and *Monitoring capacity* are consistently positive and significant across all our specifications. Substantively, the coefficient of *Private transfer offered* in Model 3 indicates that receiving the private transfer increases the likelihood that the voter votes for the clientelistic candidate from 0.06 to 0.63. Likewise, the coefficient of *Monitoring capacity* in Model 6 implies that the probability the candidate offers the transfer to the voter is 0.64 when she has strong monitoring capacity, but only 0.18 when her capacity is weak.

### 4.3 Mechanism: the informational effects of vote buying

The previous results are strongly consistent with Proposition 1’s separating equilibrium. However, we have not explored one of the model’s key insights, which is that vote buying affects voter behavior *by revealing information about the candidate’s monitoring capacity*. Thus, it is possible that the observed patterns are driven by channels other than the informational mechanism we propose. Here, we provide evidence for this mechanism by analyzing the relationship between the participants’ beliefs, which we elicited in the guessing game, and their behavior in the

Downloaded from https://www.cambridge.org/core. IP address: 86.30.115.167, on 12 Feb 2026 at 18:49:11, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/psrm.2024.54

election game. We expect the patterns described in Hypotheses 2 and 3, and documented in Table 3, to be stronger among participants whose beliefs are consistent with Proposition 1's separating equilibrium.

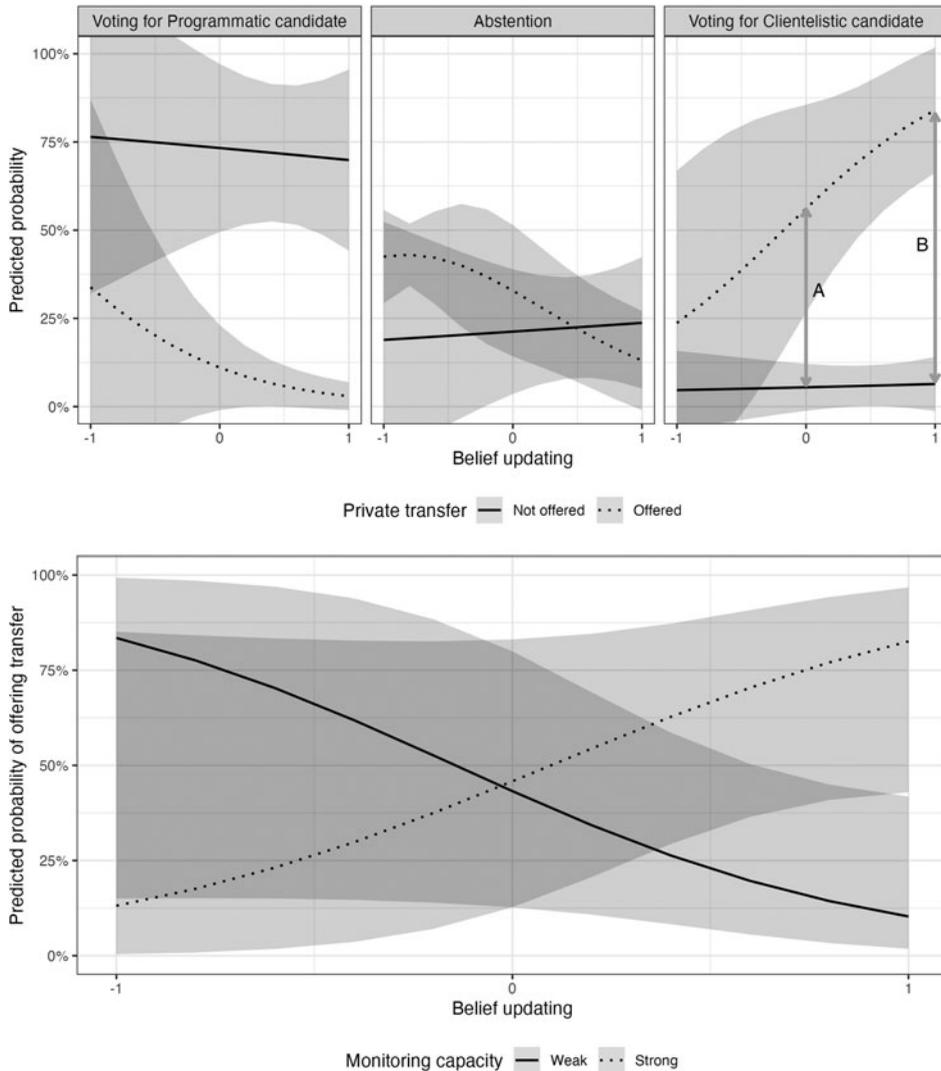
First, we use the respondents' *continuous* measure of beliefs to create a variable that captures the type of belief updating required by the separating equilibrium. This variable (named *Belief updating*) is equal to the respondent's belief under the hypothetical scenario in which the candidate offered the transfer minus their belief under the hypothetical scenario in which the candidate did not offer the transfer. It ranges from  $-1$  to  $1$ , with greater (lower) values corresponding to beliefs that are more (less) consistent with the voter's beliefs in the separating equilibrium. A score of  $1$  indicates the respondent is fully confident that the candidate is able to monitor (i.e., strong type) in the transfer scenario and unable to monitor (i.e., weak type) in the no-transfer scenario, while a score of  $-1$  means the respondent is fully confident that the opposite is the case. The score equals  $0$  when the respondent has the same belief (and confidence level) about the candidate's monitoring capacity under the two scenarios.<sup>20</sup> Appendix E.7 shows the distribution of *Belief updating*.

We replicate Table 3's models but including *Belief updating* and its interaction with the main explanatory variables, *Private transfer offered* and *Monitoring capacity*. Figure 5 shows a visual representation of our results (Appendix E.8 shows regression estimates). The top panel shows predicted probabilities for the voter's actions as a function of (1) the candidate's choice, and (2) the voter's belief updating score. Consistent with Hypothesis 2, voters are more likely to support the clientelistic candidate when they receive the transfer than when they do not. More importantly, as expected, the impact of receiving the private transfer on the probability of voting for the clientelistic candidate is increasing in the value of *Belief updating* (see dotted line in right-most panel).

Substantively, the predicted behavior of a voter with a belief-updating score equal to  $0$  serves as a benchmark. As described above, a score of  $0$  means that receiving the transfer did not change the voter's belief about the candidate's monitoring capacity. At this level, we see the voter is more likely to vote for the clientelistic candidate after they receive the private transfer; specifically, the marginal effect of *Private transfer offered* equals  $0.51$  ( $=0.56 - 0.05$ ; marked as A in the upper right panel of Figure 5). This effect cannot be explained by the informational mechanism we propose, and thus must be driven by other channels (e.g., positive reciprocity, voters assuming the candidate can always monitor their vote). We can compare this to a respondent with a belief-updating score equal to  $1$ , which represents the ideal case in which the respondent's beliefs exactly match the beliefs of the voter in the separating equilibrium. In this case, the marginal effect of the transfer on the likelihood of voting for the clientelistic candidate is  $0.78$  ( $=0.84 - 0.06$ ; marked as B in the upper right panel of Figure 5). The difference between these two marginal effects is  $0.27$  ( $=0.78 - 0.51$ ; subtracting A from B), which indicates that, even after accounting for the effect of *Private transfer offered* through any non-informational mechanisms (which we estimated to be of size  $0.51$ ), the informational mechanism we advance still has a sizable impact on voter behavior.

Similarly, the lower panel shows predicted probabilities of the candidate's actions by (1) the candidate's type, and (2) the candidate's belief-updating score. Consistent with the informational mechanism, we find the evidence for Hypothesis 3 reported in Table 3 is, in fact, fully driven by participants with high belief-updating scores. The marginal effect of *Monitoring capacity* on the probability that the candidate provides the transfer is negative (although not statistically significant) for most negative values of *Belief updating*. In contrast, this marginal effect

<sup>20</sup>Scores between  $-1$  and  $0$ , and between  $0$  and  $1$ , indicate respondents are more confident about their beliefs under one hypothetical scenario than under the other.



**Figure 5.** Predicted probabilities of players' actions according to *Belief updating*,  
 Note: Lines and gray areas show predicted probabilities and 95 percent confidence intervals, respectively.

is positive and statistically significant for high-enough belief-updating scores. The magnitude of this effect is substantial, reaching a maximum of 0.73 (=0.83 – 0.10) when *Belief updating* equals 1.

Finally, to verify the robustness of our results, we replicate the models reported in Table 3 and Figure 5 in several ways. First, we use two alternative measures of respondents' beliefs (Appendices E.9-E.10). Second, although our within-subject design reduces the concern about the use of post-manipulation variables in the analysis (Montgomery *et al.*, 2018), we replicate the models but excluding all post-manipulation variables (Appendix E.11). Finally, we address the potential that feedback from previous rounds affects participants' current choices by adding lagged variables (Appendix E.12). All results are substantively equal to those reported in the manuscript.

## 5. Conclusion

We study an unexplored mechanism that explains why, under certain conditions, vote buying can be effective even in secret-ballot elections. We argue that, by initiating vote-buying exchanges, politicians can erode voter confidence in ballot integrity, which in turn may contribute to making vote buying effective. We formalize this intuition by modeling the interaction between a clientelistic candidate and a voter as a signaling game. Our main result shows that electoral handouts serve an informational role. Specifically, receiving a handout makes the voter more likely to both believe the candidate can monitor his vote choice and comply with the vote-buying exchange.

The result described above can emerge even when the *ex ante* probability that the candidate can monitor the voter's behavior is extremely low. This has substantive implications regarding the types of real-world cases in which our informational mechanism could be at play. First, the model applies to settings in which contextual features make monitoring individual vote choices feasible, however rare this may be.<sup>21</sup> Second, the model also captures situations in which monitoring is imperfect.<sup>22</sup> These types of monitoring require a considerable amount of resources and organizational capacity (e.g., Stokes, 2005; Larreguy *et al.*, 2016). Thus, a natural interpretation of our model is that voters do not know whether a candidate's clientelistic machine has the resources and organizational strength to carry out some type of monitoring in their own polling stations.

We conclude by highlighting two avenues for future empirical research. The first has to do with the scope of our findings. In this paper, our focus is on assessing whether experiencing vote buying affects a very specific outcome: confidence in ballot secrecy. However, we acknowledge that exposure to vote buying could more generally erode voter perceptions of electoral integrity (i.e., beyond ballot secrecy). While we are not able to test this possibility with the data at hand, we believe that understanding the extent to which vote buying shapes public opinion has important practical implications. For instance, this knowledge could inform the design of educational campaigns aimed at curbing electoral malpractice.

The second one relates to the generalizability of our findings. The external validity or generalizability of experiments using online convenience and student samples is well documented in existing studies (e.g., Mullinix *et al.*, 2015; Lupton, 2019). Although it is important to acknowledge that participants in our experiments are different from people who typically participate in clientelistic exchanges in several important respects, such as income and educational attainment, we also highlight that *a priori* it is unclear whether and how our findings would change if we replicated this study with more representative samples. Thus, replicating this study in other settings could provide additional insights about the effectiveness and limits of our theory.

**Supplementary material.** The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2024.54>. To obtain replication material for this article, <https://doi.org/10.7910/DVN/XTEML4>

**Acknowledgements.** We thank Rebecca Morton, who sadly passed away in September 2020. She passionately organized a series of workshops for postdoctoral researchers at New York University Abu Dhabi's Social Science Division. This manuscript greatly benefited from the constructive and encouraging feedback she and other participants provided at these workshops. We also thank Francisco Cantú, Romain Ferrali, Ji Yeon (Jean) Hong, Adrián Lucardi, Miguel Rueda, and Jonathan Woon for comments and suggestions at various stages of this project. All remaining errors are, of course, our own.

**Competing interests.** The authors declare that they have no conflict of interest and that they did not receive any funding for conducting this study.

<sup>21</sup>Qualitative evidence shows the type of ballot used Brusco *et al.*, 2004 and the presence of party agents at polling stations (Mares and Young, 2019, 127) occasionally allow parties to monitor individual vote choices.

<sup>22</sup>For example, clientelistic parties monitor turnout at the individual level (e.g., Nichter, 2008; Larreguy *et al.*, 2016; Ascencio and Rueda, 2019) and vote choices at the group level (e.g., Schaffer and Schedler, 2007; Rueda, 2017).

## References

- Ascencio SJ and Rueda MR (2019) Partisan poll watchers and electoral manipulation. *American Political Science Review* **113**, 727–742.
- Aziz A, Castro Cornejo R, Gilas K, Isunza E, de la Peña R and Somuano F (2022) Informe país 2020. el curso de la democracia en México.
- Bratton M (2008) Vote buying and violence in Nigerian election campaigns. *Electoral Studies* **27**, 621–632.
- Brierley S and Nathan NL (2022) Motivating the machine: which brokers do parties pay?. *The Journal of Politics* **84**, 1539–1555.
- Brusco V, Nazareno M and Stokes SC (2004) Vote buying in Argentina. *Latin American Research Review* **39**, 66–88.
- Cantú F and García-Ponce O (2015) Partisan losers' effects: perceptions of electoral integrity in Mexico. *Electoral Studies* **39**, 1–14.
- Carpenter J, Graham M and Wolf J (2013) Cognitive ability and strategic sophistication. *Games and Economic Behavior* **80**, 115–130.
- Challú C, Seira E and Simpser A (2020) The quality of vote tallies: causes and consequences. *American Political Science Review* **114**, 1071–1085.
- Cox GW and Kousser JM (1981) Turnout and rural corruption: New York as a test case. *American Journal of Political Science* **25**, 646–663.
- Cruz C (2019) Social networks and the targeting of vote buying. *Comparative Political Studies* **52**, 382–411.
- Dahl RA (1971) *Polyarchy: Participation and Opposition*. New Haven, CT: Yale University Press.
- Duffy J and Tavits M (2008) Beliefs and voting decisions: a test of the pivotal voter model. *American Journal of Political Science* **52**, 603–618.
- Ferree KE and Long JD (2016) Gifts, threats, and perceptions of ballot secrecy in african elections. *African Affairs* **115**, 621–645.
- Finan F and Schechter L (2012) Vote-buying and reciprocity. *Econometrica* **80**, 863–881.
- Frederick S (2005) Cognitive reflection and decision making. *Journal of Economic perspectives* **19**, 25–42.
- Gans-Morse J, Mazuca S and Nichter S (2014) Varieties of clientelism: machine politics during elections. *American Journal of Political Science* **58**, 415–432.
- Gerber AS, Huber GA, Doherty D and Dowling CM (2012) Is there a secret ballot? Ballot secrecy perceptions and their implications for voting behaviour. *British Journal of Political Science* **43**, 77–102.
- Gingerich DW and Medina LF (2013) The endurance and eclipse of the controlled vote: a formal model of vote brokerage under the secret ballot. *Economics & Politics* **25**, 453–480.
- Kramon E (2016) Electoral handouts as information: explaining unmonitored vote buying. *World Politics* **68**, 454–498.
- Larreguy H, Marshall J and Querubín P (2016) Parties, brokers and voter mobilization. *American Political Science Review* **110**, 160–79.
- Lawson C and Greene KF (2014) Making clientelism work: how norms of reciprocity increase voter compliance. *Comparative Politics* **47**, 61–77.
- Lupton DL (2019) The external validity of college student subject pools in experimental research: a cross-sample comparison of treatment effect heterogeneity. *Political Analysis* **27**, 90–97.
- Mares I and Young LE (2019) *Conditionality and Coercion: Electoral Clientelism in Eastern Europe*. Oxford: Oxford University Press.
- Montgomery JM, Nyhan B and Torres M (2018) How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science* **62**, 760–775.
- Morgan J and Vardy F (2012) Negative vote buying and the secret ballot. *Journal of Law, Economics, and Organization* **28**, 818–849.
- Mullinix KJ, Leeper TJ, Druckman JN and Freese J (2015) The generalizability of survey experiments. *Journal of Experimental Political Science* **2**, 109–138.
- Muñoz P (2014) An informational theory of campaign clientelism: the case of Peru. *Comparative Politics* **47**, 79–98.
- Nichter S (2008) Vote buying or turnout buying? Machine politics and the secret ballot. *American Political Science Review* **102**, 19–31.
- Oliveros V (2019) Perceptions of ballot integrity and clientelism. In Lupu N., Oliveros V., and Schiumerini L. (eds), *Campaigns and Voters in Developing Democracies: Argentina in Comparative Perspective*. Ann Arbor, MI: University of Michigan Press, pp. 213–238.
- Robinson JA and Verdier T (2013) The political economy of clientelism. *The Scandinavian Journal of Economics* **115**, 260–291.
- Rosas G (2010) Trust in elections and the institutional design of electoral authorities: Evidence from Latin America. *Electoral Studies* **29**, 74–90.
- Rueda MR (2017) Small aggregates, big manipulation: vote buying enforcement and collective monitoring. *American Journal of Political Science* **61**, 163–177.

- Schaffer FC and Schedler A** (2007) What is vote buying? In Schaffer F. C. (ed.), *Elections for Sale: The Causes and Consequences of Vote Buying*. London: Lynne Rienner Publishers.
- Smith A and Bueno de Mesquita B** (2011) Contingent prize allocation and pivotal voting. *British Journal of Political Science* **42**, 371–392.
- Stokes S** (2005) Perverse accountability: a formal model of machine politics with evidence from Argentina. *American Political Science Review* **99**, 315–325.