Unpacking Ethnic Preferences: Theory and Micro-Level Evidence From North India

SAGE

Simon Chauchard[1]

Abstract

Much of the recent scholarship about ethnicity in comparative politics has focused on *why* ethnicity becomes a salient cleavage. Yet opinions still diverge as to *how* ethnicity matters. This article tests three hypotheses relevant to this question. Building on recent arguments, it first hypothesizes that voters take the ethnicity of parties *and* candidates into consideration. Second, it hypothesizes that wherever ethnicity is politically salient, it matters *beyond coethnicity*—that is, voters' decisions are guided by ethnicity even when they are choosing among non-coethnics. Third, it argues that the advantage conferred by ethnicity is mediated by non-ethnic factors. A large vignette experiment carried in 2013 in the Indian state of Uttar Pradesh provides strong support for each of these hypotheses. These results imply that the influence of ethnicity on voting behavior may be subtler and more complex than what the main theories of ethnic politics usually assume.

Introduction

Over the past decade, an important literature has sought to explain *why* ethnicity becomes a salient cleavage in diverse democracies. Authors have explored why certain ethnic cleavages become salient when others do not

[1]Dartmouth College, Hanover, NH, USA

Corresponding Author:
Simon Chauchard, Dartmouth College, 6108 Silsby Hall, Hanover, NH 03755, USA.
Email: simon.chauchard@dartmouth.edu

(Posner, 2005), the conditions under which voters are likely to choose ethnic parties (Chandra, 2004; Pepinsky, Liddle, & Mujani, 2012), and the motivations of ethnic parties and ethnic voters (Chandra, 2004; Ferree, 2011). This literature has changed the way we interpret ethnic voting. According to these authors, voters' preference for parties perceived to be favorable to their coethnics has less to do with "expressive voting" (Horowitz, 1985) than with the conditions under which elections take place. Ethnic parties win *not* because voters feel strongly about their identities but rather because the ethnic labels attached to parties constitute useful informational shortcuts—or heuristic devices—for patronage-seeking voters. The convergence toward this conclusion has been remarkable.

Yet no such convergence exists as to *how* ethnicity affects voters' choices. As a result, key works in this literature rely on assumptions that remain disputed and/or untested. This article empirically evaluates these assumptions. Namely, it tests three interrelated hypotheses. The first one has to do with the level at which the influence of ethnicity takes place. The most cited arguments about ethnic voting focus on voters' comparative evaluations of different parties (Chandra, 2004; Horowitz, 1985). Yet, in single-member districts, the idea that voters solely focus on party-level "ethnic labels"—that is, on whether a party is overall perceived to be "inclusive" or representative of a given ethnic group—may not be entirely credible. If, as already articulated by several authors (Ferree, 2011; Madrid, 2012), "ethnic labels" evolve over time as a result of parties' strategic diversification of their leadership, we should not expect voters to be indifferent to the identity of the specific candidates fielded by parties in their constituency. Building on these arguments, I accordingly hypothesize that voters' ethnic considerations are not entirely located at the party level and that voters also take the ethnicity of the specific *candidates* on the ballot into account when evaluating candidates. My second hypothesis has to do with the aspect of ethnicity that drives ethnic preferences. The comparative literature has so far mostly thought of ethnicity-based preferences in terms of a *preference for coethnics* and in terms of a binary coethnic/non-coethnic alternative, potentially overlooking more sophisticated reasoning about ethnic categories by voters. This may be limiting, insofar as voters may take ethnicity into account including where none of their coethnics are competing. Accordingly, I hypothesize that voters are likely to sort out between different types of non-coethnics. Finally, my third hypothesis relates to the malleability of ethnicity-based preferences. Most studies in this subfield imply that coethnicity confers an absolute rather than a conditional advantage in elections. By contrast, I hypothesize that the advantage conferred by ethnicity is merely *conditional* and that voters' ethnicity-based preferences are mitigated by other factors.

To test these three hypotheses—and better understand *how* ethnicity matters—this article relies on a large vignette experiment that was implemented in two separate locations across one of the most famous cases in the recent literature on ethnic parties: the Northern Indian state of Uttar Pradesh (UP). In this experiment, likely voters rated fictional candidates whose key characteristics (*ethnicity*, *party label*, *perceived probity*, and *perceived record*) had been randomly manipulated. Insofar as real-world elections imply a degree of strategy that was not present in the experiment, these ratings are not equivalent to voting decisions. They nonetheless provide us with important and causally credible insights about voters' *evaluation* of candidates and, more generally speaking, about the psychology of the median voter in North India.

These ratings provide support for all three aforementioned hypotheses. Because this design isolates the relative weight of party and candidate ethnicity in voters' evaluations, it allows me to show that the ethnicity *of candidates* matters in addition to their *party* affiliations. Second, as the experiment exposed respondents to different types of non-coethnic candidates, it provides evidence that ethnicity matters *beyond coethnicity* (i.e., whether voters make a difference among different types of non-coethnics). Finally, because factors other than ethnicity (*perceived probity* and *perceived record* of the candidates) are also manipulated, I am able to show that ethnicity-based preferences are not absolute.

This research contributes to the burgeoning literature about the effect of various candidate characteristics in India (Aidt, Golden, & Tiwari, n.d.; Berenschot, 2011; Michelutti, 2008; D. Tiwari, n.d.; Vaishnav, 2012; Ziegfeld, in press). Beyond India, it more generally contributes to a growing experimental comparative literature evaluating the relative weight placed on candidates' (ethnic and non-ethnic) characteristics by voters (Banerjee, Green, McManus, & Pande, 2014; Beaman, Chattopadhyay, Duflo, Pande, & Topalova, 2009; Conroy-Kurtz, 2013; Ziegfeld, in press). Its main contribution, however, is to the large and expending comparative literature on ethnic politics and ethnic parties. In line with the evolution of theoretical models of political selection—from the simple Downsian model of competition to more multi-dimensional models, where the identity of specific candidates also matters (Akerlof & Kranton, 2005; Besley, 2005; Chattopadhyay & Duflo, 2004)—these results imply that ethnicity affects the preferences of voters in ways that are both more complex and subtler than has so far been suggested in the comparative literature on ethnic politics. They imply that models focused on the effect of ethnicity *at the party level* (Chandra, 2004; Horowitz, 1985) in part misrepresent *how* ethnicity matters. They also support the idea that ethnic parties can attract new voters by fielding new types of candidates (Chandra, 2004; Ferree, 2011; Madrid, 2012), a finding which more generally supports the idea that the

"ethnic labels" of parties evolve over time (Birnir, 2007). Third, these results show that the impact of ethnicity—in electoral politics and beyond—cannot be reduced to a simple preference for coethnics, insofar as individuals sort out among non-coethnics at least as much as they do between coethnics and non-coethnics. Fourth and finally, these results suggest that the effect of ethnicity is often mediated by other factors, contrary to what much of the comparative literature on voting behavior in diverse societies (Chandra, 2004; Ferree, 2011; Horowitz, 1985; Pepinsky et al., 2012) has so far implied.

Theorizing Ethnic Preferences in Comparative Politics

Over the past decade, leading arguments in comparative politics (Birnir, 2007; Chandra, 2004; Ferree, 2011; Pepinsky et al., 2012; Posner, 2005) have explored the motivations of ethnic voters. *Why* do voters prefer to vote for ethnic parties? Although these explanations have relied on a variety of data and cases, they have converged toward a remarkably consistent explanation. In low-information contexts, coethnicity allows rational, self-interested voters to make crucial inferences about patronage networks.

   The most important models in this literature have made strong assumptions about the way in which ethnicity enters voters' preferences. While a handful of observational studies have hinted that the influence of ethnicity may be less strong or less automatic (Bratton & Kimenyi, 2008; Lindberg & Morrison, 2008), much of the recent literature has implied that voters are *not* affected by the identity of the specific *candidate* who competes under these party labels, or by the non-ethnic characteristics of this candidate. Furthermore, these models have also implied that voters do not differentiate among non-coethnics. Yet each of these assumptions appears problematic.

Beyond Party Labels

There are, first, good reasons to believe that ethnicity matters both *at the party level* and *at the candidate/constituency level.* In countries in which ethnic cleavages are salient, it is not surprising that voters' evaluations of candidates would derive, at least in part, from the ethnic labels attached to their parties (Chandra, 2004; Horowitz, 1985). But it seems unrealistic that voters would not simultaneously pay attention to the ethnicity of the specific candidates fielded by these parties in their constituency. Prominent works based on a variety of contexts (Chandra, 2004; Ferree, 2011; Madrid, 2012) note that ethnic parties strategically distribute party tickets to leaders from

other ethnic groups when they hope to change their image or broaden their potential bases. As a result, these works suggest that the ethnic labels attached to parties derive, at least in part, from the ethnicity of the various candidates fielded by these parties over time. If this is the case, it is especially likely that voters' individual-level evaluations are also, in part, influenced by the ethnicity of the *specific* candidate they are choosing from. If voters perceive different ethnic makeups at the party level to be more or less conducive to patronage, it is likely that they also perceive variations in the ethnicity of each specific candidate they are evaluating to have similar consequences, regardless of the ethnic label they perceive to be attached to that candidate's party. Voters' evaluations should thus be a function of the existing label attached to the party as well as of the ethnicity of the specific candidate fielded by these parties.

Insofar as the strength of partisanship tends to increase around elections (Clarke & Stewart, 1998; Lupu, 2013; Michelitch & Utych, 2015), we should expect the relative weights voters place on party labels and on candidate ethnicity to fluctuate over the course of the electoral cycle. It is, however, difficult to think of contexts or timings in which voters would be expected not to pay *any* attention to the ethnicity of the specific candidates they are evaluating. This might be the case wherever voters believe that their representative plays no role in the distribution of state-related benefits, or where representatives' decisions do not directly affect the livelihood of voters. In most ethnically diverse democracies—including India, South Africa, Zambia, and other important cases featured in this literature—this is, however, not the case. Individual representatives are usually in charge of disbursing *some* funds to their districts. More importantly, voters need the mediation of their representative to access the state and its many benefits. Given the life-or-death character that this form of constituency service can take in "patronage-democracies" (Chandra, 2004), it is unlikely that voters who otherwise perceive ethnic favoritism to be a structuring force would overlook the ethnicity of candidates.[1]

Beyond Coethnicity

The second problematic assumption relates to the way in which voters evaluate different types of *non*-coethnics. To the best of my knowledge, *all* aforementioned works in this literature imply that voters perceive ethnic differences in strictly binary terms: Coethnic parties or candidates are preferred to all others (hereafter referred to as "non-coethnics"), and voters do not distinguish among those non-coethnics.

Yet, in most contexts in which ethnicity is politically salient and in which more than two salient groups exist, there is reason to believe that voters do sort out among non-coethnics. In fractionalized settings (i.e., settings wherever there are more than *two* salient groups, none of which accounts for a majority of the population), a model based on a form of valence for coethnics should prove a limited tool to explain the role of ethnicity in the formation of political preferences. In such contexts, a better—if slightly more complex—model might consider the additional existence of a type of *negative* valence directed at *some* non-coethnics, which implies that different types of non-coethnics may be evaluated differently.

This broader, non-binary type of ethnicity-based preference may result from several dynamics. Coalition politics could first prompt voters to differentiate among non-coethnics that are seen as being allied to their coethnics within a political organization and non-coethnics that are seen as being opposed to coethnics in the current party system. More generally speaking, the history of intergroup relations may be expected to be different with different groups of non-coethnics. Antagonisms may exist between Group A and Group B, but not between Group A and Group C. In the United States, it is, for instance, likely that Whites would *not* evaluate Latino politicians and Black politicians similarly; in Kenya, a Luo voter is unlikely to have the same baseline attitude toward a Kalenjin politician and a Kikuyu politician, even though both are perceived as non-coethnics. In India, it seems even more unlikely that voters would evaluate politicians from the several dozens of locally salient caste-based and religious groups that do not qualify as coethnics.

This non-binary model of ethnicity-based preferences should be especially important if voters react to candidates more than they react to parties, as suggested in the previous subsection. Where a large number of salient groups coexist, for instance in India, it is in fact likely that many voters do not have a coethnic candidate on the ballot; in this configuration, large groups of voters effectively *have to* rank non-coethnics to choose a candidate.

An Absolute Advantage of Ethnicity?

A third problematic assumption common in the ethnic politics literature relates to the relative importance of ethnicity in voting preferences. Because they mostly focus on *why* ethnicity matters, or because they strive to present parsimonious models, arguments about ethnic voting often appear to suggest that ethnicity is the *only* or the *most significant* factor of voters' preferences, even if evidence suggesting the contrary (Bratton & Kimenyi, 2008; Lindberg & Morrison, 2008) has emerged. According to many oft-cited models, the

ideologies, probity, and record of candidates are either unimportant (Ferree, 2011; Horowitz, 1985) or irrelevant (Chandra, 2004).[2]

Assuming that voters' preferences are established at the party level, Chandra (2004) develops a reasonable rationale for this overwhelming influence of ethnicity. In many settings in which party-voters linkages are ethnic in nature, ideology and programs are unlikely to play an important role. As such polities are likely to be low-information polities, voters likely struggle to learn something useful about competing parties (see also Birnir, 2007 on this point). Information about the *overall* quality of candidates fielded by said party, about their probity, or about the overall record of the party may be difficult to access, compile, or comprehend. By contrast, the ethnic makeup of parties sends a simple, visible, and unambiguous message to the electorate. Following this reasoning, ethnicity should play a central role wherever information is scarce and voting decisions are party centric.

When voters also take candidate characteristics into account, this is, however, unlikely to be true. If voters pay attention to candidates, it is not clear why they would merely focus on their ethnicity and disregard their non-ethnic characteristics. Although it is difficult for voters to obtain aggregate information about a party's record, program, or probity, it is easier for them to obtain *some* information about the specific candidate competing in their constituencies. Patronage-seeking voters are, for one, expected to have a sense of the record of their incumbent. Besides, information concerning the reputation or the probity of that candidate should be available at the local level (such information is not readily available about parties only because it is difficult for voters to aggregate the information across many constituencies). If voters pay attention to candidate-level characteristics when forming their political preferences, and if they have all this additional information when evaluating *candidates*, then it is unlikely that they would *only* take ethnicity into account. If that was the case, we should expect the effects of ethnicity and coethnicity to be mitigated by other, non-ethnic candidate-level factors, such as perceived record or perceived probity.

Summary of Hypotheses

This brief review of the recent comparative literature on ethnic voting in sum suggests that three interrelated hypotheses should be tested to clarify *how* ethnicity matters:

> **Hypothesis 1 (H1):** Ceteris paribus, voters take the ethnicity of candidates into account (i.e., coethnic candidates should receive better evaluations than non-coethnic candidates running for the same party).

Hypothesis 2 (H2): When taking the ethnicity of candidates into account, voters distinguish between different types of non-coethnics (some non-coethnics are evaluated more positively than others).

Hypothesis 3 (H3): Non-ethnicity-related candidate characteristics mitigate the effect of ethnicity: Coethnics with a good record are preferred to coethnics with a bad record, and non-coethnics with a good record are preferred to non-coethnics with a bad record.

Research Design

Although two of these hypotheses are not entirely novel—H1 and H3 build on existing arguments—we lack causally credible tests for all three of these hypotheses. To generate such tests, I implemented a vignette experiment in which likely voters rated randomly manipulated profiles of candidates in the Northern Indian state of UP.

The Case: Ethnic Politics in UP

Given the diversity of cases described in the comparative literature on ethnic politics, several locations would have been appropriate to test these hypotheses. Focusing on the Northern Indian state of UP, however, made sense for a combination of reasons. For one, UP—with a population of more than 200 millions—is the largest Indian state. Second, it is the case on which much of the argument of Chandra (2004), a key reference in this literature, relies. Third, even if UP is often presented among Indian states as a state in which "ethnic politics" has triumphed, it is in no case the only state in India in which ethnicity factors into political preferences. Recent accounts of politics in settings as diverse as Mumbai (Hansen, 2001) and rural Rajasthan (Chauchard, in press) suggest that the politics of *most* Indian states—including many Southern Indian states—is structured around a "caste-based mosaic" that sees caste groups forge and break alliances with various political parties (Jaffrelot, 2010). Besides, as I show below and in the Online Appendix E, recent survey data suggest that the preferences of UP voters are altogether not that different from the preferences of other Indian voters. In that sense, trends among UP voters should, *at the very least*, not be entirely irrelevant to trends among other Indian voters. Insofar as one agrees with Chandra that UP and India shares key characteristics with a number of other "patronage-democracies"—relatively uninformed voters, political control over the distribution of state benefits, the persistence of a large state—the UP case should in addition be relevant to a number of other cases of ethnic politics across the post-colonial world, such as Nigeria, Zambia, or Senegal. Finally, voters in UP—as

elsewhere in India and in many representative democracies—choose among several candidates *in single-member constituencies*, making the state an especially appropriate location to test the aforementioned hypotheses.[3]

Ethnicity plays a major role in UP politics (Chandra, 2004; Pai, 2002), as religion and caste-based divisions have over the past decades become the most electorally relevant cleavages in the state (Jaffrelot, 2010; Pai, 2002). The two locally dominant parties—the Samajwadi Party (SP), currently in power at the state level, and the Bahujan Samaj Party (BSP), in power before that, until 2012—openly appeal to various and fast-changing combinations of groups within the state.[4] As a result of these partisan dynamics, all dominant political parties in the state have in recent times considered the ethnicity of candidates vying for their tickets. The two locally dominant parties—the SP and the BSP—have engaged in "social engineering" in an effort to build around their original caste base (*Yadavs* for the SP, *Chamars* for the BSP). But other players—including two national parties, the Indian National Congress and the Bharatya Janta Party—the latter of which won 71 out of 80 constituencies in the 2014 general elections—have also strategically recruited leaders and pandered to specific groups during electoral campaigns. News reports from recent electoral campaigns in UP suggest that the ethnicity of candidates is expected to play a central role in voters' choices ("UP Assembly Elections," 2012; Khan, 2012; Srivastava, 2012; M. Tiwari & Pandey, 2013), independently of these candidates' party labels.[5] Besides, some candidates are described in these accounts as "polarizing," and voters from different groups often appear to be coordinating *against* a candidate from a particularly disliked group, hence suggesting the existence of a form of negative valence.

Importantly, the operative ethnic categories on which these parties have based their strategies in recent times do *not* correspond to the larger categories to which some recent works on ethnic voting in India (Banerjee et al., 2014) refer, but rather to smaller, locally salient groups, organized on the basis of subcastes (*jatis*) among Hindus, and on the basis of religion among non-Hindu minorities (Suryanarayan & Huber in press; M. Tiwari & Pandey, 2013). This nuance is important: As parties appeal to these more specific subunits instead of appealing to larger caste group categories, they frequently target specific communities within one of these larger categories while ignoring, or failing to target, others.[6]

Given the complex setup of the experiment—described at length below—and given the large size of UP, data collection could only realistically take place in a handful of locations across the state. Because I wanted rural voters to constitute the majority of the sample—the population of UP remains 76.79% rural—I first *randomly* selected Gonda district among the list of districts whose headquarter city counted fewer than 150,000 inhabitants. I

subsequently selected the district of Faizabad as an additional case to evaluate across-district variation.[7] While limiting myself to two locations allowed for better monitoring of the experimental procedure of the study—hence enhancing the internal validity of the study—this by definition limited the representativeness of my results. In light of this, it seems important to highlight that Gonda and Faizabad districts are not obvious outliers within UP. As can be seen from the broad range of aggregate statistics and public opinion data available on the two districts (Online Appendix A), the two districts are not particularly dissimilar to the rest of the state.

The Vignette Experiment

To evaluate voters' preferences, I exposed a sample of likely voters to profiles of *fictional*—but realistic—politicians presented as likely candidates in the (then upcoming) general elections. Relying on *fictional* politicians provided me with an important advantage: It allowed me to randomly vary several of their characteristics, which would have been impossible and unethical with real politicians. Because many voters know the identity of likely candidates in their district, these fictional politicians were presented as likely contenders *in other districts* of UP.[8] Because some voters may have known the identity of candidates from other districts immediately before the election, the experiment was run in between elections, that is, in the summer of 2013, 10 months before the general elections of May 2014.[9]

Although the analyses presented below focus on a subset of these, no fewer than eight manipulations were inserted in each vignette. As noted by Hainmueller, Hopkins, and Yamamoto (2014), complex vignettes allow for richer profiles and for a more realistic experiment. This appeared especially important in the context of this study, insofar as the amount of information revealed can change respondents' propensity to develop ethnicity-based preferences (Conroy-Kurtz, 2013).[10]

Four of these pieces of information were *not* of any direct theoretical interest to this study: the face of the politician (i.e., one of the pictures presented in Figure 1), his previous experience in politics, his family status, and the district in which he planned to run. Although these manipulations were randomly varied across the three profiles viewed by each respondent, I did not expect them to have an effect on the outcomes of interest—the results presented in the Online Appendix B confirm that they did not.[11] Accordingly, I do not comment on the effect of these manipulations in the rest of this article. I instead focus on the effect of the four remaining manipulations, which—by contrast—were directly relevant to the three hypotheses outlined above. These manipulations concern the following:

Figure 1. The three fictitious politicians.

1. The politician's ethnicity: Vignettes exposed respondents to a fictitious candidate that belong to one of three possible groups: (a) the respondent self-identified community, (b) a group that the respondent had *previously* identified as a "friendly outgroup," or (c) a group that the respondent had *previously* identified as a "rival outgroup."[12]

2. His party affiliation: Each fictitious candidate presented in these vignettes also belonged to only one of two possible types of political parties: either their favorite party or their *third favorite party* (from a ranking of the four main parties active in UP: the SP, the BSP, the India National Congress [INC], and the Bharatiya Janata Party [BJP]).[13]

3. His record in terms of development: The politician was either presented as having significantly contributed to the development of his constituency during his term in office or as having done little.[14]

4. The number of serious crimes mentioned on his affidavit:[15] The criminal reputation of politicians took one of two possible values: The politician either had a "clean" reputation (which was conveyed by the fact that "no serious crimes were reported on his affidavit") or a "criminal" one ("four serious crimes" were reported on his affidavit).

Restrictions were implemented to avoid or limit repetitions on each of these manipulations across the three vignettes consecutively evaluated by each respondent. When the choice set was binary, each choice could only appear on two out of the three vignettes. When there were three possible options, each of the three choices appeared only once across the three vignettes. In keeping with attempts by UP's main parties to co-opt elites from

different groups, all possible combinations of caste and party were allowed, to the exception of one (*Muslim* and *BJP*).[16]

A Two-Stage Survey

Because this design required the research team to have information on the respondent's ethnicity, on his or her attitudes toward a variety of non-coethnics, and on his or her partisan preferences *prior to the vignette experiment*, this research took place in two stages.

The first stage of this design implied interviewing a random sample of respondents *at their residence*. In this face-to-face "recruitment interview," respondents were asked a number of background questions. As in most surveys in India, they were asked a series of demographic questions, including an open-ended question about their self-perceived ethnicity. This interview also included a number of questions about respondents' attitudes toward non-coethnics. Namely, respondents were asked to cite two communities with whom their own community has "traditionally maintained good relations" as well as two communities with whom their community has "traditionally had tens*er* relations with." These perceptions were later retrieved to determine ethnic groups that each respondent considered as coethnics, as "friendly outgroups," or as "rival outgroups."[17] Finally, the instrument also included a question that required respondents to rank the four main parties of the state in order of preference. As detailed below, this information was used to determine each respondent's preferred and third-ranked party.

At the end of this interview, respondents were invited to participate, a few days later, to a "completely unrelated" study. They were given the address of a community center in which the research team was based and told about the financial incentives for participating to the study.[18] Upon visiting this location a few days later, they participated to a study during which they evaluated three randomly varied profiles.

To maximize the diversity of respondents that would visit our "lab," "recruitment interviews" took place in 40 polling booths randomly sampled in a perimeter of 20 kilometers around each of the two "labs" in which the research team was established (there were two "labs," one in Gonda and one in Faizabad). Using random walk techniques from the center of each polling booth area, an average of 98 likely voters were then sampled and interviewed within each polling booth.[19] A total of 7,767 interviews were completed in this fashion (4,030 in Gonda and 3,737 in Faizabad) over a period of 5 weeks. A total of 2,860 of these 7,767 persons recruited at their residence subsequently visited one of our "lab" and took part to the lab-based vignette experiment.

As labs were located on the outskirts of each city, and stood in each case very close to a well-known bus station, both rural and urban voters participated to the study (24% of participants to the lab study lived in an urban polling booth). More importantly, these locations allowed the research team to sample from *three* different assembly constituencies in each of the two districts, hence ensuring that the respondents in our sample were represented by a diverse set of politicians.

Both the initial random sample of respondents and the subset of these respondents that subsequently made it to the lab were diverse and roughly representative of their district. Table 1 presents descriptive information for both sets of respondents. As can be seen from the table, the subset of respondents that made it to the lab was very similar to the random sample of the initial "recruitment survey." This unusual two-stage procedure thus does not appear to have led to major self-selection bias.

Lab Procedure

The procedure used to elicit respondents' attitudes at our "lab" limited the potential for misreporting or for influence by both interviewers and bystanders (Chauchard, 2013). Eliciting respondents' attitudes in a lab-like setting first allowed the research team to isolate respondents. Each "lab interview" took place in a separate room, which ensured that no bystanders were present during the interview.[20] Second, the procedure minimized respondents' awareness that partisan and caste-related attitudes were being assessed. Respondents were told that they would participate to an experiment on "how much they know about politics in UP." They were quickly shown a series of pictures of extremely prominent UP politicians from all major parties and asked whether they recognized them (they overwhelmingly did). They were then presented with "a more difficult task." At this point, they were presented the photocopies of bogus news articles (Figure 2), each of which profiled "a likely candidate in the upcoming parliamentary elections in UP." In keeping with the previous task, respondents were *not* invited to read the article but rather simply asked whether they knew whom the pictured politician was. As they could not recognize the man portrayed,[21] nor decipher the article,[22] the interviewer summarized the content of the article to the respondent. At this point, the interviewer provided respondents with the following description (random manipulations italicized):

> Here is what the article is saying: this politician *is currently/previously was* a Member of the State Legislative Assembly. He was elected in (*district name*). He is married and has (*two/three*) children. His name is (*group-specific name*)

Table 1. Comparing the Characteristics of Respondents in the Recruitment and Lab-Based Surveys.

Variable	Value among respondents of *recruitment survey*	Value among participants to *vignette experiment (Lab-study)*
% declare having voted in 2012 elections	92.29	93.97
% declared voting for SP in 2012	46.06	43.89
% cited BJP as preferred party	43.18	44.39
% cited INC as preferred party	14.46	14.68
% cited SP as preferred party	21.74	19.77
% cited BSP as preferred party	20.48	21.10
% urban	23.56	23.97
% married	88.35	88.32
% declared *never* reading news in paper	65.42	63.18
% declared *never* listening news in radio	37.84	35.52
% declared *never* watching news on TV	54.77	56.29
% knows MLA name	86.59	88.19
% knows MP name	74.11	76.08
Age	39.93	40.95
Years living in that location	39.46	40.31
% never went school	20.05	19.05
Average years of schooling	8.65	8.49
% living in *pakka* house	57.30	55.84
Number of rooms in household	3.21	3.08
Number of people in household	7.57	7.39
Household monthly income (in rps)	5,178.74	4,717.16

SP = Samajwadi Party; BJP = Bharatiya Janata Party; INC = India National Congress; BSP = Bahujan Samaj Party; MLA = member of legislative assembly; MP = member of parliament.

and he is the (*caste/community name*) face for the (*party name*). The article also made a thorough investigation of what this politician achieved while in office and it says that villages of his district received (*a lot/very little)* in terms of development when he was in office. Finally, it reports the information provided on the affidavit that candidates have to fill in order to run, and it says that (*he is accused of no major crime/he has four cases pending for major criminal violations*).[23]

Once he had read this description to the respondent, and summarized it a second time, the interviewer used two survey questions—none of which referred to ethnicity—to elicit the respondent's evaluation of the profiled politician. The first item sought to measure respondents' *general* perceptions of the quality of the politician represented in the vignette (a dimension I here-after refer to as *general expectations*):

> Politicians' job is to address their constituents' problems. On a scale from 1 to 5, how good a representative do you think this politician would be in that respect?

The second question aimed to tape into respondents' perceptions of the politician's willingness to help them personally, over other groups and individuals:

> Politicians sometimes help some of their constituents more than they help others. On a scale from 1 to 5, how good a representative do you think this politician would be *for you personally*?

In both cases, responses were on a scale from 1 (*worse*) to 5 (*best*). Overall, the procedure described in this section was repeated three times—that is, each respondent was successively exposed and reacted to three different vignettes.

Tests

To provide intuitive and visual tests of my hypotheses, I take advantage of the randomized nature of these manipulations—which I check in the Online Appendix C[24]—and compare the average rating of different types of profiles.[25] In each case, I focus on responses to the first vignette rated by each respondent, which implies that I rely on the most conservative standard errors possible.[26] In addition, as responses to these two questions were highly correlated (.82),[27] my analyses in the body of this article simply focus on responses to the first item (*general expectations*).[28]

Testing H1: Does the Ethnicity of Candidates Matter?

To test whether respondents took the ethnicity of candidates into account, Figure 3 compares the average rating of all profiles in which the candidates were coethnics to the average rating of profiles in which the candidates were not coethnics. Figure 4 subsequently breaks these results down by party.

Figure 2. Example of a vignette.

These analyses provide clear support for H1: As can be seen from Figure 3, coethnic candidates were rated much more favorably than non-coethnic candidates. This was true regardless of the party label attached to these candidates, as can be seen in Figure 4. In both figures, the 95% confidence intervals around the mean ratings for coethnics and non-coethnics do not
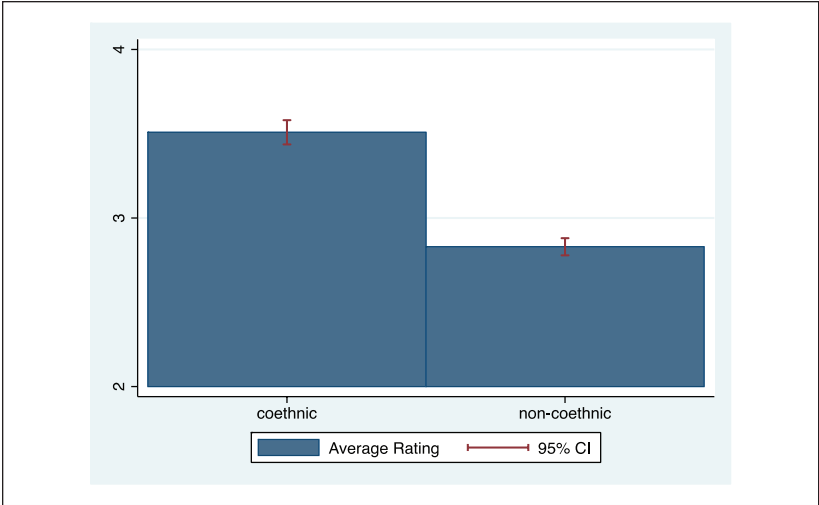
Figure 3. Average *general expectations* by ethnicity (binary).
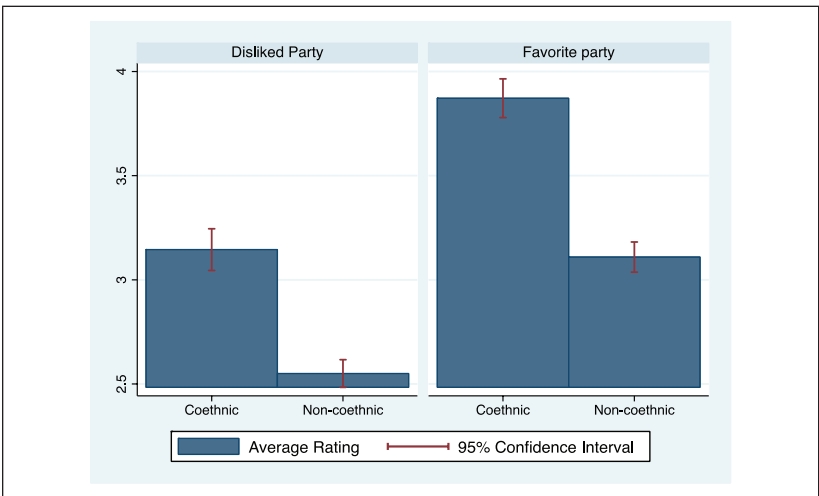CI = confidence interval.

Figure 4. Average *general expectations* by ethnicity (binary) and party.

overlap. This implies that respondents took the ethnicity of candidates into account and that they did so even when the candidate in the vignette belonged to their favorite party.
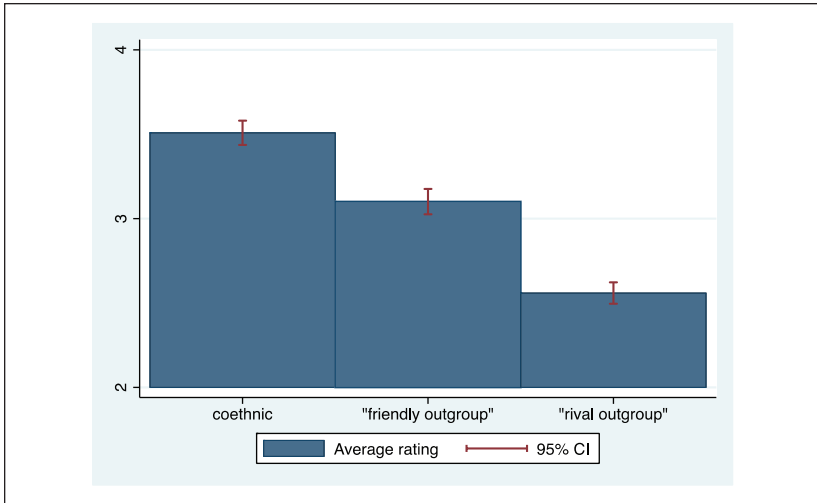
Figure 5. Average *general expectations* by ethnicity.
CI = confidence interval.

These results are remarkable insofar as the design of this experiment did *not* particularly prompt respondents to do so: The amount of information they were presented with did not obviously drive them to react to the *ethnicity* of candidates. Yet, when asked to rate a candidate about whom they were given no fewer than eight pieces of information, respondents strongly reacted to this specific manipulation about ethnicity. Besides, the experiment was *not* presented as an experiment about ethnicity-related attitudes, and respondents could *not* have guessed that the research team was interested in testing this specific hypothesis. The two questions to which they responded after they were exposed to this first vignette were the first questions in the entire lab-based interview that related to political preferences (most of the interview until this point had been about "political knowledge," that is, about how much people *knew* about politicians in UP).

Testing H2: Do Respondents Differentiate Between Non-Coethnics?

To test whether voters differentiate between non-coethnics (H2), Figure 5 breaks down the non-coethnic category into two subcategories and compares the mean rating of all profiles in which the candidates were coethnics with the mean rating of profiles in which the candidates were from "friendly

outgroups" or from "rival outgroups." These analyses provide clear support for H2: While both types of non-coethnics were rated lower than coethnics, non-coethnic candidates from "friendly outgroups" were rated much higher than non-coethnic candidates from "rival outgroups," and none of the confidence intervals around my estimates for these three categories overlap. This confirms that the ethnicity of candidates matters, but it also suggests non-coethnics were not all equivalent in the eyes of respondents. While evaluating non-coethnic candidates, ethnicity-related considerations played a big role in the deliberations of respondents. This was true regardless of the party label under which these candidates competed, as shown in Figure D1 (Online Appendix D).

Altogether, Figures 3 to 5—as well as the corresponding average marginal component effects reported in the Online Appendix B1—suggest that the ethnicity of candidates *strongly* affected evaluations. The size of these candidate-ethnicity effects was large in absolute terms. As can be seen from the estimates presented in Table B1 (Online Appendix B), candidates from "rival outgroups" were, for instance, rated about .85 point lower than coethnics. This is a large and significant difference, suggesting that participants paid a great deal of attention to the ethnicity *of candidates*. While both types of non-coethnics were rated significantly lower than coethnics, the large difference between them in turn implies that the category "non-coethnic" may be substantively meaningless, insofar as the difference between "friendly outgroups" and "rival outgroups" (.54) was significantly larger than the difference between coethnics and "friendly outgroups" (.32). Finally, it is remarkable that the size of these effects was large in comparison with the size of the effects of the other manipulations included in the vignette experiment. This suggests that respondents sometimes weighed the ethnicity of candidates *more* than they weighted the party labels under which those candidates ran. While candidates from a "rival outgroup" were rated .85 point lower than coethnics (Table A1 in the Online Appendix A), candidates described as belonging to the respondents' "third preferred party" (i.e., a *disliked* party) were only rated .56 point lower on average than candidates from the respondent's favorite party (Table A1). This suggests that the ethnicity of *some* candidates mattered at least as much as the party labels under which they reportedly ran.

Testing H3: Do the Non-Ethnic Characteristics of Candidates Matter?

Finally, to test whether the non-ethnic characteristics of candidates matter, I break down the average rating of profiles associated with each type of

Figure 6. Average *general expectations* by ethnicity and perceived record.
CI = confidence interval.

ethnicity (coethnic, "friendly outgroup," "rival outgroup") by their *perceived record* and *perceived criminality* (both of which were randomly manipulated in the vignette experiment). Focusing on the interaction between candidate ethnicity and *perceived record*, Figure 6 shows that participants' ratings evolved as a function of the perceived record of the candidate and that this was true regardless of the ethnicity of said candidate. For both coethnic and non-coethnic candidates, the ratings were significantly lower if the politician was presented as having a bad record (the confidence intervals do not overlap). Focusing on the interaction between ethnicity and criminality, Figure 7 displays a comparable pattern: Candidates presented as being currently charged in four criminal cases (not an uncommon situation in UP) were rated significantly lower than candidates presented as facing no charge at all (the confidence intervals do not overlap).[29] These graphs, as well as row 7 and 8 of Table B1 provide clear support for my third hypothesis: the non-ethnic characteristics of candidates strongly affected voters' preferences. Altogether, these results imply that factors *additional* to ethnicity did drive respondents' evaluations.
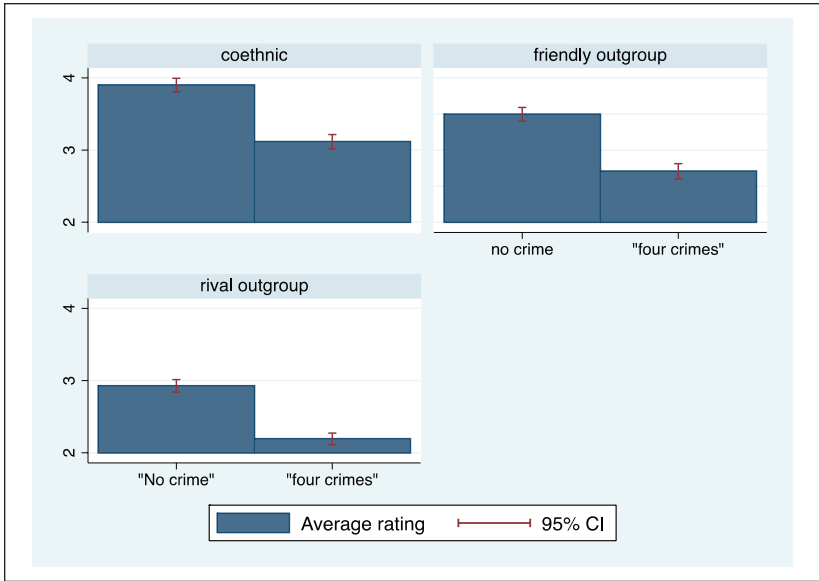
Figure 7. Average *general expectations* by ethnicity and perceived criminality.
CI = confidence interval.

Discussion

These analyses challenge some of the key assumptions of the comparative literature on ethnic parties, as they show that ethnicity was—at least in the context of this study—factored into voters' preferences in a way that is both more complex and less automatic than what frequently cited arguments in this literature (Chandra, 2004; Horowitz, 1985) so far suggest. To better explore the implications of these findings, one, however, needs to assess their validity beyond this context. Accordingly, the following subsections discuss the extent to which the protocol of the study might have affected its conclusions, before exploring the relevance of its findings for other contexts.

Do These Preferences Reflect UP Voters' Preferences?

Insofar as they were expressed in the context of a lab-based survey-experiment, the preferences expressed here cannot be equated to the preferences on which voters rely in actual elections. Several elements, however, suggest that the preferences captured by this study should resemble real-world preferences.

It is, first, important to note that the experiment did *not* prompt respondents to react as strongly as they did to the ethnicity of candidates. Because respondents were faced with a wide range of information about politicians, it is difficult to argue that they strongly reacted to ethnicity because this was the only information about candidates available to them (Conroy-Kurtz, 2013). However, social desirability concerns suggest that the effects of ethnicity I detect should be seen as *underestimates*, especially when it comes to the effect of "rival outgroup" candidates. This is due to a practical aspect of my design: When respondents were asked as part of their recruitment interview to name groups with whom their own group has maintained tense relationships with—a sensitive question asked in a face-to-face interview—efforts to limit response editing were limited. As a result, it is likely that some respondents misrepresented their views and named groups other than those they might have named in private, or with a more confidentiality-enhancing administration mode. This implies that some of the "rival outgroup" politicians likely do not belong to groups that respondents truly consider to be "rival outgroups," which in turn implies that the results above likely *under*estimate the effect of ethnicity. This, importantly, means that the lab-based survey constituted a hard test for H1 and H2.

These results are, in addition, unlikely to be the product of two other features of my research design. The fixed order in which manipulations appeared in the vignette is unlikely to have inflated my results. This is, first, because the order was *in practice* not as fixed as it might have been if respondents had self-administered the survey. After interviewers described the vignette to respondents a first time, interviewers checked that respondents had heard and understood the information included in the vignette. As they asked respondents to summarize what they had heard, the information about candidates was often repeated in a different order, as respondents often struggled to remember some of the seven randomly varied pieces of information contained in the vignette. Second, and more generally speaking, it is not clear that "row-order effects" always exist in complex experiments of this type: Hainmueller, Yamamoto, and Hangartner (2015), for instance, do not detect any.[30] Third, potential row-order effects would likely run *against* my main finding. In the vignette, respondents heard about the ethnicity of candidates *first*. This should lead us to expect that they were more strongly primed by the other treatments than by the candidate-ethnicity one. If that were the case, it would imply that my design potentially underestimates the effect of candidate ethnicity and that support for H1 and H2 might have been even larger than it already is in the results presented here.

These findings are similarly unlikely to owe to the structure of my dependent variable. Vignette experiments can be of two types: "rating tasks"

involving one profile at a time (as in this experiment) and "choice tasks" requiring respondent to pick one of several randomly varied profiles presented simultaneously. While choice tasks are seen by some as being closer to real-world decision making, in which respondents must cast a single ballot between competing candidates who vary on multiple dimensions, they do not encompass any of the strategic aspects of voting in the real world. Besides, recent works (Hainmueller et al., 2014; Hainmueller et al., 2015) suggest that the estimates retrieved from either type of tasks are often comparable. This may be seen as evidence for the fact that preferences retrieved using either type of tasks do not differ and/or that they both significantly differ from preferences in real-world decision making.

It is, finally, unlikely that respondents reacted to the record and the criminal status of candidates because they felt compelled to reject "bad types" in front of an interviewer. Several elements point in that direction. The way respondents reacted to the ethnicity treatments suggests, first, that respondents may not have been exceedingly worried by social desirability as they rated these candidates. Caste is a sensitive topic in India and one that is prone to editing by respondents (Chauchard, 2013). Yet my results show that respondents had no problem endorsing coethnics and rejecting non-coethnics in front of their interviewer, who usually had a sense of their own caste, based on their name and appearance. Second, the large number of treatments included in the experiment would have made it easy for respondents to endorse criminals without necessarily appearing to do so, as interviewers would not necessarily have been able to deduce which characteristics (among a total of eight) drove the respondent's evaluation. Because interviewers specifically did not ask *why* respondents were rating these candidates the way they were rating them, it is not clear that respondents should have felt any pressure at all to reject candidates with a criminal background. Third, it would have been relatively unproblematic for respondents to disregard information about criminality and record, as neither the veracity of this information nor the impartiality of these charges was clear. Yet they overwhelmingly reacted to these prompts. These estimates may be seen as conservative, insofar as the effect of a positive development record might have been larger had respondents had a personal experience with these politicians—which they did not—or had respondents been able to verify that the information in the newspaper article was accurate—which they were not able to do.[31]

Implications

While I see no obvious reason to believe that the protocol of the study would have exacerbated respondents' reactions to candidate ethnicity, or prompted

them to penalize politicians with unattractive pedigrees, two important caveats are in order. The *timing* of the study may first have helped increase respondents' emphasis on candidate-level characteristics. Because the study took place *in between* elections, party preferences likely were in motion at the time the study took place. It is thus possible that the study underestimated the effect of party labels and/or that the voters interviewed in the study weighed candidates more heavily than they might have closer to the actual election. Further studies would accordingly need to measure the evolution of the weights voters place on candidate-level characteristics and party preferences *over the course of the electoral cycle*. As party preferences crystallize, and as the "labels" of parties become clear to voters, it is possible that voters weigh candidates less. Given the size of the effects of candidate ethnicity detected here, it is, however, unlikely that these effects would entirely disappear as elections get closer. Second, because I chose to focus on two mostly rural districts within the state of UP (a logical choice given that the state's population is 76% rural), the preferences retrieved in this study may not be entirely representative of the preferences of urban voters, who tend to be more educated, more informed, and overall less influenced by ethnicity in elections.

In spite of these limitations, these results are relevant to a range of contexts. The voters who participated to this study do not constitute outliers in India. The state of UP—the largest Indian state, with a population of more than 200 millions—has over the last decades acquired the reputation of a state riddled with high levels of "ethnic politics" and with particularly polarized intergroup relations. Yet this does not necessarily imply that ethnicity is not relevant to elections taking place in other Indian states (Chandra, 2004; Jaffrelot, 2010), nor that other Indian voters would place more emphasis on party-level factors than the voters of UP. As can be seen from the Online Appendix E, data from the 2009 National Election Study (the most recent available data based on a national sample of voters[32]) show that the attitudes of UP voters were altogether not that different from the attitudes of other Indian voters in these respects: They similarly declared weighing parties more than candidates (Table E1 in the Online Appendix E); they similarly did not feel particularly close to any political parties (Table E2); besides, as can be seen from Tables E3 to E6 (and even more clearly from Figures E1 to E4), UP voters were not the only Indian voters to endorse openly anti-minority or "casteist" statements;[33] finally, as shown in Tables E7 and E8, UP voters were as likely as others to declare that they would pick a corrupt or a criminal representative. Altogether, this additional data suggest that UP voters are unlikely to weigh the various characteristics of candidates in a way that radically differs from the way the median Indian voter weighs them. While this is limited data, and while further studies will need to confirm this, this suggests that the

aforementioned results might not have been radically overturned had I sampled across North India—rather than in just two locations.

These results also likely inform the behavior of voters from a variety of ethnically diverse democracies *beyond India*. Ethnicity often becomes a heuristic device when relatively uninformed voters hope to access state benefits controlled by elected politicians (Chandra, 2004; Posner, 2005). In such contexts, voters similarly select representatives on the basis of single-member constituencies—this is the case in most commonwealth nations. The logic articulated above is likely to hold in these contexts: "ethnic voters" are likely to weigh both party-level and candidate-level factors as they evaluate candidates. Moreover, because the fragmented ethnic context of North India is not unique—many African nations are for instance as ethnically fragmented—these voters are as likely to differentiate between different types of non-coethnics as Indian voters.

In light of these elements, these results have implications for the comparative literature on ethnic voting and ethnic parties. While popular theories of ethnic voting (Chandra, 2004; Horowitz, 1985) focus on the influence of ethnicity *at the party level*, voters may react to the ethnicity of candidates as much as they reacted to the party labels under which these candidates run. This implies that ethnicity-related considerations enter voting decisions in several ways and that current models focused on the effect of ethnicity *at the party level* misrepresent *how* ethnicity matters. This further suggests that a change of candidate at the constituency level can change the electoral fortunes of a party, a finding coherent with arguments about African and Latin-American politics suggesting that ethnically identified parties can alter their image and attract new voters by fielding new types of candidates (Ferree, 2011; Madrid, 2012).

Second, these results show that the impact of ethnicity is not limited to a simple preference for coethnics. Voters do prefer coethnics to other types of candidates, but they also clearly differentiate between non-coethnics, rating some of them nearly as positively as their own coethnics and strongly penalizing others—suggesting that ethnicity matters *positively* as well as *negatively*. This has far-reaching implications—beyond electoral politics—for how we should think of intergroup relations in ethnically fragmented societies: In the absence of a simple, binary dichotomy between coethnics and non-coethnics, we should expect voters' decisions and more generally speaking citizens' decisions, to rely on ethnic considerations in a much larger set of situations, including those in which none of their coethnics is involved.

Finally, contrary to what both "primordialist" (Horowitz, 1985; Rabushka & Shepsle, 1972) and constructivist models (Chandra, 2004) have often implied, these results suggest that ethnicity is unlikely to provide candidates

with an *absolute* advantage: When coethnic candidates have problematic personal characteristics, voters—*ceteris paribus*—rate them more negatively, hence suggesting that non-ethnicity-related candidates-level factors matter, as for their American counterparts (Campbell, Green, & Layman, 2011; McDermott, 1997, 1998; Popkin, 1991). Altogether, these findings thus imply that the voters of new and post-colonial democracies engage in complex, multi-dimensional calculations about their political representatives. While models of ethnic voting should remain parsimonious, they should better acknowledge this degree of sophistication.

Acknowledgments

Declaration of Conflicting Interests

Funding

Notes

1. Similarly, voters may not consider the ethnicity of candidates when parties tend to be ethnically homogeneous. If almost all candidates from a given party belonged to the same group, voters may assume the ethnicity of candidates or consider it relatively unimportant. However, as the aforementioned works on the diversification of ethnic parties suggest, situations of this type may not be set in stone (Chandra, 2004; Ferree, 2011; Madrid, 2012).
2. A dissonant voice on this is Madrid (2012), which suggests that programs matter.
3. While voters should similarly be expected to take the ethnicity of candidates into account in multi-members constituencies, such a setting might have required more complex tests.
4. While the Samajwadi Party (SP) originally grew as a party led by Yadavs and catered to the interests of Yadavs, it is now seen as a party supported by a number of other groups, especially so Muslims. Similarly, while the Bahujan Samaj Party (BSP) has historically been built around the large Chamar population of Uttar Pradesh (UP), it has in recent times pandered to other groups.

5. Recent evidence also suggests that candidates, more generally speaking, play a bigger role than what has so far been expected in India (Ziegfeld, in press).

6. While the SP has over the years established itself as a party serving the interests of Yadavs, it has for instance not established the same reputation among Kurmis, also classified as "Other Backward Castes." As a result, studies that evaluate the strength of ethnic voting based on these other dimensions—as Banerjee, Green, McManus, and Pande (2014)—can only reveal very little about patterns of ethnic voting.

7. The results are extremely similar in both locations, which suggest that these results may not be the product of idiosyncratic features of my research site. While I could have selected a less proximate district as a second location for the experiment, this would have limited my ability to monitor the experiment, as data collection took place simultaneously in the two locations to avoid potential time effects.

8. No data allow me to evaluate the extent to which this feature of my design affected my results, as *all* vignettes profiled candidates who were fictional and were presented as contenders *in other districts*. This is admittedly a limitation of this design, as voters might react differently to politicians from their local area. Assuming that respondents' evaluations of politicians from their district and politicians from other districts are strongly correlated, the benefits of focusing on fictional politicians—namely, my ability to randomly vary these profiles—in my mind, however, overweigh this limitation.

9. Given that partisanship tends to increase in electoral periods (Michelitch & Utych, 2015), this means that these results may underestimate the importance of party labels. I return to this question in the "Discussion" section below.

10. As the vignettes provide respondents with many types of information about the politician, it is difficult to argue that respondents reacted to the ethnicity of the purported candidates because it was the sole salient information about them.

11. Both in the case of the picture and in the case of the district, these sets of alternative were chosen based on the result of a large pretest in which these alternatives created no significant differences in the outcome of interest. The picture was randomly chosen among a set of three possible pictures, reproduced in Figure 2. The politician's district was randomly chosen among a set of three possible districts of UP: *Jalaun*, *Pratapgarh*, and *Unnao*. The politician's family situation was either "married with two children" or "married with three children." Finally, the politician was described as being either a *current* or a *former* Member of the State Legislative Assembly.

12. As politically salient ethnic groups are fairly small groups in UP, a large number of politically salient ethnic categories existed. Our respondents, drawn from only two districts, altogether cited 45 different caste or religion-based groups when answering a series of open-ended questions asking them to mention groups that were playing a role in politics. In light of this, randomly varying candidate ethnicity would have required an extremely large number of experimental treatments and hence an extremely large *N*. To address this issue, I limited

the categories of candidate ethnicities and candidate party affiliations that the respondents would be exposed to these three categories. The following section clarifies how "friendly" and "rival" outgroups were concretely measured.

13. Respondents' *third favorite party* clearly is a very *disliked* party in the context of the polarized politics of UP. Because a majority of respondents was unable to rank more than three parties—the India National Congress (INC) being a non-player in both districts—we could not measure respondents' fourth-ranked party. The following section also clarifies how partisan preferences were measured.

14. Namely, the script said, "villages of his district received (*a lot/very little*) in terms of development when he was in office."

15. All of these elements are frequently reported in recent profiles of political candidates in both the English and vernacular press.

16. A pretest suggested that many respondents would have doubted the credibility of such a profile, even if the Bharatiya Janata Party (BJP) fields a small number of Muslim candidates in UP.

17. The research team in each case used the category each respondent had referred to in this recruitment interview and did not constrain or reformulate this choice to fit the respondent's self-reported identity in a pre-established list. In each case, a credible group-specific name was created and adjoined to that caste or religious community. The same name was subsequently used for all profiles of politician from the same group.

18. To incentivize respondents to visit our lab, they were offered 50 rupees and a ticket to a lottery in which a basket of more substantial gifts was at stake. We also organized a shuttle service from their neighborhood on the day they were asked to visit the lab.

19. Investigators attempted to interview each respondent three times before they were "replaced" by another randomly selected respondent.

20. Organizing interviews in a lab setting also allowed the research team to have full control over the protocol followed by interviewers, which is an underestimated but important advantage, especially in the case of a complex vignette experiment.

21. He was *not* an actual representative, and only 1 out of 2,861 respondents claimed to know the name of one of these politicians.

22. Because the photocopies were deliberately low quality and the title of the article was blackened, respondents could not read the text of the article. While the vignette experiment, thus, relied on realistic news article to enhance respondents' perceptions that the profiled politicians were real politicians, the only information from these articles that respondents could properly analyze was the picture of the politician.

23. The information contained in this summary and the phrasings used were adapted from actual profiles of politicians collected in the local Hindi press. Such profiles frequently mention the community of the candidate, in a way that is coherent with the vignette.

24. Table C1 (Online Appendix C) shows that the randomization procedure actually produced well-balanced groups by regressing a series of key respondent

characteristics (age, education, income, etc.) on two key attributes featured in the vignettes: The fact that none of these coefficients comes even close to usual significance levels suggests that the randomization strategy was effective.

25. Following Hainmueller, Hopkins, and Yamamoto (2014), this design also allows me to retrieve the *average marginal component effect* (AMCE) of all eight manipulations, as detailed in the Online Appendix B.

26. All results hold if I include all three vignettes. Standard errors become much *smaller*.

27. There are two possible reasons—among which I am unable to differentiate—as to why responses to these two questions were so highly correlated. The first one is that patronage may be so ingrained that there are almost no voters for whom the public and particularistic interests diverge. The second—and in my view more likely—explanation, however, has to do with a "coherence effect." After they had been asked the more general question, a large proportion of respondents likely felt compelled to provide a similar evaluation to the second item, despite interviewers' efforts to differentiate between the two items.

28. Analyses based on *expectations of targeted benefits* can be found in Table B2 of the Online Appendix B. As can be seen by comparing Table B1 to Table B2, the effects described here are slightly smaller compared with the effects I detect on this second item.

29. The fact that respondents overall penalized candidates with a criminal background does *NOT* imply that they systematically rated such candidates lower than "clean" politicians. As detailed in a companion paper relying on the same data (Chauchard, 2015), and as can be seen from Table B1, respondents weighed considerations about criminality significantly *less* than they weighed considerations about ethnicity. Voters penalized both *non-coethnic* candidates and *criminal* candidates in their evaluations; non-coethnics from "rival outgroups" were, however, penalized significantly more strongly than "criminals." Because of this, respondents rated some of these criminal candidates more favorably than some of the "cleaner" candidates they were presented with, *even though they overall penalized candidates with criminal backgrounds*. This suggests that the criminalization of assemblies in India does not necessarily owe to voters' indifference to or preference for criminals, but also potentially to the fact that Indian voters are trapped in ethnicity-based decisions: Voters dislike criminal candidates, yet sometimes prefer these candidates to other, "cleaner" ones because they weigh ethnicity more than they weigh probity.

30. As noted by Hainmueller et al. (2014), "experimental tests in the literature on consumer choice . . . have found that for most consumers the quality of integrated decision making decreases only once the choice task includes more than ten attributes." (p.26)

31. Note, in addition, that the question did not ask respondents to like or endorse a politician. It more specifically asked them *how good they thought a politician would be at addressing their constituents' problems*. Hence, respondents who

agreed that a candidate with a criminal background would serve his constituents effectively would not necessarily have been seen as endorsing that candidate.

32. As of June 2015, this is the latest National Election Study that is publicly available.

33. This is true of all four relevant items that were included in this survey. While UP voters were in some cases among the most likely to endorse these statements (see Figures E3 and E4 especially), they never were outliers in this respect, with many other Indian voters—especially in large North Indian states such as Rajasthan or Bihar—endorsing these statement at the same rates.

References

Aidt, T., Golden, M. A., & Tiwari, D. (n.d.). *Criminality in the Indian legislature*. Retrieved from http://polisci2.ucsd.edu/dtiwari/research.html

Akerlof, G. A., & Kranton, R. E. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives*, *19*, 9-32.

Banerjee, A., Green, D. P., McManus, J., & Pande, R. (2014). Are poor voters indifferent to whether elected leaders are criminal or corrupt? A vignette experiment in rural India. *Political Communications*, *31*, 391-407.

Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R., & Topalova, P. (2009). Powerful women: Does exposure reduce bias? *Quarterly Journal of Economics*, *124*, 1497-1540.

Berenschot, W. (2011). On the usefulness of Goonda's in Indian politics: "Moneypower" and "musclepower" in a Gujarati locality. *South Asia: Journal of South Asian Studies*, *34*, 255-275.

Besley, T. (2005). Political selection. *Journal of Economic Perspectives*, *19*(3), 43-60.

Birnir, J. K. (2007). *Ethnicity and electoral politics*. Cambridge, UK: Cambridge University Press.

Bratton, M., & Kimenyi, M. S. (2008, March). *Voting in Kenya: Putting ethnicity in perspective* (University of Connecticut: Department of Economics Working Paper Series). Mansfield: University of Connecticut.

Campbell, D., Green, J. C., & Layman, G. C. (2011). The party faithful: Partisan images, candidate religion, and the electoral impact of party identification. *American Journal of Political Science*, *55*, 42-58.

Chandra, K. (2004). *Why ethnic parties succeed*. Cambridge, UK: Cambridge University Press.

Chattopadhyay, R., & Duflo, E. (2004). Women as policy makers: Evidence from a India-wide randomized policy experiment. *Econometrica*, *72*, 1409-444.

Chauchard, S. (2013). Using MP3 players in surveys: The impact of a low-tech self-administration mode on reporting of sensitive attitudes. *Public Opinion Quarterly*, *77*, 220-231.

Chauchard, S. (2015). *Is "ethnic politics" responsible for "criminal politics"? A vignette-experiment in North India* Unpublished manuscript, Dartmouth College, Hanover, NH.

Chauchard, S. (in press). *Political representation and intergroup relations in rural India: The psychological impact of ethnic quotas*. Cambridge, UK: Cambridge University Press.

Clarke, H., & Stewart, M. (1998). The decline of parties in the minds of citizens. *Annual Review of Political Science*, *1*, 357-378.

Conroy-Kurtz, J. (2013). Information and ethnic politics in Africa. *British Journal of Political Science*, *43*, 345-373.

Ferree, K. (2011). *Framing the race in South Africa: The political origins of racial census elections*. Cambridge, UK: Cambridge University Press.

Hainmueller, J., Hopkins, D. J., & Yamamoto, T. (2014). Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis*, *22*, 1-30.

Hainmueller, J., Yamamoto, T., & Hangartner, D. (2015). Validating vignette and conjoint survey experiments against real-world behavior. *Proceedings of the National Academy of Sciences*, *112*, 2395-2400.

Hansen, T. B. (2001). *Wages of violence: Naming and identity in postcolonial Bombay*. Princeton, NJ: Princeton University Press.

Horowitz, D. L. (1985). *Ethnic groups in conflict*. Berkeley: University of California Press.

Jaffrelot, C. (2010). Caste and political parties in India: Do Indians vote their caste—While casting their votes. In A. Kohli & P. Singh (Eds.), *Routledge handbook of Indian politics*. London, England: Routledge.

Khan, A. (2012, February 25). Cane and caste politics in Western Uttar Pradesh. *The Hindu*. Retrieved from http://www.thehindu.com/news/national/other-states/cane-and-caste-politics-in-western-uttar-pradesh/article2929587.ece

Lindberg, S., & Morrison, K. C. (2008). Are African voters really ethnic or clientelistic? Survey evidence from Ghana. *Political Studies Quarterly*, *123*, 95-122.

Lupu, N. (2013). Party brands and partisanship: Theory with evidence from a survey experiment in Argentina. *American Journal of Political Science*, *57*, 49-64.

Madrid, R. (2012). *The rise of ethnic politics in Latin America*. New York, NY: Cambridge University Press.

McDermott, M. L. (1997). Voting cues in low-information elections: Candidate gender as a social information variable in contemporary United States elections. *American Journal of Political Science*, *41*, 270-283.

McDermott, M. L. (1998). Race and gender cues in low-information elections. *Political Research Quarterly*, *51*, 895-918.

Michelitch, K., & Utych, S. (2015). *Electoral cycle fluctuations in partisanship intensity: Global evidence* [Southern Political Science Association Conference Paper].

Michelutti, L. (2008). *The vernacularisation of democracy: Politics, caste and religion in India*. New Delhi, India: Routledge.

Pai, S. (2002). *Dalit assertion and the unfinished democratic revolution: The Bahujan Samaj Party in Uttar Pradesh*. New Delhi, India: SAGE.

Pepinsky, T., Liddle, R. W., & Mujani, S. (2012). Testing Islam's political advantage: Evidence from Indonesia. *American Journal of Political Science*, *56*, 584-600.

Popkin, S. L. (1991). *The reasoning voter: Communication and persuasion in presidential campaigns*. Chicago, IL: University of Chicago Press.

Posner, D. (2005). *Institutions and ethnic politics in Africa*. Cambridge, UK: Cambridge University Press.

Rabushka, A., & Shepsle, K. (1972). *Politics in plural societies*. Columbus, OH: Charles E. Merrill.

Srivastava, T. (2012, February 1). Caste-ing the vote in Uttar Pradesh. *Sify News*. Retrieved from http://www.sify.com/news/caste-ing-the-vote-in-uttar-pradesh-news-columns-mcbr1tdachf.html

Suryanarayan, P., & Huber, J. (in press). Ethnic inequality and the ethnification of political parties: Evidence from India. *World Politics*.

Tiwari, D. (n.d.). *Quality is in the eye of the beholder: Electoral competition and candidate quality in India*. Retrieved from http://polisci2.ucsd.edu/dtiwari/research.html

Tiwari, M., & Pandey, R. (2013). *Battleground U.P.: Politics in the land of Ram*. New Delhi, India: Tranquebar Press.

UP assembly elections 2012: The caste matrix in Awadh. (2012, February 2). *Economic Times*. Retrieved from http://articles.economictimes.indiatimes.com/2012-02-02/news/31017349_1_caste-avadh-rae-bareli

Vaishnav, M. (2012). *Social divisions, credibility and criminality: Political selection in India*. Unpublished manuscript.

Ziegfeld, A. (in press). Candidate characteristics in Indian elections: Who wins votes? *Asian Survey, 25*.

Author Biography

Simon Chauchard is an assistant professor of political science at Dartmouth College and the author of *Political Representation and Intergroup Relations in Rural India: The Psychological Impact of Ethnic Quotas* (in press, Cambridge University Press).