Trump is not a (Condorcet) loser! Social choice analysis of the 2016 Republican presidential nomination

Jonathan Woon[*], Sean Craig[†], Amanda Leifson[‡], Matthew Tarpey[§]

September 15, 2018

Abstract

Many commentators argued that if elites and voters had coordinated on an alternative candidate, Donald Trump could have been defeated for the 2016 Republican presidential nomination. Conventional pre-election polls, however, do not provide enough information about voters' preferences to assess the plausibility of this claim. Relying on novel data to construct individuals' complete preferences over the set of leading Republican candidates, we find that no other candidate strictly defeats Trump in head-to-head majority rule comparisons and, far from being a Condorcet loser, that Trump is a member of the majority rule core. Our results call into question the plausibility of the coordination narrative, as Trump's support was wider than initially believed. His support came from a broad swath of the Republican primary electorate rather than from a small but intense minority. (WORD COUNT: 3647)

Keywords: Preferences, collective choice, majority rule core

[*]Professor, University of Pittsburgh, woon@pitt.edu
[†]PhD Candidate, University of Pittsburgh, sean.craig@pitt.edu
[‡]PhD Candidate, University of Pittsburgh, aml185@pitt.edu
[§]PhD Candidate, University of Pittsburgh, mmt52@pitt.edu

When Scott Walker withdrew from the Republican presidential race in September 2015, he called for unity, urging other candidates to withdraw and support a single contender who could defeat Donald Trump.[1] If #NeverTrump Republicans split their votes, Walker worried, then Trump might win despite opposition from a majority of the party's voters. The remaining candidates ignored the warning as Trump continued to earn narrow victories against a crowded field. Many commentators later endorsed Walker's logic, attributing Trump's primary success to the Republican establishment's failure to coordinate on a frontrunner.[2] This coordination narrative, however, rests on an implicit assumption about the underlying structure of Republican primary voters' preferences: that Trump could have been defeated in head-to-head, or *pairwise*, contests.[3] How accurate was this assumption about Republican primary voters' preferences in 2016?

We take an empirical social choice approach to answering this question, relying on an original dataset of Republican voters' preferences to examine the outcomes of pairwise majority rule contests. In contrast to conventional pre-election polls that ask respondents only about their likely vote, we measure voters' binary preference for pairs of leading candidates.[4] With these data, we investigate whether the structure of voters' preferences satisfied the conditions necessary for the coordination narrative to have been plausible. In addition, we characterize the extent of Trump's support (or lack thereof) in social choice terms: specifically, whether Trump was a Condorcet winner, a Condorcet loser, or a member of the

[1] Jim Newell. "Why Scott Walker Dropped Out" *Slate* (September 21, 2015), `http://www.slate.com/articles/news_and_politics/politics/2015/09/scott_walker_drops_out_why_the_wisconsin_governor_left_the_gop_primary.html`

[2] Julia Azari, "A For Effort? Republican Elites Tried to Coordinate But Never Quite Got There." *Vox* (May 19, 2016), `https://www.vox.com/mischiefs-of-faction/2016/5/19/1171261/republican-elites-coordination`; Alexander Burns and Jonathan Martin, "Republican Leaders Map a Strategy to Derail Donald Trump" *The New York Times* (March 19, 2016); Eli Stokols, "GOP Elites Line Up behind Ted Cruz" *Politico* (March 23, 2016), `https://www.politico.com/story/2016/03/ted-cruz-republican-establishment-elites-221174`.

[3] Maskin and Sen analyze hypothetically how Trump could win primaries while losing individual match-ups against each of his opponents ("How majority rule might have stopped Donald Trump" *New York Times* (May 1, 2016), `https://www.nytimes.com/2016/05/01/opinion/sunday/how-majority-rule-might-have-stopped-donald-trump.html`.

[4] A handful of existing studies examine the structure of voters' preferences empirically (e.g., Feld and Grofman 1988, Niemi and Wright 1987). Notable examples examining pairwise preferences are Brady and Ansolabehere (1989) and Radcliff (1993).

majority rule core.

Our results suggest that elite coordination was unlikely to be successful. Trump either wins or ties against every other candidate in pairwise contests, which means that not only is Trump not a Condorcet loser but no other candidate had enough support to defeat him. While Trump's support was not so broad to make him a Condorcet winner, we do find that it was enough for him to earn a place in the majority rule core. Thus, in some sense, we can say that Trump was a "rational choice" of the Republican primary electorate.

Winners, Losers, and the Core

Plurality rule is a straightforward way to make a collective choice over many candidates. Simply count the number of voters, and the candidate with the most votes wins. While easy to implement, with more than two candidates, this process may not produce the "best" social choice because a majority of voters might oppose the plurality winner. As the number of candidates increases, such opposition becomes increasingly likely (Wright and Riker 1989). However, to determine whether this is indeed a problem, sometimes referred to as Borda's paradox (Nurmi 1983), we must know how each voter ranks every possible candidate.[5] When each voter's ordering is known, the set of orders constitute a *preference profile*.

Two hypothetical preference profiles are presented in Table 1. In both profiles, Trump is the plurality winner with 40% of the vote, while Cruz and Rubio split the remainder of the vote with 30% each. However, the two profiles differ starkly in how Trump is ranked relative to the others. In Profile A, Cruz voters (Group 2) and Rubio voters (Group 3) rank Trump last. This means that either Cruz or Rubio would defeat Trump with a decisive majority (by a 60% to 40% margin) in a pairwise contest. Since Trump is defeated by all other candidates in pairwise majority rule contests (a decision procedure advocated by the Marquis de Condorcet), we would say that he is a *Condorcet loser*. This situation illustrates the Borda paradox, in which plurality rule can select a Condorcet loser, and it is a context in

[5]We focus on complete preference orders here for ease of exposition. It is sufficient to know how each voter compares each pair of candidates, which does not necessarily produce an ordering.

Table 1: Hypothetical Preference Profiles

Profile A: Trump is a Condorcet loser

	Group 1	Group 2	Group 3
1st choice	Trump	Cruz	Rubio
2nd choice	Cruz	Rubio	Cruz
3rd choice	Rubio	Trump	Trump
Share of voters	40%	30%	30%

Profile B: Trump is a Condorcet winner

	Group 1	Group 2	Group 3
1st choice	Trump	Cruz	Rubio
2nd choice	Cruz	Trump	Trump
3rd choice	Rubio	Rubio	Cruz
Share of voters	40%	30%	30%

Profile C: No Condorcet winner or loser

	Group 1	Group 2	Group 3
1st choice	Trump	Cruz	Rubio
2nd choice	Cruz, Rubio	Trump, Rubio	Trump, Cruz
Share of voters	40%	30%	30%

which coordination by Republican candidates could have prevented Trump's nomination.[6]

By contrast, in Profile B, Cruz and Rubio voters place Trump in the middle of their orderings. In this second scenario, Trump would win pairwise contests against both Cruz and Rubio (70% to 30%). Thus, Trump is majority-preferred to every other candidate and is, instead, a *Condorcet winner*. In this case, Trump has enough support that neither coordination nor strategic voting would have prevented his nomination. Importantly, these simple examples call attention to the necessity of knowing, empirically, how voters rank candidates beyond their top choice.

In practice, many voters may lack a *strict* preference for one candidate over another and may instead be *indifferent* between them.[7] That is, these voters would view two candidates as

[6]This profile is also a situation in which the Condorcet loser could be defeated by strategic voting (Forsythe et al. 1993).

[7]A potential source of indifference may be uncertainty in the voter's preferences, as in Potthoff and Munger (2005), who examine the effects of uncertainty about candidates' positions in a one-dimensional

equally good (or equally bad). While admitting this possibility allows us to consider a richer and more empirically plausible set of preference profiles, it leads to some complications (Jones et al. 1995). Consider Profile C at the bottom of Table 1 in which every voter is indifferent between the candidates ranked below their top choice. In each pairwise comparison, neither candidate has the support of a majority of voters. For example, in the head-to-head contests between Trump and Cruz, a 40% plurality supports Trump, a 30% minority supports Cruz, and the remaining 30% (Rubio supporters) are indifferent.

If these indifferent voters were forced to choose between Trump and Cruz, voting for either candidate would be, strictly speaking, rational. How should we determine the majority preference? In our analysis, we consider a candidate to be the majority preference if and only if a majority of voters strictly prefer that candidate. This definition of majority preference (following Austen-Smith and Banks, p. 27) is the most conservative in the sense that it involves the fewest assumptions about how indifferent voters would behave in practice.[8] If neither candidate garners an absolute majority, then we consider voters to be indifferent in the aggregate.

From a social choice perspective, Profile C differs subtly, but importantly, from Profile A. Collectively, voters are indifferent between Trump and Cruz, between Trump and Rubio, and between Cruz and Rubio. No candidate strictly defeats any other candidate, and so there is no Condorcet winner. But neither is there a Condorcet loser because no candidate strictly loses to any other candidate. Nevertheless, social choice theory provides a definition of what it means to be "best" according to preferences for the entire group: an alternative is a member of a set called the *core* if it is "at least as good as" (i.e., weakly better than) every

spatial model on whether the preference profile of the entire electorate satisfies the single-peaked property. In our analysis, we do not impose any spatial or single-peaked structure on voters' preferences, nor do we take a stance regarding the source of preferences. We posit only that preferences exist and impose as little structure on those preferences as possible.

[8] Formally, the *majority rule* preference is xPy if and only if $|P(x, y; \rho)| > n/2$, where x needs to be strictly preferred by a majority. In contrast, *plurality rule* (as a binary preference aggregation rule) defines the group preference to be xPy if and only if $P(x, y; \rho) > P(y, x; \rho)$, requiring that more voters strictly prefer x to y than y to x. Whereas plurality rule ignores indifferent voters altogether (as if they abstain), we use majority rule because it treats indifferent voters as if they were, in practice, "swing voters" who could be persuaded (and therefore might still vote for either candidate).

other alternative according to some preference aggregation rule. In Profile C, all candidates are equally good according to pairwise majority rule, so all are members of the core and are "best" (albeit in a weak sense), whereas in Profile B, Trump is the unique member of the core, and in Profile A, Trump is not in the core (instead, Cruz is the unique member). If we find that Trump is, indeed, in the core then his nomination would reflect the rationality of Republican voters' group choice.

Data

Our analysis relies on two independent sources of data. First, we conducted an original survey to elicit complete preferences over the six leading candidates in early 2016 (Jeb Bush, Ben Carson, Chris Christie, Ted Cruz, and Donald Trump) from likely Republican voters ($N = 755$) through Amazon's Mechanical Turk platform (Berinsky, Huber and Lenz 2012).[9] We fielded our surveys between January 21 and February 8, 2016, a period corresponding to the week before the Iowa Caucus until the eve of the New Hampshire Primary. Since our MTurk sample may not be representative of the Republican primary electorate, we also conduct parallel analysis using the 2016 ANES Pilot Survey (Republican-leaning voters' preferences over the same candidates except Christie, $N = 440$). These data were generated concurrent with our MTurk data, from January 22 to 28, 2016.

In both datasets, the distribution of respondents' first choice candidates were comparable to publicly available commercial polls fielded during the same period.[10] Remarkably, the proportion of Trump supporters in our MTurk sample is quite close to both the percentage of Trump supporters in the ANES Pilot sample and to the polling average. The percentage of MTurk respondents who supported each of the other candidates was also well within

[9]Identifying primary voters is not entirely straightforward, as there is no consensus among commercial and academic polls regarding how to identify respondents' likelihood of voting in a primary election. Out of 26 news polls conducted concurrently with our MTurk surveys, 60% included all Republican-identifying registered voters in their samples, whereas 35% further restricted the sample to self-reported likely primary voters. We used self-reported likelihood of voting in the Republican primary as our screening criterion.

[10]Table A1 in the appendix compares our MTurk and ANES Pilot samples with polls between January 16 and February 15. We obtained the data for the comparison from HuffPost Pollster, `http://elections.huffingtonpost.com/pollster/2016-national-gop-primary`.

the range of contemporary polls, with the exception of greater support for Rand Paul.[11] These comparisons suggest that we can be reasonably confident in the validity of our MTurk sample.

Results

The goal of our analysis is to understand the structure of Republican voters' preferences for or against Donald Trump relative to other candidates. Was he a Condorcet winner? Condorcet loser? A member of the core? To answer these questions, we aggregate the preferences for Trump against the field of candidates elicited from the MTurk and ANES pilot samples in two ways. First, we compare the proportion of voters who strictly prefer Trump to each of the other candidates. Using this criterion, we say that Trump is strictly preferred by majority rule (the majority winner) if the percentage of voters who strictly prefer Trump exceeds 0.5. If the proportion strictly favoring the other candidate is greater than 0.5, then that other candidate is strictly preferred (Trump is the majority loser). If neither candidate is strictly preferred by a majority, then we consider the aggregate preference to be one of indifference (a majority tie). Second, we account for uncertainty by using the 95% confidence interval instead of the point estimate. Under this interval criterion, we can say that Trump is the majority winner if the entire confidence interval for the percentage of respondents supporting Trump exceeds 0.5. Similarly, if the entire interval falls below 0.5, then Trump is the majority loser, and if the 0.5 is contained in the interval, then there is a majority tie.

Figure 1 presents Trump's level of support compared to the support for each of the other candidates for our two data sets. Across both samples, Trump's share of supporters is higher than that of every other candidate, although he does not always have a majority. In our MTurk sample, shown in the left panel, Trump is strictly preferred to both Jeb Bush (0.54 to 0.38) and Chris Christie (0.55 to 0.36) according to the point estimate criterion.

[11] As is typical, our MTurk sample is somewhat less representative of the population in demographic terms. MTurk respondents tend to be younger, poorer, more educated, and more likely to be female. See Appendix Table A2 for details.

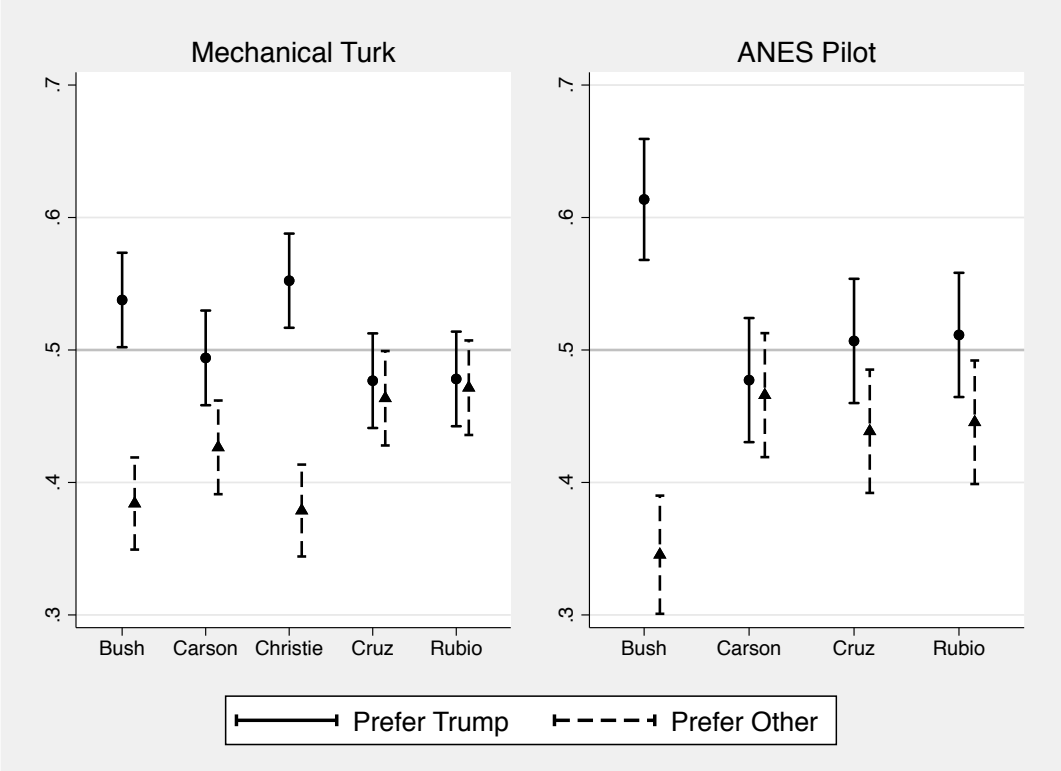

Figure 1: Pairwise preferences for and against Donald Trump

While he has a larger share of supporters than Ben Carson (0.49 to 0.42), Ted Cruz (0.48 to 0.46), or Marco Rubio (0.48 to 0.47), he does not have a majority of strict supporters against any of these other candidates. Applying the interval criterion to our MTurk sample yields identical conclusions regarding aggregate preferences. Since Trump defeats two of the five other candidates and is tied with the other three, he is neither a Condorcet winner nor Condorcet loser, but he is in the core and therefore a rational Republican group choice.

The data from the ANES Pilot sample, shown in the right panel, also support the conclusion that while Trump is neither a Condorcet winner nor loser, he is in the core.[12] Looking at the point estimates, Trump is majority preferred to Bush (0.61 to 0.35), Cruz (0.51 to 0.44), and Rubio (0.51 to 0.45), but is majority tied with Carson (0.48 to 0.47).

[12]We obtain similar results if we use plurality rule (ignoring indifferent voters) instead of majority rule. If anything, using plurality rule suggests stronger support for Trump in that he emerges as a Condorcet winner in both samples using the point criterion. When we factor in uncertainty, however, Trump is in the core but not a Condorcet winner using the interval criterion.

When we examine the confidence intervals, which are wider for the ANES sample than for the MTurk sample, we see that Trump is majority preferred only to Bush. For the remainder of the candidates, we cannot be sufficiently confident that his support exceeds 0.5. Nevertheless, we find that no other candidate majority defeats Trump, so Trump is in the core.[13]

We find consistent support for our conclusion that Trump is in the majority rule core even when we examine subsamples of our MTurk data to check the robustness of our results.[14] This analysis, summarized in Table 2, also suggests that the results in Figure 1 may underestimate Trump's support. In our MTurk surveys, we elicited preferences using either pairwise comparisons or ranked choice, and when we restrict the sample to respondents who gave pairwise preferences, the results are the same as for the sample as a whole. Interestingly, Trump's support appears to be greater when we look only at the subsample of respondents who described their preferences by ranking candidates: the point estimates from ranked choices suggest that Trump is a Condorcet winner, although the intervals yield the same conclusions as the pairwise comparison data.

Similar findings obtain when we account for the potential unrepresentativeness of our MTurk sample. When we drop (overrepresented) Rand Paul supporters from the sample and re-estimate support for Trump, we find stronger support for Trump over Ben Carson (Trump defeats Carson using the interval criterion), but the majority preference between

[13]It is interesting to contrast our results to those of Kurrild-Klitgaard (2017). In the time period closest to ours (January 2016), Kurrild-Klitgaard's results are similar to ours, as he does not find any statistically significant differences in head-to-head matchups between Cruz and Trump and Trump and Rubio, suggesting aggregate indifference. However, in two surveys later than ours (March 2016), his results differ substantially. In both, he finds that Cruz is strictly preferred to Trump. In the first of these, Trump's support is lower but not statistically different from Rubio's, while in the second, Trump's support is also significantly less than that of Kasich or Rubio, which makes Trump a Condorcet loser. There are key differences between Kurrild-Klitgaard's data and ours that may account for the conflicting findings. Importantly, the survey questions he uses do not allow for indifference as a response option, so it is possible that a sizeable portion of the respondents in his data are truly indifferent but nevertheless express preferences against Trump. Furthermore, the poll in which Trump is found to be a Condorcet loser (NBC/WSJ March 3-6) may be an outlier: of the 10 polls fielded between February 22 and March 11, Cruz's support in this particular poll is the highest (27% compared to an average of 22.2%) and Trump's support is the lowest (30% compared to an average of 40.4%).

[14]Our results are also corroborated by bootstrap simulations that we used as an additional robustness check. For details, see the Appendix.

Table 2: Comparing Preferences in Subsamples

Opponent	MTurk Sample	ANES Pilot Sample	MTurk Pairwise	MTurk Ranked	MTurk (No Paul)	MTurk Weighted
Bush	Trump*	Trump*	Trump*	Trump*	Trump*	Trump*
Carson	——	——	——	Trump	Trump*	Trump*
Christie	Trump*	n/a	Trump*	Trump*	Trump*	Trump*
Cruz	——	Trump	——	Trump	——	Trump
Rubio	——	Trump	——	Trump	——	Trump
Condorcet Winner	✗	✗	✗	✓	✗	✓
Condorcet loser	✗	✗	✗	✗	✗	✗
Core	✓*	✓*	✓*	✓*	✓*	✓*

Cell entry indicates the preferred candidate

—— indicates tie

* indicates satisfies interval criterion

Trump and Cruz and between Trump and Rubio remain the same (indifference). We also adjust for the skew in the demographic characteristics of our sample (that it is younger, more educated, and more female than Republican samples in election studies) by creating post-stratification weights based on gender, income, age, and education to match the sample of Republican voters in the 2016 ANES Time Series Study.[15] The weighted sample provides the greatest support for Trump's nomination. As with the ranked choice subsample, Trump is a Condorcet winner according to the point estimates, and he defeats three of the five candidates (Bush, Carson, and Christie) using the interval criterion.

Conclusion

By eliciting and examining voters' complete preferences over leading Republican candidates in 2016, our empirical social choice analysis yields a picture of Donald Trump's support that challenges the conventional coordination narrative. Had Trump been a Condorcet

[15]See the appendix for details of our weighting procedure.

loser, defeated in pairwise contests by every other candidate, elite coordination could have prevented his nomination. Yet since this was not the case, neither coordination among elites nor strategic voting in the electorate would have made much difference. To put it plainly, Trump did not win because Republicans simply split their votes. Indeed, our data suggest no other candidate could have defeated Trump in a head-to-head match-up once the primaries began. Our findings suggest instead that Trump drew support from a sufficiently broad swath of Republican voters who, in the aggregate, found him to be "at least as good as" any other candidate. This earned him a place in the majority rule core, making him a rational choice of the Republican electorate.

While our results suggest that it is unlikely that Trump's nomination could have been stopped by the time the primaries began, Trump's nomination might have been stopped in other ways. For example, Azari and Masket (2017) note that Republican Party's nomination rules were too democratic, allowing participation by too many candidates in the debate process and broad participation by voters. Stronger parties in which elites controlled the set of candidates, they argue, could have prevented Trump from receiving the Republican nomination. We agree that party elites with greater control could have nominated another candidate, but the widespread rank-and-file support found in our data suggests that blocking Trump's nomination would have required shutting him out of Republican primaries and caucuses altogether.

It is also possible that Trump's nomination could have been forestalled by elite coordination much earlier in the process, such as when Trump announced his candidacy. Because our data provide only a snapshot of voters' preferences at a particular point in time, it is possible that voters' preferences were less developed in the summer of 2015. Coordinating on a strong establishment Republican candidate and cohesive messaging by party elites might have been able to influence the formation of voters' preferences, but this is pure speculation.

References

Azari, Julia and Seth Masket. 2017. "Intraparty Democracy and the 2016 Election." *Conventional Wisdom, Parties, and Broken Barriers in the 2016 Election* p. 137.

Berinsky, Adam J, Gregory A Huber and Gabriel S Lenz. 2012. "Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk." *Political Analysis* 20(3):351–368.

Brady, Henry E and Stephen Ansolabehere. 1989. "The nature of utility functions in mass publics." *American Political Science Review* 83(1):143–163.

Feld, Scott L and Bernard Grofman. 1988. "Ideological consistency as a collective phenomenon." *American Political Science Review* 82(3):773–788.

Forsythe, Robert, Roger B Myerson, Thomas A Rietz and Robert J Weber. 1993. "An experiment on coordination in multi-candidate elections: The importance of polls and election histories." *Social Choice and Welfare* 10(3):223–247.

Jones, Bradford, Benjamin Radcliff, Charles Taber and Richard Timpone. 1995. "Condorcet winners and the paradox of voting: probability calculations for weak preference orders." *American Political Science Review* 89(1):137–144.

Kurrild-Klitgaard, Peter. 2017. "Trump, Condorcet and Borda: Voting paradoxes in the 2016 Republican presidential primaries." *European Journal of Political Economy* .

Niemi, Richard G and John R Wright. 1987. "Voting cycles and the structure of individual preferences." *Social Choice and Welfare* 4(3):173–183.

Nurmi, Hannu. 1983. "Voting procedures: a summary analysis." *British Journal of Political Science* 13(2):181–208.

Potthoff, Richard F and Michael C Munger. 2005. "Voter uncertainty can produce preferences with more than one peak, but not preference cycles: a clue to the fate of Ross Perot?" *The Journal of Politics* 67(2):429–453.

Radcliff, Benjamin. 1993. "The structure of voter preferences." *The Journal of Politics* 55(3):714–719.

Wright, Stephen G. and William H. Riker. 1989. "Plurality and runoff systems and numbers of candidates." *Public Choice* 60:155–175.

Appendix

Data and Preference Variables

We collected our MTurk data in three waves. We fielded Wave 1 ($N = 91$) on January 21, 2016, Wave 2 ($N = 476$) from January 25 to 28, 2016, and Wave 3 ($N = 188$) from February 5 to 8, 2016. For context, the second wave took place a few days prior to the Iowa Caucus and the third wave just before the New Hampshire Primary.

We implemented the survey in Qualtrics and solicited respondents using Amazon's Mechanical Turk ("MTurk"). We elicited preferences from MTurk respondents using two methods, pairwise comparisons and rank orderings. Respondents in Wave 1 and Wave 2 were randomly assigned to one of the two methods. All respondents in Wave 3 provided pairwise comparisons. In the pairwise elicitation method, respondents ($N = 469$) gave their preferences among all possible combinations of the six highest-polling Republican candidates at the time: Jeb Bush, Ben Carson, Chris Christie, Ted Cruz, and Donald Trump. Between all three waves, the top six candidates remained the same. The first one to drop out of the top six candidates was Chris Christie, who dropped out of the race after the New Hampshire primary on February 10th. Pairwise survey items followed the format in Figure A1.

Figure A1: Example of Pairwise Preference Item

Respondents in the rank ordering condition ($N = 286$) assigned each of these same six candidates a rank between 1 and 6, with one representing the most-preferred candidate and ties permitted. Figure A2 shows the survey interface. We code pairwise preferences based on this ordering. For example, if a respondent gives candidate A a rank of 2, then the respondent prefers A to candidate B if B is ranked 3 or greater. We code the respondent as indifferent if both candidates are placed in the same rank.

To code preferences from the ANES Pilot data, we use feeling thermometer scores (a method used by Niemi and Wright (1987) and Radcliff (1993)) for five of the six candidates covered in our MTurk data: Bush, Carson, Cruz, Rubio, and Trump. (ANES respondents did not rate Christie.) We generated pairwise preferences by comparing feeling thermome-

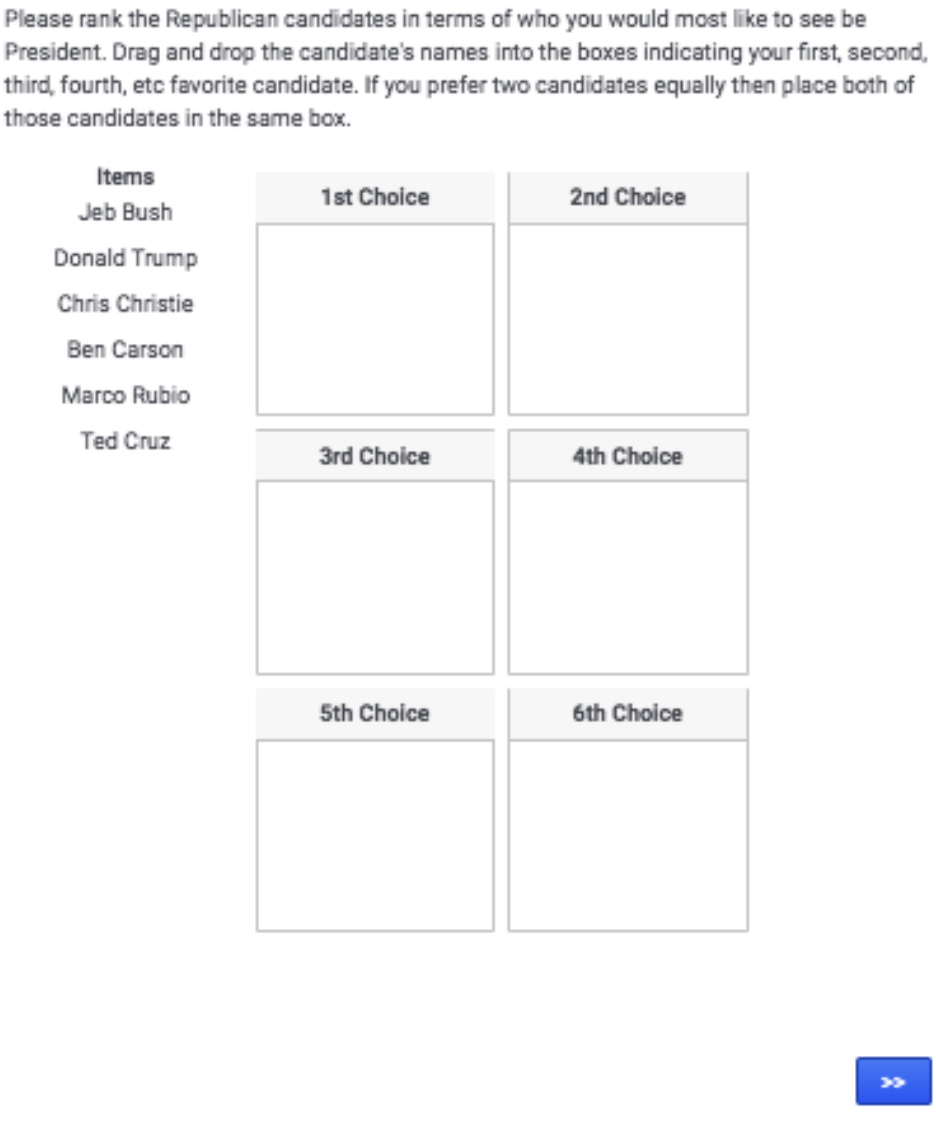

Figure A2: Rank Ordering Preference Item

ter scores for each possible combination of candidates. Equal scores indicate indifference, whereas a higher score for one candidate indicates a strict preference for that candidate.

Like most polls, perfect identification of primary voters is impossible because we cannot use voting records to validate turnout. Pollsters and scholars typically deal with this problem by relying on self-reported participation or intended participation. This is the approach we took in our MTurk survey, asking voters how likely they were to vote in the Republican primary on a five-point scale ranging from "Not very likely" to "Very likely." We screen out survey respondents who indicated that they were unlikely to vote in the Republican primary and, where applicable, those who indicated that they were likely to vote in the *Democratic* primary. The ANES Pilot data contain no items related to voting intent. Instead, we take the respondent's party identification as a proxy, including only respondents who identified with the Republican Party, including "leaners."

Sample Characteristics

As discussed in the text, Table A1 shows that the distribution of top choice candidates in our MTurk data and in the ANES Pilot sample closely match the polling averages during our survey period. Demographic characteristics of both samples are reported in Table A2. Compared to the ANES Pilot sample, our MTurk sample is much younger (70.4% under 33 in MTurk, 40.7% in ANES), more highly educated (46.7% with bachelor's or graduate degrees in MTurk, 29.1% in ANES), somewhat poorer (21.1% with incomes over $80,000 in MTurk, 34.4% in ANES), and more female (55.8% in MTurk, 50% in ANES).

Table A1: Comparison of Top Choice Preferences

Top Choice	MTurk sample	ANES Pilot sample	Average Poll	Range
Bush	6.6	4.3	5.4	(3-10)
Carson	8.5	8.6	8.0	(4-14)
Christie	3.7	3.2	3.0	(1-5)
Cruz	13.6	16.8	17.7	(11-23)
Fiorina	2.1	3.0	2.1	(1-5)
Kasich	3.1	3.0	3.8	(1-11)
Paul	9.4	4.8	2.9	(1-5)
Trump	35.3	36.4	36.7	(25-43)
Rubio	12.1	10.0	13.2	(8-21)
Other	2.9	1.4	1.7	(0-6)
Undecided	2.8		7.1	(1-15)

Poll Average: 16 January - 15 February 2016, 26 Surveys

Post-Stratification Survey Weights

To account for the unrepresentativeness of our sample, we generated post-stratification weights based on four demographics: gender, income, age, and education. Each of the variables was dichotomized in order to limit the number of categories and ensure that there

Table A2: Sample Demographics

	MTurk		ANES Pilot	
	Count	Percent	Count	Percent
Age				
18-24	107	14.2	34	7.7
25-33	424	56.2	145	33.0
45-64	204	27.0	171	38.9
65+	20	2.7	90	20.5
Education				
Less than High School	7	0.9	37	8.4
High School Diploma	93	12.3	161	36.6
Some College	205	27.2	85	19.3
Associate's Degree	97	12.9	29	6.6
Bachelor's Degree	260	34.4	79	18.0
Graduate/Professional Degree	93	12.3	49	11.1
Income				
Under $30,000	188	24.9	109	24.8
$30,000 to $50,000	199	26.4	89	20.3
$50,000 to $80,000	208	27.6	90	20.5
Over $80,000	159	21.1	151	34.4
Sex				
Male	334	44.2	220	50.0
Female	421	55.8	220	50.0

were no empty cells, for a total of 16 strata. We created these strata for our MTurk sample and for the 2016 ANES Time Series Sample. We use likely Republican primary voters in the ANES Time Series data as a proxy for the target population, as there is no other demographic information on this population of which we are aware. To create the weights, we divided the proportion in each stratum of the ANES Time Series data by the percentage we observed in the MTurk sample. Thus, the weights represent the inverse probability of being selected into our sample. When we perform pairwise comparison using the weighted ANES-weighted MTurk, our key findings remain. Moreover, Trump's performance improves substantially in both the top-choice selections and pairwise comparisons. Indeed, as Table A8 illustrates, Trump comes the closest to Condorcet victory when we apply these survey weights.

Robustness Checks

Tables A3 and A4 present the numerical results corresponding to Figure 1 in the text and are the basis for the first two columns in Table 2. The remaining tables provide the numerical estimates that support the robustness findings in Table 2: Table A5 (column 3), Table A6 (column 4), Table A7, (column 5), and Table A8 (coumn 6).

Table A3: MTurk Sample, Unweighted

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	53.8 [50.2, 57.3]	38.4 [34.9, 41.9]
Carson	49.4 [45.8, 53.0]	42.6 [39.1, 46.2]
Christie	55.2 [51.7, 58.8]	36.0 [34.4, 41.3]
Cruz	47.6 [44.1, 51.3]	46.4 [42.8, 49.9]
Rubio	47.8 [44.2, 51.4]	47.2 [43.6, 50.7]

95% confidence interval in brackets

Table A4: ANES Pilot Sample

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	61.3 [56.8, 65.9]	34.5 [30.1, 39.0]
Carson	47.7 [43.0, 52.4]	46.6 [41.9, 51.3]
Cruz	50.6 [46.0, 55.4]	43.9 [39.2, 48.5]
Rubio	51.1 [46.4, 55.8]	44.5 [39.9, 49.2]

95% confidence interval in brackets

Table A5: MTurk Sample, Pairwise Only

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	51.8 [47.3, 56.4]	36.9 [32.5, 41.3]
Carson	46.1 [41.5, 50.6]	42.2 [37.7, 46.7]
Christie	51.2 [46.6, 55.7]	39.0 [34.6, 43.5]
Cruz	45.0 [40.5, 49.5]	46.5 [42.0, 51.0]
Rubio	45.0 [40.5, 49.5]	47.8 [43.2, 52.3]

95% confidence interval in brackets

Table A6: MTurk Sample, Rank Orderings Only

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	57.0 [51.2, 62.8]	40.9 [35.2, 46.6]
Carson	54.9 [49.1, 60.7]	43.4 [37.6, 49.1]
Christie	61.9 [56.2, 67.6]	36.0 [30.4, 41.6]
Cruz	52.1 [46.3, 57.9]	46.2 [40.3, 52.0]
Rubio	52.4 [46.6, 58.3]	46.2 [40.3, 52.0]

95% confidence interval in brackets

Table A7: MTurk Sample, No Paul Supporters

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	56.0 [52.3, 59.7]	37.6 [33.9, 41.2]
Carson	51.4 [48.0, 55.5]	41.1 [37.4, 44.8]
Christie	58.0 [54.3, 61.7]	36.0 [32.4, 39.6]
Cruz	49.9 [46.1, 53.6]	44.4 [40.7, 48.2]
Rubio	49.7 [46.0, 53.4]	45.6 [41.9, 49.4]

95% confidence interval in brackets

Table A8: MTurk Sample, Weighted

	Pct Strictly Preferring Trump	Pct Strictly Preferring Other
Bush	65.6	27.4
	[59.8, 71.4]	[22.1, 32.6]
Carson	62.2	30.6
	[56.2, 68.3]	[25.1, 36.0]
Christie	65.9	30.9
	[60.2, 71.7]	[25.3, 36.5]
Cruz	55.6	37.6
	[49.1, 62.0]	[31.6, 43.6]
Rubio	56.3	38.5
	[50.0, 62.5]	[32.5, 44.4]

95% confidence interval in brackets

Plurality Rule

As noted in footnote 8, majority rule counts indifferent voters while plurality rule ignores them. Here, we analyze social preferences in favor of Trump using plurality rule instead of majority rule. Table A9 presents the results. If anything, throwing out indifferent voters strengthens the appearance of support for Trump. Using the point criterion, Trump is a Condorcet winner in both the unweighted MTurk and unweighted ANES samples. However, using the interval criterion, the main result is qualitatively the same as in our main analysis. There is enough uncertainty about the aggregate preference between Trump and Carson, between Trump and Cruz, and between Trump and Rubio that Trump is a member of the core but not a Condorcet winner.

Bootstrap Simulations

We also check the robustness of our conclusions by using simulation methods. We draw 1,000 bootstrap samples of 1,000 Republican voters using the ANES pilot data, then compute the percentage of samples for which Trump is a Condorcet winner or a member of the core. Trump is a Condorcet winner in only 17% of our bootstrapped samples but a member of the majority rule core in 95% of them. This discrepancy matches our findings in Figure 1 and is due to the fact that there is a non-trivial portion of respondents who are indifferent between Trump, Cruz, and Rubio.

Using plurality rule, Trump is a Condorcet winner in 63% of our bootstrapped samples and a member of the core in 65%. The small difference is due to the fact that ties are rare (the electorate is indifferent only when the number of voters that strictly prefer Trump is equal to the number that strictly oppose him).

Strikingly, Trump is *never* a Condorcet loser. For Trump to be a Condorcet loser, he must have strictly less support than *every* other candidate in pairwise contests. In our simulations, Trump may sometimes lose to another candidate, but never to all other candidates.

Table A9: Plurality Rule

	Pct Strictly Preferring Trump	
	MTurk	ANES
Bush	58.3	63.7
	[54.7, 62.0]	[59.1, 68.3]
Carson	53.7	50.1
	[50.0, 57.4]	[45.3, 54.9]
Christie	59.3	–
	[55.7, 63.0]	–
Cruz	50.7	53.1
	[47.0, 54.4]	[48.3, 57.9]
Rubio	50.3	52.8
	[46.7, 54.0]	[48.1, 57.6]

95% confidence interval in brackets

Even if we restrict our attention to Ted Cruz and Marco Rubio as the best candidates to defeat Trump, in 1,000 bootstrapped samples, Trump strictly loses to Rubio once and never to Cruz (and hence never loses to both).