Candidate Inconsistency and Voter Choice

Michael Tomz
Stanford University
tomz@stanford.edu

Robert P. Van Houweling
University of California, Berkeley
rpvh@berkeley.edu

Draft: August 2009

Preliminary. Comments welcome!

1. Introduction

The consistency of the positions that candidates take over their careers often becomes an issue in political campaigns. George W. Bush famously accused John Kerry of "flip-flopping" from supporting to opposing the Iraq war as public opinion turned against the conflict. Many pundits viewed the attack on Kerry's inconsistency as successful because it diverted attention from the relatively popular position he had recently adopted.[1] Others argued, however, that Kerry would have received even fewer votes if he had clung to his outmoded position, instead of reversing his stance on the Iraq war.

The Bush-Kerry controversy raises several general questions about the consequences of candidate inconsistency. How much attention do voters pay to the past positions of candidates? To what extent do past positions influence expectations about what candidates will do in the future? Are voters less likely to support candidates simply because they have changed their positions over time? And if there is a cost to inconsistency per se, is it large enough to prevent politicians from abandoning positions that people no longer support? The existing political science literature does not resolve these fundamental questions.

Understanding the consequences of inconsistency is important not only for assessing candidate strategies and tactics, but also for evaluating the effectiveness of representation in democracies. Voters who punish candidates for changing their positions may unwittingly deter candidates from adapting as new policy-relevant information comes to light. Moreover, penalties for inconsistency may contribute to polarization and legislative gridlock. Office holders could be wary of compromising with political opponents, for fear that even voters who agree with the content of the compromise will nevertheless punish candidates for veering from their past commitments.

We designed and fielded a survey experiment to shed light on the consequences of candidate inconsistency. Our experiment shows that voters consider not only the present but also the past positions of candidates. Candidate inconsistency influences voter behavior in two ways: by triggering a negative reaction to inconsistency in itself, and by leading citizens to discount the current-day promises of candidates who advocated different policies in the past. Both mechanisms reduce the attractiveness of flip-flopping. Nevertheless, we identify situations in which the electoral gains from reversing course outweigh the electoral costs. Overall, our experiments expose the penalties for changing positions, while also illuminating when doing so can be an optimal strategy.

In the remainder of the paper, we develop several hypotheses about the effects of candidate inconsistency. We then describe our experimental procedure, present our findings, and discuss the implications for democratic politics.

[1] "Despite Bush Flip-Flops, Kerry Gets Label," *The Washington Post*, September 23, 2004.

2. Theory and Hypotheses

Political scientists disagree about both the prevalence and the effects of candidate inconsistency. Some scholars argue that candidates are typically consistent over time. In their examination of congressional roll-call records, for example, Poole and Rosenthal (1997) uncover few examples of ideological shifts in the voting patterns of individual representatives. Other scholars add that when party positions on issues such as civil rights, abortion, national defense, and foreign trade have shifted, the change has occurred primarily because of the arrival of new representatives, not because of changes in the voting patterns of current representatives (e.g. Carmines and Stimson 1989, Adams 1997). In a recent reexamination of the data, however, Karol (2009) finds substantial intertemporal inconsistency across a wide range of issues. Thus, the frequency of inconsistency remains controversial.

Scholars disagree not only about the frequency but also about the consequences of taking inconsistent positions. Some emphasize the negative, by arguing that inconsistency entails high reputational costs (Downs 1957; Poole and Rosenthal 1997; Tomz 2007) and creates uncertainty that repels risk-averse voters (Bernhardt and Ingbergman 1985). Others emphasize the potential benefits of inconsistency. Karol (2009), for example, contends that office holders regularly abandon their past positions in order to reap the benefits of popular new policy stances.

In this paper, we focus on the second controversy. We propose and test the following hypotheses about the effects of candidate inconsistency.

> Hypothesis 1: Voters dislike inconsistency in itself, and therefore punish candidates for changing positions.
>
> Hypothesis 2: Voters discount the current policy pronouncements of candidates who have taken different positions in the past.
>
> Hypothesis 3: Despite the limitations implied by the previous hypotheses, inconsistency is sometimes electorally beneficial, because the gains associated with moving toward popular positions can outweigh the costs of being perceived as a flip-flopper.

3. Experimental Design

We designed an experiment to assess the effect of candidate inconsistency on voter choice and embedded it in an Internet survey of a nationally representative sample of U.S. adults. Our experiment involved three steps. First, we measured respondents' preferences about taxes on wealthy Americans, defined as people who make more than $250,000 per year. Second, using the same issue, we described the policy statements of two candidates, who varied randomly in the consistency with which they espoused positions, and asked which candidate the respondent preferred. Finally, we asked respondents to state their expectations about what one of the candidates would do if elected.

3.1 Measuring Respondent's Policy Preferences

The survey began by asking respondents about a key political issue: taxes on wealthy Americans. We noted that "Some people think the government should increase taxes on wealthy Americans, defined as people who make more than $250,000 per year. Other people think taxes on wealthy Americans should be kept at their current level. And still other people think the government should decrease taxes on wealthy Americans." We then asked whether the respondent thought the government should increase taxes on wealthy Americans, keep taxes on wealthy Americans at their current level, or decrease taxes on wealthy Americans (Figure 1a).

[Figure 1 about here]

The positions were described verbally rather than numerically, though for convenience we sometimes use integers as shorthand for the three fully labeled options. In those instances, the integer one represents the most liberal position (increase taxes on wealthy Americans), two represents an intermediate position (keep taxes the same), and three represents the most conservative position (decrease taxes on wealthy Americans).

3.2 Measuring Preferences about Candidates

We then asked respondents to choose between two candidates, who had previously expressed opinions about taxes on wealthy Americans. "Non-partisan groups often survey candidates about tax rates for wealthy Americans," we noted. "We would like your views about two candidates, whose names will remain confidential. They are Candidate A and Candidate B" (Figure 1b). By denoting candidates with letters, we were able to test theories about the effects of inconsistency in their purest form, with out the potentially confounding effects of party or other candidate attributes.

We reported what the candidates had said on two occasions: two years ago and this year. On each occasion, a candidate either called for increasing taxes on wealthy Americans or called for decreasing taxes on wealthy Americans. Denoting these policy positions as I and D, respectively, each candidate's history of statements on the issue can be summarized as (I,I), (D,I), (I,D), or (D,D), where the first letter represents what the candidate said two years ago, and the second letter represents what the candidate said this year. Candidates with a history (I,I) have consistently advocated higher taxes on the rich, whereas candidates with a history (D,D) have repeatedly proposed lower taxes, and candidates with histories (D,I) or (I,D) have espoused inconsistent positions.

After eliminating cases in which the candidates had identical histories, and could therefore be expected to tie, we were left with six possible candidate pairings, which are displayed in Table 1. We randomly assigned each respondent to receive one of these six pairings, and randomly labeled one of the candidates as A and the other as B. Respondents were then asked, "On this issue, which candidate do you prefer?"

[Table 1 about here]

3.3 Measuring Expectations

The final phase of the experiment investigated how citizens form expectations about the likely positions of candidates. Each respondent was reminded of the opinions candidate A had expressed: (I,I), (D,I), (I,D), or (D,D). "If you had to guess," we asked, "what do you think Candidate A would try to do if elected." The response options were increase taxes on wealthy Americans, keep taxes on wealthy Americans at their current level, or decrease taxes on wealthy Americans (Figure 1c).

4. Data and Findings

The experiments discussed in this paper were administered by Knowledge Networks, an Internet-based polling firm. Knowledge Networks uses random digit dialing to recruit participants and provides Internet access to households, resulting in a nationally representative sample of adults. The interviews took place in February–March 2009, and 825 people (about 60% of invitees) agreed to take the survey.

4.1 The Average Effect of Inconsistency

According to Hypothesis 1, voters dislike inconsistency and therefore punish candidates for taking inconsistent positions. As first step toward testing this hypothesis, we computed the proportion of times that respondents voted for an inconsistent candidate (either D,I or I,D) in pairwise competition against all possible opponents, both consistent and inconsistent. Similarly, we computed the proportion of times that respondents preferred a consistent candidate (one with history I,I or D,D) in races against all conceivable opponents. The difference between these two values represents the average effect of inconsistency.

As noted earlier, we did not present respondents with scenarios in which the two anonymous candidates had identical histories, because respondents would have found no basis for preferring one over the other. The choice between Candidate A and Candidate B would have reduced to the flip of a coin. In actual political competition, of course, candidates sometimes take identical positions, so a complete set of opponents would include those with policy histories exactly like one's own. We reincorporated these ties into our analysis by assuming that, if candidates had competed against their clones, each would have received 50 percent of the vote.

Table 2 confirms our hypothesis that voters dislike inconsistency. On average, inconsistent candidates received 43 percent of the vote, whereas consistent ones received 57 percent. The estimated cost of inconsistency was, therefore, 14 points. The 95 percent confidence interval around this estimate ranged from –17 to –10, enabling us to reject the null hypothesis that voters are indifferent about inconsistency. On the contrary, inconsistency entails a heavy political cost.

[Table 2 about here]

4.2 How Inconsistency Shapes Expectations

Inconsistency not only evokes a negative reaction from voters, but also shapes expectations about the policies a candidate will pursue. We hypothesized that voters look skeptically upon the current policy positions of candidates who espoused different policies in the past. To test this hypothesis, we asked each respondent what they thought Candidate A would do if elected.

Our evidence strongly supports the discounting hypothesis. Some candidates in our study consistently advocated higher taxes on the rich. When asked about such candidates, 81 percent of respondents thought the candidate would try to increase taxes. Only 16 percent felt the candidate would keep taxes at the same level, and a miniscule 3 percent guessed that the candidate would try to reduce taxes. On a scale from 1 to 3, where one represents an effort to increase taxes and three represents an effort to decrease taxes, the average expectation about candidate (I,I) was (.81 x 1) + (.16 x 2) + (.03 x 3) = 1.22. Figure 2 summarizes this expectation. The dot below the label (I,I) marks the average position of 1.22, and the associated horizontal line represents a 95 percent confidence interval around the mean.

[Figure 2 about here]

Voters were significantly more skeptical of candidates who switched from advocating lower taxes two years ago to calling for higher taxes today. Only 53 percent of citizens who evaluated this type of candidate expected him to strive for higher taxes. Fully 36 percent guessed that the candidate would maintain the status quo, and 11 percent expected the candidate to lobby for tax cuts. The mean placement of candidates with history (D,I) was, therefore, 1.58. Clearly, voters had less faith in the pro-tax intentions of candidates who only recently embraced that position than of candidates who had proposed it repeatedly.

Expectations about the other candidates were just as consistent with the discounting hypotheses. On average, voters placed candidates with histories (D,D) at 2.48, while locating candidates with histories (I,D) at 2.01. Thus, candidates who called for lower taxes this year were seen as less credible if they had previously advocated higher taxes. The estimates in Figure 2 confirm, therefore, that voters doubt the sincerity of candidates whose current platforms contradict their previous ones.

4.3 When Inconsistency can be Advantageous

Although voters dislike inconsistency in general and look skeptically upon the recent promises of inconsistent candidates, flip-flopping may nonetheless be beneficial in certain circumstances. Hypothesis 3 states that candidates may be able to improve their lot by switching to a position that better resonates with voters. The benefits of switching to a now-popular option may, in fact, outweigh the costs of being branded as a flip-flopper, even if voters do not fully believe the sincerity of the new position.

To evaluate this hypothesis, we first need to establish what kinds of policies voters would like to see. Table 3 shows that, in our sample, more than 63 percent of respondents think the

government should increase taxes on wealthy Americans, while only 7 percent want the government to decrease taxes on wealthy Americans. Thus, notwithstanding the popular disdain of flip-floppers, candidates who proposed lower taxes two years ago may do better by abandoning their unpopular stance, instead of reiterating a position that rubs most voters the wrong way.

[Table 3 about here]

Table 4 confirms this hypothesis. The bold numbers in the table represent the average support each type of candidate received in competition against all possible opponents. The arrows denote changes from consistent to inconsistent positions, and the italicized numbers give the estimated effect of each policy reversal. As the table shows, a candidate who called for lower taxes two years ago and clung to the same position this year could expect to receive 42 percent of the vote. If the same candidate flip-flopped, though, by advocating higher taxes this year despite advocating lower taxes previously, he could secure around 51 percent of the vote. Overall, flip-flopping would yield a 9-point gain, enough to transform the candidate from a certain loser into a potential winner. The confidence interval around this effect ranges from 5 to 15, giving us high assurance that flip-flopping would be a rewarding strategy.

[Table 4 about here]

If switching to a popular position makes sense, switching to an unpopular one does not. Table 4 shows that candidates who advocated higher taxes two years ago and remained steadfast this year would receive 72 percent of the vote. Switching positions would reduce the candidate's support to 35 percent, resulting in a 37-point loss. Thus, flip-flopping can be beneficial or harmful, depending on whether the policy reversal moves the candidate nearer or farer from voters' preferences.

Table 4 summarizes the effects of two other changes that shed additional light on the effects of inconsistency. The right column compares the popularity of candidates who are now calling for higher taxes. Candidates who consistently articulated that position get 72 percent of the vote. Those who historically called for tax cuts, on the other hand, receive only 51 percent of the vote. Two mechanisms contributed to the 21-point decline: voters reacted negatively to inconsistency per se, and they felt less sure that the inconsistent candidate would actually take the popular step of increasing taxes on the rich.

Finally, the left column of Table 4 compares candidates who now propose to decrease taxes. Those who consistently advocated tax cuts could expect 42 percent of the vote, whereas those who previously called for higher taxes could get 35 percent of the electorate. In this scenario, the mechanisms in Hypotheses 1 and 2 work at cross purposes. Voters penalize the inconsistent candidate for being a flip-flopper, but they take some comfort in the fact that the inconsistent candidate might not carry out his current (and highly unpopular) promise to decrease taxes on the rich. The negative mechanism dominates the positive one, but the loss of support is only 7 percentage points.

4.4 Distinguishing the Effects of Valence and Proximity Effects of Changing Course

As noted above, inconsistency not only evokes a negative valence reaction against flip-flopping, but also—in expectation—moves candidates closer to or farther from the ideal points of voters. Previous research shows that most voters prefer candidates who take positions closer to their own (Tomz and Van Houweling 2008). In this section, we use subgroup analysis to more effectively isolate the valence and proximity implications of changing course.

Specifically, we divide the sample of voters into three groups: those who think the government should increase taxes on the wealthy, those who want tax rates to remain the same, and those who think taxes on the wealthy should decrease. Within each group of voters, we note the actual performance of two candidates—one consistent and the other inconsistent—when each faces a common opponent. We then compare the actual performance with the proximity benchmark: how voters would have behaved if they experienced no valence reaction against flip-flopping, but instead relied exclusively on the proximity decision rule of choosing the candidate closest to their own positions. For the purpose of assessing proximity, we assumed that candidates stood at the locations in Figure 2. If voters had a negative valence reaction, then espousing an inconsistent position instead of a consistent one should have entailed higher costs than implied by pure proximity theory.

We present this type of analysis in Figures 3. The vertical axis in the figure measures the percent of the total vote, and the horizontal axis at the bottom of the figures define the issue space, ranging from increase taxes (the leftmost point) to decrease taxes (the rightmost point). The hollow circles in the interior of the plots represent the sizes of the three subgroups, and therefore mark the maximum percent that each subgroup could contribute to the total vote. The group that supports an increase in taxes comprises 63 percent of the electorate. The group that wants to keep taxes the same makes up 30 percent of the electorate, and the group that favors lower taxes rounds out the remaining 7 percent. The circles stand above their appropriate positions in the issue space, with “increase taxes” at one extreme, “decrease taxes” at the opposite extreme, and “keep taxes the same” in the middle.

[Figure 3 about here]

Figure 3 shows the consequences of espousing inconsistent position (I,D) instead of consistent position (I,I) when facing an opponent at position (D,D). We mark the positions of these candidates on the horizontal axis. The consistent candidate stands at the beginning of the horizontal arrow that hovers just above the axis, the inconsistent candidate stands at the tip of the arrow, and a box surrounds the common opponent. The vertical arrows show how candidates (I,I) and (I,D) fared within each group of voters. The beginning of each vertical arrow denotes the vote for the consistent candidate (I,I) in competition against (D,D), whereas the tip of each vertical arrow denotes the vote for the inconsistent candidate (I,D) in a race against the same opponent. In this case, all three vertical arrows point down, implying that each group of voters liked the inconsistent candidate less than the consistent one, such that the switch from (I,I) to (I,D) led to lower support across the board.

The figure then compares these observed losses with what one would have expected if voters had followed a pure proximity rule that included no penalty for flip flopping. The top horizontal line in Figure 3 spans the portion of the issue space that candidate (I,I) would be expected to win in a proximity contest against (D,D). The line overlaps the group of voters who want to increase taxes, but does not overlap voters who want to maintain or decrease taxes. Thus, a pure proximity rule would give candidate (I,I) the support of all voters who want higher taxes (63 percent of the electorate), while giving opponent (D,D) the remainder. The second horizontal line at the top of the figure shows the proximity-predicted support for candidate (I,D). This candidate should win not only the 63 percent of voters who want to increase taxes, but also the 30 percent who want to keep taxes the same, while losing the remaining 7 percent.

Thus, if voters focused only on proximity, a candidate who espoused inconsistent position (I,D) instead of consistent position (I,I) could gain among voters who want to keep taxes the same, while maintaining the same support from the other voter groups. Graphically, the middle arrow corresponding to the "keep same" subgroup would point up, and the other two arrows would not show a noticeable amount of upward or downward movement.

The figure reveals a strikingly different reality, however. Contrary to the predictions of pure proximity theory, all three arrows point down, implying that flip-flopping triggered losses among all three groups, including the "keep same" subgroup where the candidate was expected to register major gains among proximity voters. As summarized in the title of the figure, pure proximity theory predicted that a candidate could gain 30 points by embracing the inconsistent position (I,D) instead of the consistent position (I,I). In fact the reverse occurred: the candidate lost 34 points by disappointing all three segments of the electorate.

We now consider three other examples that help distinguish proximity effects from valence effects. Figure 4 shows the consequences of espousing inconsistent position (D,I) instead of consistent position (D,D) when facing an opponent at (I,I). Pure proximity theory predicts that the switch from (D,D) to (D,I) would not affect the distribution of votes: the candidate would continue to enjoy the support of voters who want to maintain or decrease taxes, while support from voters who want to increase taxes would remain elusive.

[Figure 4 about here]

In our experiment, though, the switch from (D,D) to (D,I) led to losses among voters who to maintain or decrease taxes. These losses were partially offset by unexpected gains among voters who preferred to increase taxes. However, on net, consistent candidates at (D,D) outperformed inconsistent candidates at (D,I) by 8 percentage points, even though both should have done equally well on proximity grounds alone.

Figure 5 explores a different situation, in which even proximity theory counsels against changing course. The figure shows the consequences of espousing inconsistent position (D,I) instead of consistent position (I,I) in an election against an opponent at (I,I). Under proximity theory, the contest between (I,I) and a clone candidate who also espoused (I,I) would result in a tie. We convey this with the dashed line running across the top of the figure, which signals that each candidate would receive support from half the members of each subgroup of voters. Proximity

theory implies a different outcome for candidate (D,I); he would receive no support from any voters who want higher taxes, but would garner the votes of all people who want to maintain or decrease taxes. Graphically, then, proximity theory predicts that the first vertical arrow should point down, whereas the second and third vertical arrows should point up.

[Figure 5 about here]

Once again, though, the actual costs of inconsistency exceed what one would have expected if voters had cared only about proximity. An actual contest between two anonymous candidates with identical histories—both having espoused (I,I)—would result in a tie. We did not actually administer this scenario, but for the purpose of analysis we assumed that each candidate would have received half of the vote from each subgroup of the electorate. Thus, the three vertical arrows in Figure 5 begin at 63/2 = 31.5 percent, 30/2 = 15 percent, and 7/2 = 3.5 percent. From those starting points, all three arrows point down, showing that the inconsistent candidate lost support not only where expected (among those who want to increase taxes), but also where proximity theory predicted strong gains (among those who wanted to maintain or decrease taxes). Overall, proximity theory predicted a loss of 13 points, but the inconsistent candidate actually suffered a 34-point electoral blow. This excess loss again confirms that voters are not only thinking about proximity, but also reacting negatively to flip flopping.

Finally, we consider a case in which proximity theory would argue in favor of changing course. Figure 6 shows the consequences of espousing inconsistent position (D,I) instead of consistent position (D,D) when facing an opponent at (D,D). Pure proximity theory predicts that the switch from (D,D) to (D,I) would lead to losses among voters who want to decrease taxes, but result in gains among the other electoral subgroups. Graphically, the first and second arrows should point up, while the third arrow should point down. On balance, such a move would be electorally advantageous: adopting the inconsistent position should help the candidate capture the full support of the two largest groups of voters, while losing ground among the minor faction who would like to see lower taxes on the rich. Support should leap from 50 to 93 percent, for a net gain of 43 points.

[Figure 6 about here]

The actual behavior of voters only partially coincided with proximity theory, however. By switching from (D,D) to (D,I), the candidate did indeed gain among voters who wanted higher taxes, but support in this subgroup did not rise to the unanimity implied by pure proximity theory. Instead, a substantial share of voters who wanted higher taxes actually preferred a consistent candidate who repeatedly called for lower taxes, instead of embracing an inconsistent candidate. The switch from (D,D) to (D,I) also failed to bring the expected gains among voters who wanted to keep taxes the same. In fact, as the downward arrow indicates, the inconsistent candidate lost ground in this intermediate group. Finally, the switch from (D,D) to (D,I) resulted in the anticipated evaporation of support among voters who want to decrease taxes. Overall, changing course helped the candidate gain 15 points—enough to win the election, but only a fraction of what the candidate would have gained if voters had not reacted negatively to flip flopping.

Overall, the preceding figures illustrate the importance of accounting for both spatial and valence consequences in specific competitive contexts when assessing the electoral implications of inconsistency. All four figures (and twelve others not presented for reasons of space) confirm the existence of valence losses, while also showing that flip-flopping can sometimes be a profitable strategy.

5. Conclusion

To our knowledge, this paper represents the first experimental analysis of the effects of candidate inconsistency. Our research confirms three hypotheses. First, voters dislike inconsistency in itself, and therefore punish candidates for changing positions. Second, voters discount the current policy pronouncements of candidates who have taken different positions in the past. Finally, despite the limitations implied by the other two hypotheses, inconsistency is sometimes electorally beneficial, because the gains associated with moving toward popular positions can outweigh the costs of being perceived as a flip-flopper.

In future research, we plan to study how responses to inconsistency vary with the individual attributes of voters. Voters who are very certain of their positions or who have themselves held consistent positions over time, for example, may be less tolerant of flip-flopping. We also plan to study how the attributes of candidates and the context of elections affect the consequences of inconsistency. Candidates with party affiliations might evoke affective judgments from voters with partisan attachments, but such candidates also have partisan policy reputations that might condition responses to their current and past policy positions. Moreover, candidates may influence the reactions of voters by attempting to excuse their own inconsistency (e.g. explaining to voters that circumstances have changed) or by criticizing the inconsistency of their opponents (e.g. claiming their opponent is not really committed to the new position). This paper offers initial answers to several fundamental questions about the politics of inconsistency, while also providing an experimental template for future research about a topic at the heart of democratic politics.

References

Adams, Greg D. 1997. “Abortion: Evidence of an Issue Evolution.” *American Journal of Political Science.* Vol. 41 No.3

Bernhardt, Michael D. and Daniel E. Ingberman. 1985. “Candidate Reputations and the Incumbency Effect.” *Journal of Public Economics* 27.47-67.

Carmines, Edward G. and James A. Stimson. 1989. *Issue Evolution: Race and the Transformation of American Politics.* Princeton: Princeton University Press.

Downs, Anthony. 1957. *An Economic Theory of Democracy.* New York: Harper Collins.

Karol, David. 2009. *Party Position Change in American Politics: Coalition Management.* Forthcoming: Cambridge University Press.

Poole, Keith T. and Howard Rosenthal. 1997. *Congress: a Political-Economic History of Roll Call Voting.* New York: Oxford University Press.

Tomz, Michael. “Domestic Audience Costs in International Relations: An Experimental Approach.” *International Organization* 61, no. 4 (Fall): 821-40.

Tomz, Michael and Robert P. Van Houweling. 2008. “Candidate Positioning and Voter Choice.” *American Political Science Review* 102, no. 3 (August): 303-18.

Table 1: Experimental Treatments

Pairing of Candidates	Number of Consistent Candidates
(I,I) versus (D,D)	2
(I,I) versus (D,I)	1
(I,I) versus (I,D)	1
(D,D) versus (I,D)	1
(D,D) versus (D,I)	1
(D,I) versus (I,D)	0

Note: Each candidate's history of positions is denoted by two letters, enclosed by parentheses. The first letter indicates the position the candidate took two years ago, whereas the second letter indicates the position the candidate took this year. I = increase taxes on wealthy Americans; D = decrease taxes on wealthy Americans.

Table 2: Estimated Effect of Inconsistency

	Vote Share (%)	95% C.I.
Inconsistent	43	[41 to 46]
Consistent	57	[55 to 59]
Difference	-14	[-17 to -10]

Note: The first row gives the average level of support for all inconsistent candidates in competition against all possible opponents, whether consistent or inconsistent. The second row gives the average support for all consistent candidates in competition against all possible opponents. The difference of 14 points represents the average cost of inconsistency.

Table 3: Voter Preferences regarding Tax Policy

Policy Preference	Percent	N
Increase taxes on wealthy Americans	63.4	523
Keep taxes on wealthy at their current level	29.4	243
Decrease taxes on wealthy Americans	7.2	59

Note: The table gives the percentage of voters who favored each of the three policy alternatives. Total sample size was 825.

Table 4: Estimated Effects of Changes in Position

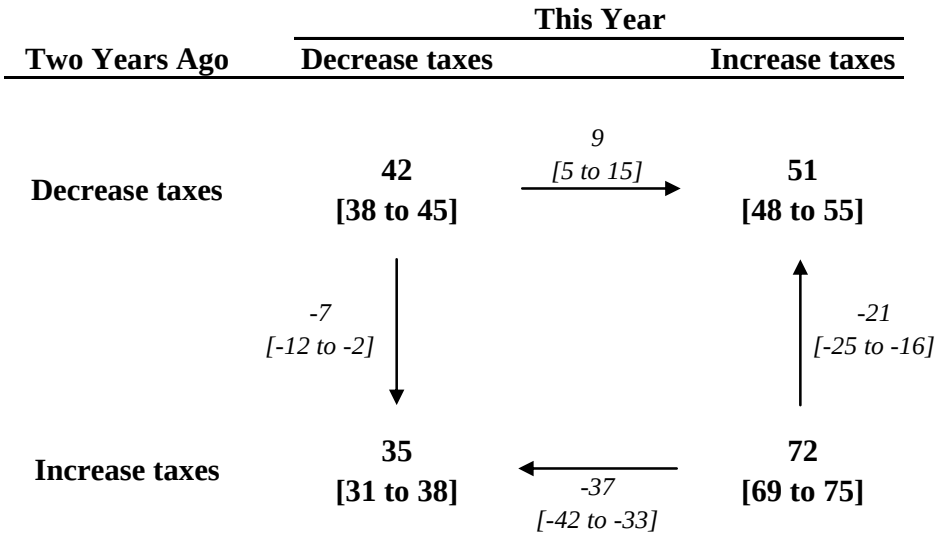

Note: The bold numbers represent the average support that each type of candidate received in competition against all possible opponents. The arrows denote changes from consistent to inconsistent positions, and the italicized numbers give the estimated effect of each change. 95 percent confidence intervals appear in square brackets.

Figure 1: Measuring Preferences and Expectations

(a) Preferences about Policy Options

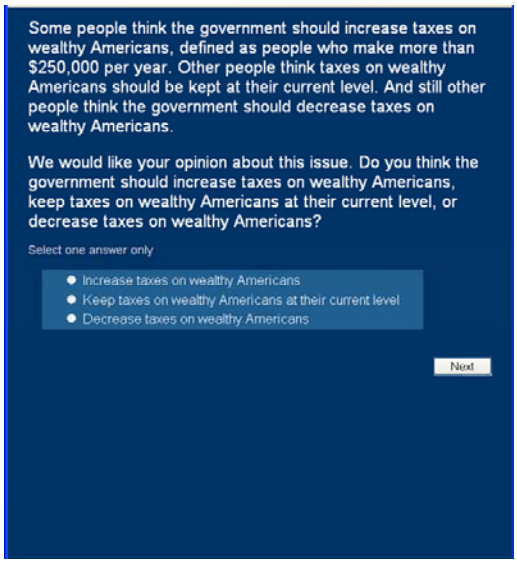

(b) Preferences about Candidates

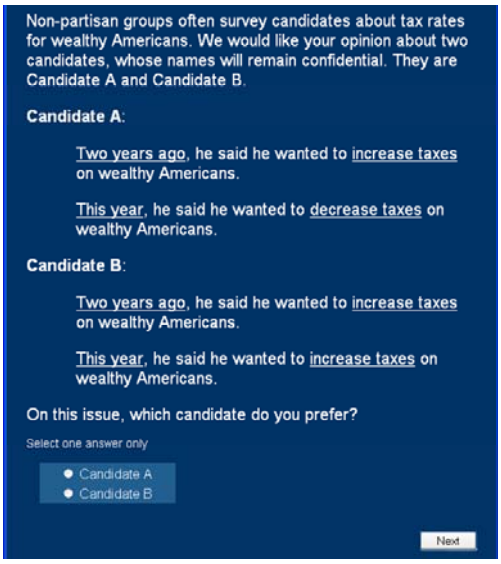

(c) Expectations about Candidates

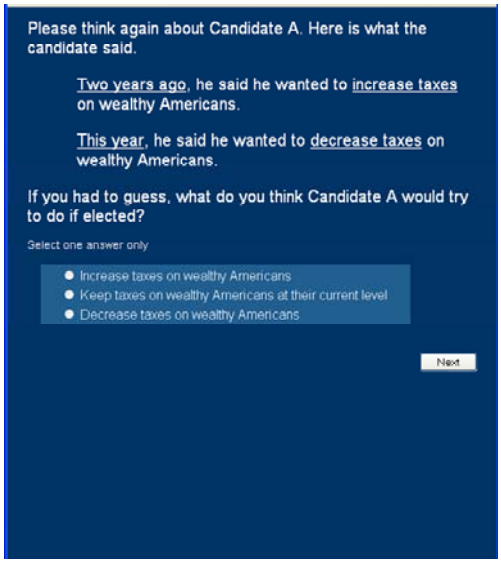

Figure 2: Expected Behavior of Four Types of Candidates

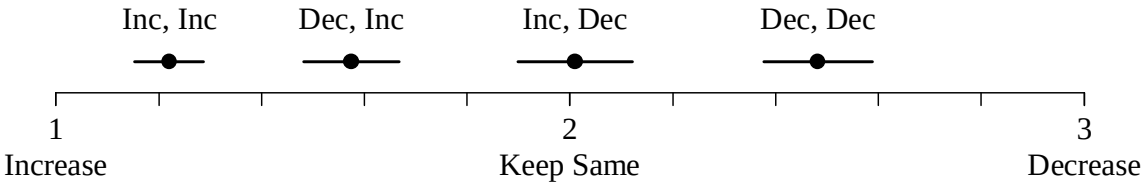

Note: The dots indicate average expectations about how each type of candidate would behave on a scale from 1 (increase taxes on wealthy Americans) to 3 (decrease taxes on wealthy Americans). The horizontal lines are 95 percent confidence intervals. In our sample, 209 respondents reported their expectations about (Inc,Inc); 214 reported expectations about (Dec,Inc); 205 reported expectations about (Inc,Dec); and the remaining 197 reported expectations about (Dec,Dec).

Figure 3: Proximity theory predicts that the inconsistent candidate will gain 30 points, but the candidate loses 34 points.

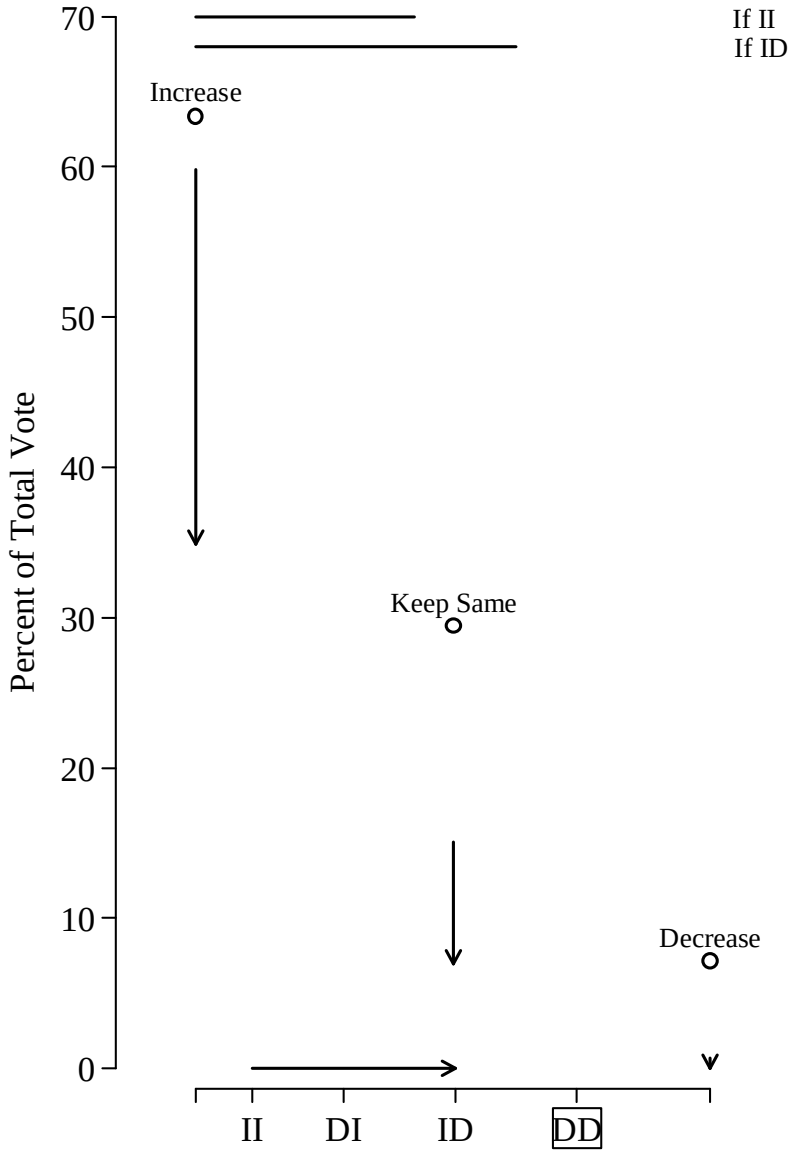

Note: The figure shows the consequences of espousing inconsistent position (I,D) instead of consistent position (I,I), when facing an opponent at position (D,D). Pure proximity theory predicts that, by switching from (I,I) to (I,D), a candidate could gain the support of all voters who want to keep taxes the same, without altering his level of support among the other electoral groups. In our experiment, though, the switch from (I,I) to (I,D) led to losses in all three voter groups, as shown by the three downward-pointing arrows. Overall, pure proximity theory predicted that a candidate could gain 30 points by espousing (I,D) instead of (I,I). But in reality the candidate who made that move lost 34 points because voters punished the candidate for being inconsistent.

Figure 4: Proximity theory predicts no change in support for the inconsistent candidate, but the candidate loses 8 points.

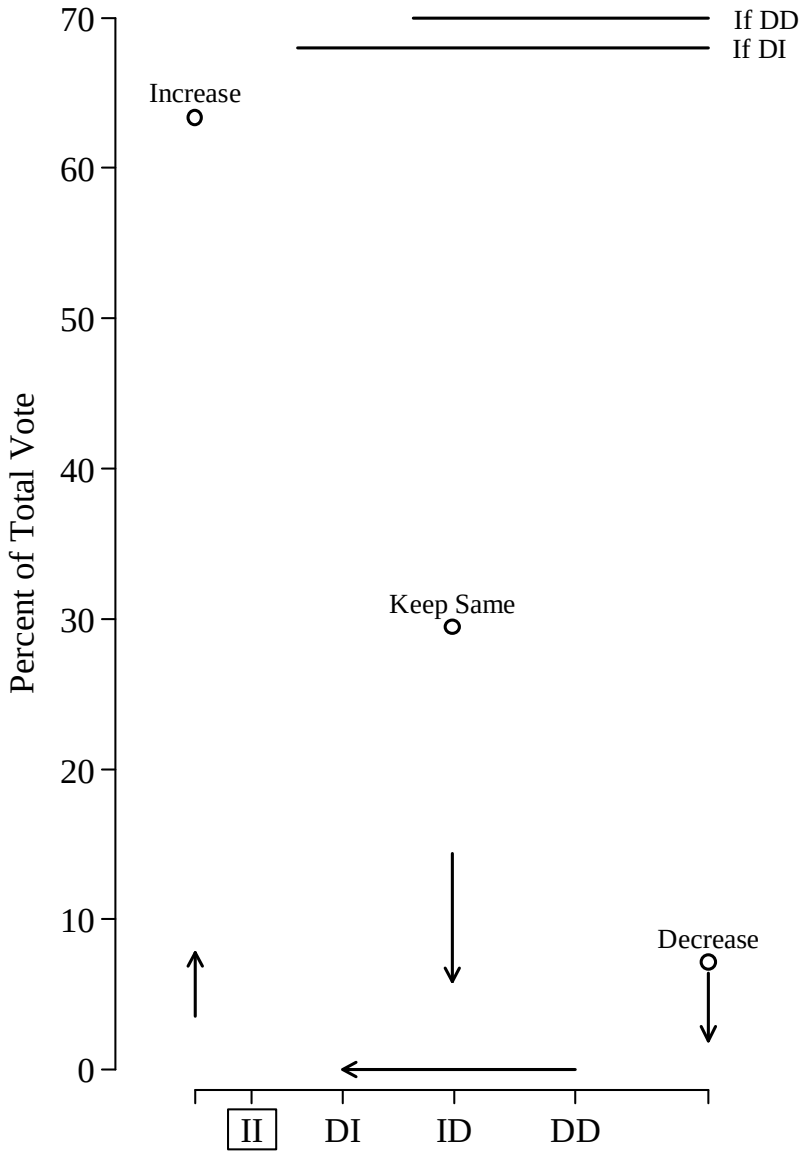

Note: The figure shows the consequences of espousing inconsistent position (D,I) instead of consistent position (D,D), when facing an opponent at position (I,I). Pure proximity theory predicts that the switch from (D,D) to (D,I) would have no effect on the outcome of the election. In our experiment, though, the switch from (D,D) to (D,I) led to losses among voters who wanted to keep taxes the same or decrease taxes, which outweighed the small gain among voters who wanted to increase taxes. Overall, pure proximity theory predicted that a candidate could neither gain nor lose by espousing (D,I) instead of (D,D). In reality, though, the candidate who made that move lost 8 points because voters punished the candidate for being inconsistent.

Figure 5: Proximity theory predicts that the inconsistent candidate will lose 13 points, but the candidate loses 34 points.

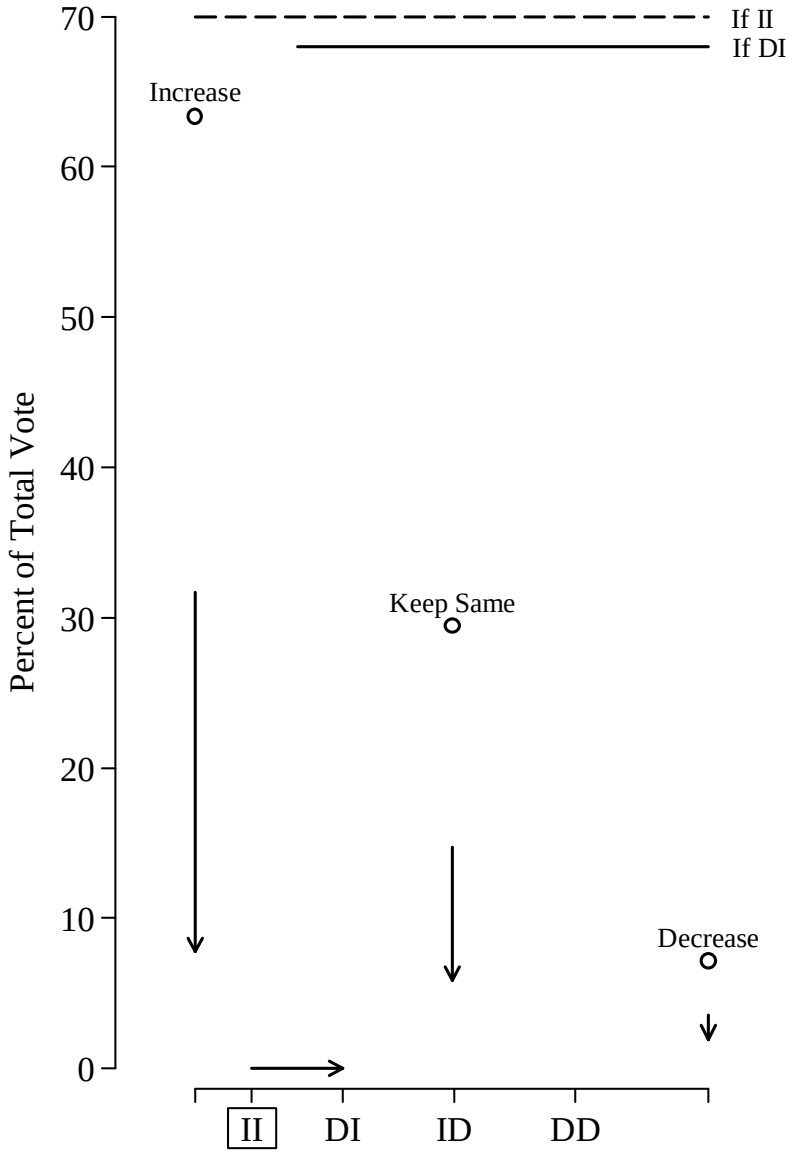

Note: The figure shows the consequences of espousing inconsistent position (D,I) instead of consistent position (I,I), when facing an opponent at position (I,I). Pure proximity theory predicts that the switch from (I,I) to (D,I) would lead to losses among voters who want to increase taxes, but lead to gains in the other electoral subgroups. In our experiment, though, the switch from (I,I) to (D,I) led to losses among all three groups of voters. Overall, pure proximity theory predicted that a candidate would lose 13 points by espousing (D,I) instead of (I,I). In reality, though, the candidate who made that move lost 34 points because voters punished the candidate for being inconsistent.

Figure 6: Proximity theory predicts that the inconsistent candidate will gain 43 points, but the candidate gains only 15 points.

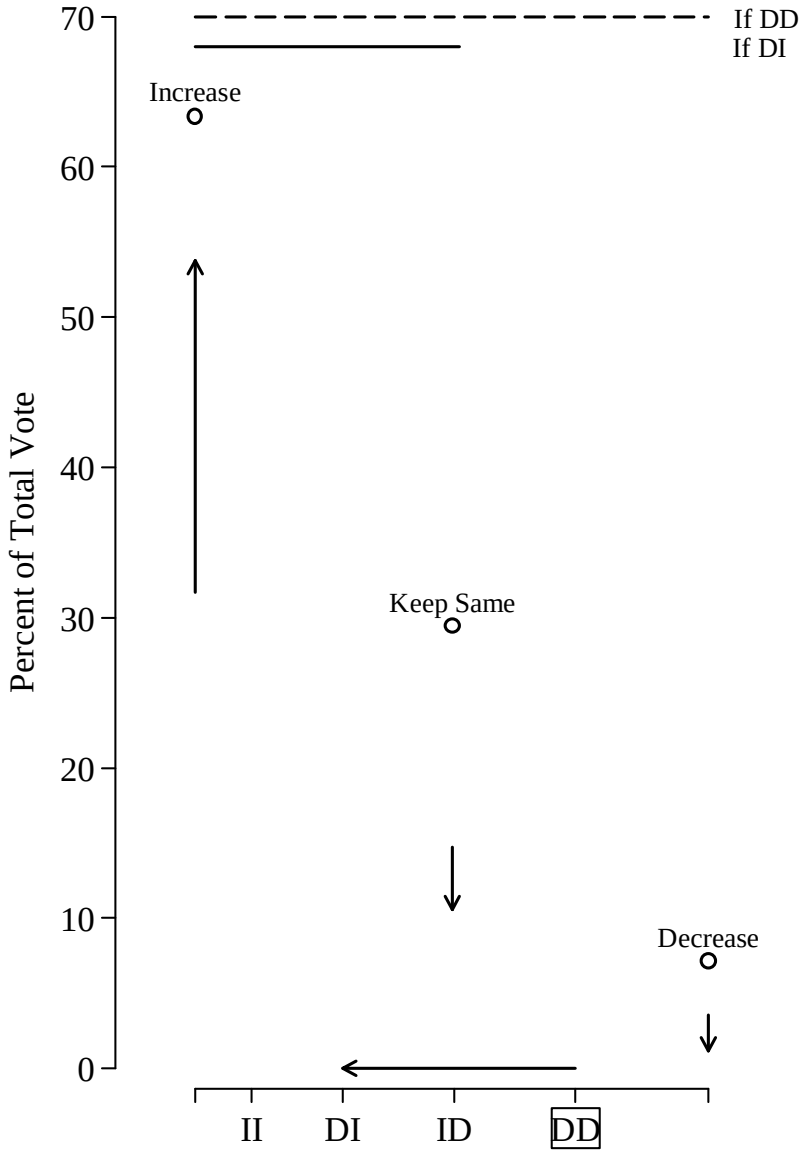

Note: The figure shows the consequences of espousing inconsistent position (D,I) instead of consistent position (D,D), when facing an opponent at position (D,D). Pure proximity theory predicts that the switch from (D,D) to (D,I) would lead to losses among voters who want to decrease taxes, but lead to gains in the other electoral subgroups. In our experiment, the switch from (D,D) to (D,I) led to gains among those who want to increase taxes, but resulted in losses among the other groups. Overall, pure proximity theory predicted that a candidate would gain 43 points by espousing (D,I) instead of (D,D). In reality, though, the candidate who made that move gained only 15 points. The candidate did not realize the full gain of 43 points because voters punished the candidate for being inconsistent.