

10 A Demarcation Problem for Political Discourse

*David Killoren, Jonathan Lang,
and Bekka Williams*

One of the main roles for a citizen in a democratic society is to engage in *political discourse*, i.e., to engage in (written or oral) conversation about political issues. Political discourse often (though not always) takes the form of debate, in which the participants take up a position on a disputed question and try to persuade other participants of their position. People from all walks of life take part in these debates, in all sorts of venues—from television and radio, to Facebook and Twitter, to workplaces and neighborhood bars. In this debate-style form of political discourse, a wide range of different questions can be up for debate (and, as the discourse unfolds, the questions themselves can evolve). But during the run-up to an important election, the simplest and most common question is: *Whom should we elect?* In the United States at present (late May, 2016) there is a lot of discussion about whom “we” (the American people) should elect as President of the United States in the upcoming presidential election.

In this paper, we introduce a problem for political discourse surrounding that simple question—*Whom should we elect?* We’ll take political discourse about the American presidential election of 2016 as our main example.

Normally, a participant in political discourse supports a given candidate and opposes others; it is unusual to get involved in political discourse without supporting any particular candidate. To *support a candidate* seems to have at least two components: (1) to support a given candidate typically includes playing a role of advocacy in political discourse—presenting arguments in favor of that candidate, responding to objections against that candidate, raising objections against opposing candidates, etc.; and (2) to support a given candidate is to be prepared to act in a certain way—e.g., to be prepared to cast a vote for that candidate when the election occurs.

But how do we—and how should we—go about selecting a candidate to support? One method that quickly presents itself is what we’ll call the *two-stage procedure*.

In the first stage, the *option-identification* stage, one produces an initial list of candidates—the *option list*. The option list is (roughly) a list of candidates whom one could support; alternatively, it is a list of candidates who are up for consideration. In 2016, any voter’s option list is likely to include Hillary Clinton

and Donald Trump, as they are the presumptive nominees of the two major parties. The list may also include various third-party candidates—e.g., Gary Johnson, who will be the Libertarian Party candidate for President, or Jill Stein, who is poised to be the Green Party candidate. And the list may include Bernie Sanders, who as of this writing is still trying to win the Democratic nomination. There may be others appearing on any given voter's option list as well.

In the second stage, which we'll call the *evaluation stage*, one compares the various candidates in the option list in order to reach a judgment about which candidate is best. Thus, for instance, a utilitarian voter may try to choose the candidate who would as President do the most to promote general happiness, whereas a libertarian voter may try to choose the candidate who would as President do the most to promote individual freedom and protect individual rights. Politically conservative voters may proceed with the view that the best candidate will be the one who will keep things more or less as they are (or will restore things to the way they used to be). Other voters may rely mainly on considerations about virtue, with the thought that the best candidate for President will be the one with the best character. Some voters may eschew general principles altogether and rely entirely on gut-level feelings to decide which candidate is the best one. In short, different voters will approach the evaluation stage in different ways.

We conceive of the two-stage procedure as a procedure that each individual participant in political discourse deploys for herself in order to make up her own mind about whom to support. Thus different voters may have different actual option lists: e.g., a longtime Green Party activist will almost certainly include Jill Stein in her option list, whereas Stein will be absent from the option list of a voter who has never even heard of the Green Party. Of course, voters can also influence one another regarding their option lists. Thus, the Green Party activist may, in conversation, persuade a less engaged voter to at least consider Jill Stein for President—that is, to add Jill Stein to his option list. Likewise, at the evaluation stage, voters can profoundly influence one another, even though each voter must execute the evaluation stage on her own to make up her own mind about whom to support.

We do not claim that all or even most voters consciously or explicitly deploy the two-stage procedure. However, it seems plausible that most voters approach the question of whom to support in a way that *functionally amounts to* our two-stage procedure. In real-world political discourse, much of the discussion appears to occur at (what we are calling) the evaluation stage: voters in political discourse typically seem to be involved in discussion and debate about which candidate is best. But one cannot even discuss which candidate is best unless one first knows which candidates are up for consideration—i.e., it seems that one cannot enter the evaluation stage without first having passed through the option-identification stage. Thus, even if participants in political discourse never consciously set out to assemble an option list before they begin to think about which candidate to support, it appears that they must somehow come to possess such a list.

In any real-world discourse—especially discourses that transpire over the course of months or years, as is the case in a Presidential election—it seems likely that the two-stage procedure will be deployed in a dynamic, ongoing way. For instance, when a new candidate enters the race and begins to gain supporters, this may result in an expansion of the option list. When a candidate drops out of the race, this may result in a contraction of the option list. And so we might say that the option list is, in effect, curated on an ongoing basis by the voter, rather than decided once and for all. And when the option list changes, the voter may then have to re-enter the evaluation stage and may thus reach a new verdict about which candidate is best. Similarly, when new considerations come to light—e.g., one’s favorite candidate turns out to have an unsavory past, or one’s least-favorite candidate puts in a very persuasive debate performance—then a voter may choose to re-enter the evaluation stage.

We suspect that many voters move back and forth between stages—continually modifying their option list and revising their views in light of new information and new arguments about which candidate in the option list is best. Thus the two-stage procedure is useful as a model of voter decision-making only if we remember that it does not result in a single, once-and-for-all judgment, but is instead a pattern of reasoning that can be revised and redeployed continually.

The central question for this paper is then: What is the right way (and what are some of the wrong ways) to assemble an option list? This is a demarcation problem; it is the problem of how a voter can justify *drawing a line* that separates candidates to be considered and evaluated from candidates to be set aside and ignored.

This problem is important. Candidates have good reason to work very hard to earn a place on voters’ option lists. Bernie Sanders, for example, is currently fighting to remain in consideration among Democratic primary voters; and his opponent, Hillary Clinton, and her supporters are doing whatever they can to persuade voters to give up on Sanders (i.e., to persuade voters to drop Sanders from their option lists). This behavior is not the least bit mysterious. If a candidate is not on voters’ option lists, then he cannot win the election.

We’ll proceed as follows. Throughout the paper, we’ll consider a series of different criteria that can be used to select an option list. In the first section, we’ll consider a criterion based on media prominence, according to which a candidate earns a place on our option list by being prominently discussed in mainstream news sources. In the next section, we’ll consider a criterion based on legal electability, according to which a candidate earns a place on our option list simply by being legally electable. In the third section, we’ll consider a criterion based on probability, according to which a candidate earns a place on our option list if her election is (at least) minimally probable. In the fourth section, we’ll consider a pair of consequentialist criteria, according to which a candidate earns a place on our option list if (roughly) the consequences of including her in the option list are sufficiently good. In the fifth section, we’ll consider (and reject) a contrastivist solution to the demarcation problem. And in the final section, we’ll consider a criterion based on practical possibility,

according to which a candidate earns a place on our option list as long as we (the electorate) can elect her.

As we indicated above, we'll focus throughout the paper on the demarcation problem as it applies particularly to the political discourse in presidential politics in the United States in 2016. We think our remarks here can apply fairly straightforwardly to a wide range of other political discourses in the United States and elsewhere; but we will not attempt to explicitly draw general lessons. In the end, we do not intend to finally solve our demarcation problem for political discourse. Our primary aim is to show that this problem is important and difficult, to survey some of the main available solutions to it, and to develop some of the arguments for and against these solutions.

The Media Prominence Criterion

In the next several sections, we want to consider various criteria, or rules, that a voter could use to determine which candidates to include in her option list and which ones to exclude. Here's the first one that we'll discuss:

The Media Prominence Criterion: A candidate X is to be included in the option list if discussion of X is prominent in mainstream news media.

A voter who relies on the Media Prominence Criterion will certainly include both Hillary Clinton and Donald Trump in her option list, as they are both obviously prominent in news coverage of the election. She is likely also to include Bernie Sanders in her option list, at least for now, because Sanders is still receiving a lot of mainstream media attention. She will not consider Jill Stein (the likely Green Party candidate) or Gary Johnson (the likely Libertarian Party candidate) because these individuals have received almost no discussion in the press thus far. And there are a great many others she'll exclude, as well. For instance, she will exclude Vermin Supreme, a bizarre individual who wears a boot as a hat and has run for President in elections since 2004, and has (unsurprisingly) received very little media attention.

We suspect that few voters would openly endorse or defend the Media Prominence Criterion, but it seems that this criterion is in fact implicit in many voters' approach to political discourse. For example, Donald Trump has by many measures garnered more media attention than any other primary candidate (in either party),¹ and this disparity appears to partially explain Trump's high level of support in the polls as well as the stagnating support seen for candidates who struggle for coverage.² In response, Bernie Sanders has claimed that there has been a "Bernie Blackout" in political journalism³ and that Donald Trump's rise is "an indictment of the media."⁴ Whether such claims are true is up for debate, but it seems clear that the quantity of candidates' media coverage causally contributes (for better or worse) to the range of candidates that voters are willing to consider; that is, it appears that voters rely heavily on the media to tell them which candidates are worthy of consideration and which are not.

Thus, we suggest, many voters behave in political discourse *as if* they rely on the Media Prominence Criterion, even if they do not endorse that criterion in a conscious or explicit way.

Yet it is very difficult to see how the Media Prominence Criterion could be justified. The mere fact that prominent journalists and pundits are discussing a candidate does not seem like a reason, on its own, to think that the candidate is worthy of voters' consideration; likewise, the fact that a candidate *isn't* being prominently discussed does not seem like a reason to think the candidate *isn't* worthy of discussion.

To support this point, consider other contexts. Imagine a recent college graduate who is trying to decide on a career. Such a person would be ill-advised to consider all and only those careers that are given prominent attention in mainstream media. After all, if our college graduate relies on such a criterion, she may overlook an unusual, little-discussed, or simply unpopular career that will suit her perfectly. Or consider a more trivial decision-making context: a couple who are trying to decide on a place to go out for dinner. If they want to consider the best dining options available to them, they should not confine their attention only to those restaurants that are very prominent in, say, online restaurant listings; such an approach could exclude many out-of-the-way restaurants that might be every bit as good, or better than, the most obvious locations.

Thus, it seems that versions of the Media Prominence Criterion applied in certain other decision-making contexts are hard to justify. Despite this, of course, it could well be that the Media Prominence Criterion *as applied to political discourse* is defensible. As we have emphasized, however, the mere fact that a given candidate has achieved (or has failed to achieve) media prominence does not, on its own, seem like any kind of reason to think that voters should give (or not give) serious consideration to that candidate. In other words, the Media Prominence Criterion seems like an utterly arbitrary answer to our demarcation problem.

Yet it could well be that media prominence of a given candidate is correlated with some other feature (or set of features) which does (or do) give voters a reason to consider that candidate. Of course, in order to determine whether this is true, we'll need to identify the feature(s) in question—that is, we'll need to identify an altogether different, more fundamental criterion. Once we do that, then it might turn out that the Media Prominence Criterion is justifiable as a *reliable heuristic* for that more fundamental criterion. Hence, even if we may ultimately want to defend the Media Prominence Criterion as a reliable heuristic, we'll still need to have an alternative, more basic criterion in terms of which to defend it.

The Legality Criterion

Here is another approach that a voter could use in deciding on her option list:

The Legality Criterion: A candidate X is to be included in the option list if that candidate could be legally elected and sworn in as President.

A first problem for the Legality Criterion is that it does not exclude very many candidates from consideration—although it does exclude *some* candidates. The Constitution states:

No person except a natural born citizen, or a citizen of the United States, at the time of the adoption of this Constitution, shall be eligible to the office of President; neither shall any person be eligible to that office who shall not have attained to the age of thirty-five years, and been fourteen years a resident within the United States.

(Article II, Section 1)

These restrictions mean that we cannot legally have Arnold Schwarzenegger, for instance, as President (since he is not a natural born citizen). Nor can we have a President Justin Bieber (since he is not thirty-five years old). Plausibly, we also cannot legally have Huckleberry Finn or Abraham Lincoln as President (assuming that fictional people and dead people both fail the Constitution's residency requirement). If we restrict our option list to those who can be legally elected and sworn in, then we are able to exclude *some* individuals from consideration. But there is still an enormous number of potential candidates left: that is, there is an enormous number of individuals in the United States who meet all of the constitutional (and other legal) requirements for being elected and sworn in as President.

This means that if voters rely on the Legality Criterion, they will need to consider a massive collection of different individuals in political discourse. They'll certainly have to consider Hillary Clinton and Donald Trump. And they'll have to consider third-party candidates such as Gary Johnson and Jill Stein. They'll also need to consider extremely weird candidates as well, such as Vermin Supreme. And they'll need to consider literally *millions* of others. It bears noticing that some of those millions of people would make very good presidents. The list of legally electable presidents includes not just professional politicians, but also humanitarians, philosophers, economists, historians, scientists, activists, artists, spiritual leaders, etc.—people who have devoted their lives to the advancement of human rights and the common good, and who could be enormously beneficial to the United States and to the world if they were given the power of the American presidency.

Given that there are millions of legally electable potential presidents, if voters were to rely on the Legality Criterion, then this would utterly transform political discourse. At present, political discourse focuses on a very small, very manageable set of different candidates. For many people, that set includes *just two* individuals—the Republican nominee and the Democrat nominee. For others, the set includes a few others—usually a handful of third-party candidates. With such small option lists, voters typically find it relatively easy to perform the evaluation stage, the second stage of the two-stage procedure that we described in the Introduction. After all, even a casually attentive voter is normally able to compare and contrast two or three different candidates. But if

the option list were to include millions of individuals, the evaluation stage would become a full-time job. Indeed, it might not be possible for a single voter, or even a large team of voters, to meaningfully compare millions of different potential presidents.

Therefore, the first problem for the Legality Criterion is a practical one: it would produce unmanageably huge option lists. A second problem for the Legality Criterion is that the mere legality of electing and swearing in a candidate does not, on its own, seem like a positive reason to think that that candidate ought to be considered as a candidate for president (even though a candidate's legal electability may count as a *necessary* condition of her inclusion in one's option list). Consider again the couple from the previous section who are trying to decide where to go for dinner: it would be a bit odd for them to approach their problem by first listing every restaurant anywhere in the world that they could legally choose for dinner, and then to attempt to compare all of those restaurants. One might reasonably ask them why they are using legality as a way to construct their initial list. And it would be reasonable to ask the same question of any voter who relies on the Legality Criterion.

The Probability Criterion

One of the reasons why the Legality Criterion seems problematic is that it requires voters to consider candidates who almost certainly will not be elected as president. Consider Vermin Supreme—the candidate who wears a boot as a hat. Many people will say that Vermin Supreme should be excluded from consideration because he is sure not to win. Of course, Vermin Supreme also has a number of other disadvantages as a candidate. No one, not even Vermin Supreme, believes that Vermin Supreme would make a good President. But there are many individuals who *would* perform very well as President, and could be legally elected as President, but seem unworthy of consideration simply because they are almost certain not to be elected. Consider Russ Shafer-Landau, the moral philosopher, who is by all accounts exceedingly wise and virtuous. Assume, at least for the sake of argument, that Shafer-Landau would perform well as President. Assume even that he would be better as President than Hillary Clinton or Donald Trump. Nevertheless, many people will say that Clinton and Trump deserve a place in our option lists, whereas Shafer-Landau does not, simply because the odds of Shafer-Landau being elected are so remote.

Thus, as an alternative to the Legality Criterion, we may wish to consider a criterion founded on probability, such as the following:

The Probability Criterion: A candidate X is to be included in the option list if the election of X is (at least) *minimally probable*.

The first problem for defenders of the Probability Criterion will be to present and defend a minimum probability cut-off. How probable does a candidate's election need to be in order to count as “minimally probable”? If a candidate's

odds are 17 percent, is that sufficient? What about 1 percent? The challenge is to identify a probability cut-off that excludes *just enough* candidates. Here there may be intramural disagreement among defenders of the Probability Criterion. Some may take a relatively permissive approach. They may say that certain quite unlikely candidates such as Bernie Sanders, Jill Stein, and Gary Johnson have good enough chances to merit inclusion—although presumably even the most permissive advocates of the Probability Criterion will want to say that *extremely* unlikely candidates, such as Russ Shafer-Landau, do not have good enough chances. Others may take a relatively restrictive approach. They may say that candidates like Sanders, Stein, and Johnson might be excellent candidates, better even than the presumptive nominees of the major parties, but their odds are just not good enough to warrant consideration. One problem, for defenders of the Probability Criterion, will be to explain why a given probability cut-off is to be used; presumably, the cut-off should not just be chosen at random.

A deeper problem for the Probability Criterion will be to explain why probability should be taken into account at all. For, in many contexts, probability seems utterly irrelevant to the question of which options one should consider.

Consider the following case. Moe has made a promise to meet Susan for coffee at noon, and it is now 11:45 a.m. Moe can still keep his appointment if he leaves right away. In this situation, there is a wide range of things that Moe can do: he can leave right away to keep the appointment; he can take a nap; he can jump out the window—and so on. Now suppose that Moe is extremely lazy and thus is very unlikely to choose to keep the appointment. (Odds are he'll take a nap instead.) In this case, the mere fact that Moe is very unlikely to fulfill his promise does not seem like any reason whatsoever to exclude that option from consideration. In this case, if we are interested in the question of what Moe ought to do, we should not confine our attention merely to the things that Moe probably will do. Likewise, if we are interested in the question posed at the outset of this paper—*Whom should we elect?*—then it may seem like a basic error to exclude a given candidate from consideration just because we are not likely to elect him or her.

To respond to this analogy, defenders of the Probability Criterion need to present a relevant difference between the case of Moe and the case of political discourse. That is, they need to explain why probability is relevant in deciding which options to consider in political discourse, even though probability (as against mere possibility) is irrelevant in deciding which options to consider in a case of individual decision-making (such as Moe's). We'll return to this issue below.

Consequentialist Criteria

It is possible to approach the construction of an option list in a consequentialist way. We'll consider two possibilities along these lines.

Let's say that the *locally optimific option list*, for any given individual in any given circumstance, is the option list whose adoption by that individual would

have the best consequences in that circumstance. By way of illustration, suppose you are engaged in political discourse with a friend who is considering two candidates—Ronald Stump and Gerald Clump. As it happens, Stump and Clump are both terrible candidates, although Stump is slightly better than Clump. Your friend is currently leaning toward Clump, the worse option. You think your friend should vote for neither Stump nor Clump; you have a third alternative in mind. However, you also know that if you even mention any candidate other than Stump or Clump, your friend will dismiss you completely. The only hope you have of reasoning with your friend is to concentrate on an option list that includes only Stump and Clump. And if you do that, then you have a good chance of persuading your friend to vote for Stump, the lesser of the two evils. In this case, it appears that adopting an option list containing only Stump and Clump would have the best consequences—that is, Stump and Clump comprise the locally optimific option list in this circumstance.

We want to consider a criterion according to which the option list in any given circumstance should be locally optimific:

The Act-Consequentialist Criterion: A candidate X is to be included in the option list if X would be included in the locally optimific option list.

We grant that in some cases, at least, the fact that a given option list is locally optimific looks like a reason in favor of it. But the Act-Consequentialist Criterion nonetheless fails to describe an acceptable approach to political discourse in general.

To begin, notice that the Act-Consequentialist Criterion looks like a straightforward application of act-consequentialism—the view that, in general, in any given circumstance of choice, one should always do whatever would have the best consequences. The Act-Consequentialist Criterion is, specifically, an application of act-consequentialism to one particular task in political discourse—the task of assembling an option list. This is but one of several different tasks that we must perform in the course of engaging in political discourse. As we'll now argue, however, a broadly act-consequentialist approach to political discourse is hard to defend; and this causes problems for the Act-Consequentialist Criterion.

The first problem with a broadly act-consequentialist approach to political discourse is an obvious one: it simply will not sit well with non-consequentialists. Consider the evaluation stage, the second stage of our two-stage procedure—the stage at which one has to decide which candidate in one's option list to support. To be sure, *some* people will apply consequentialist principles in order to make that decision, but certainly not all of us will do so. For instance, imagine a libertarian non-consequentialist who believes that we should support the candidate who has the most respect for property rights, individual liberty, and so on. Suppose this libertarian non-consequentialist is ready to grant that respect for rights and liberties does not always produce the greatest good, but she thinks that questions about the greatest good are simply beside the point.

She thinks that we must support respect for rights and liberties just because that is the right thing for us to do, regardless of whether this will have especially good consequences. Anyone who takes that sort of view, or any other non-consequentialist view, will be unhappy with a broadly act-consequentialist approach to political discourse.

But there is a more fundamental problem with the broadly act-consequentialist approach to political discourse. Notice that participation in political discourse rarely has very good consequences. Most of us are unable to persuade very many people (if any) to change their minds about politics; most conversations about politics simply have no significant consequences at all. Thus, we suspect, whenever one is engaged in political discourse, there is almost always something else that one could be doing instead that would have better consequences, such that an act-consequentialist approach to political discourse would mean that nearly all of us should radically curtail the extent of our participation in political discourse. Therefore, anyone who thinks that political discourse is a civic duty, or even an acceptable pastime, for most ordinary citizens ought to reject the act-consequentialist approach to political discourse.

In response to this, the defender of the Act-Consequentialist Criterion could say that she is not recommending that we take a broadly act-consequentialist approach to political discourse. Rather, she is recommending only that we take an act-consequentialist approach to *one specific task* in political discourse—the task of constructing an option list. Once the option list is constructed, then we are free to approach the other tasks in political discourse (including the task of deciding whether to participate in political discourse at all) according to any non-consequentialist approach we choose. But this is not plausible. If we are going to approach most of our activities in political discourse in a non-consequentialist way, then it is quite difficult to see why we should accept an act-consequentialist approach for the task of constructing an option list.

Let us now consider an alternative to the Act-Consequentialist Criterion. To describe the alternative we have in mind, we'll need to introduce a new bit of terminology. Let's say that the *globally optimific option list*, for any given individual in any given circumstance, is the option list whose universal adoption by everyone *in that individual's political society* would have the best consequences. So, for instance, if it would be best for our society if we were to continue to consider Bernie Sanders for President, even though he is presently very unlikely to win the Democratic nomination, then Bernie Sanders belongs on the globally optimific option list. Likewise, given that it would not be best for our society to consider Vermin Supreme as a candidate for President, Vermin Supreme has no place in the globally optimific option list.

With that piece of terminology in place, let us now consider the following criterion:

The Rule-Consequentialist Criterion: A candidate X is to be included in the option list if X would be included in the globally optimific option list.

We are calling this a “rule-consequentialist” criterion because it is a direct application of rule-consequentialism—the view that, in general, in any given circumstance of choice, one should always act in accordance with whatever rule or set of rules would have the best consequences if they were universally adopted by everyone⁵ in one’s society.

One of the advantages of the Rule-Consequentialist Criterion is that it, unlike the Act-Consequentialist Criterion, yields a procedure by which members of a political society can coordinate a common option list. The *locally* optimific option list will vary from individual to individual, simply because individual circumstances can vary; but the *globally* optimific option list will, by definition, be shared by all of the individuals in a given political society. Thus, if we adopt the Rule-Consequentialist Criterion, then we can all reach a point where we are considering and discussing the same collection of candidates.

But the Rule-Consequentialist Criterion is hard to defend—simply because rule-consequentialism is hard to defend. We won’t rehearse all of the many well-known problems for rule-consequentialism here.⁶ We’ll just point out that rule-consequentialism suffers from a basic problem of motivation. It is simply unclear why anyone should be required to act in accordance with a set of rules that *would* have the best consequences if they were universally adopted, given that those rules *aren’t* universally adopted. Rule-consequentialism thus comes across as a kind of oblivious idealism—an idealism that requires us to indulge the fantasy that we are living in an ideal society even though our society is in fact very far from the ideal. Given that rule-consequentialism is very hard to motivate as a general principle, it is hard to see why we should accept rule-consequentialism as applied to the narrow task of constructing an option list.

Contrastivist Solutions

To this point, it has proved difficult to arrive at a defensible criterion by which to identify an option list. We’ll return to the search for a defensible criterion in just a moment. But first we want to consider a more radical solution, one which claims that the selection of the option list is, in a certain way, beside the point. According to this solution, anyone who properly undertakes the evaluation stage will be able to reach a correct judgment—regardless of the option list that they choose during the option-identification stage.

At first glance it seems very counterintuitive to think that one’s choice of option list is completely irrelevant to ascertaining what an agent should do. To see this, recall the case of Moe from page 185. Imagine you are engaged in the deliberative process of deciding what it is that Moe should do. Then it seems all-important that you should first arrive at an option list that, at the very least, contains the option that Moe *keep his promise to Susan*. If at the outset of the deliberative process that option is excluded, then you will not even consider the possibility that Moe ought to keep his promise to Susan. This in turn will lead you to (erroneously) judge that Moe ought to do something other than fulfill his

promise. Hence, it seems reasonable to conclude that failure to carefully select our option list can lead to incorrect judgments about what ought to be done.

Yet an alternative view is available. According to this view, the very meaning of judgments about what ought to be done encodes the option list used to reach those judgments. For example, suppose that a person who begins deliberating about what Moe should do exits their option-identification stage with a short option list that excludes the option of Moe keeping his promise—e.g., the two-option list {have a sandwich, burn down a hospital}. And suppose this same person then uses that list to go about the business of deciding what Moe ought to do (i.e., she progresses through what we have called the “evaluation stage”). After reaching her verdict, imagine that the person confidently utters:

- 1 Moe ought to have a sandwich.

This seems the wrong result, since it would presumably be better for Moe to keep his promise rather than have a sandwich. However, for proponents of the alternative view under consideration, the resulting judgment above is *not* wrong. It only seems wrong, they will argue, because we are not interpreting the meaning of the judgment correctly. The full meaning of the above judgment doesn’t reveal itself until we take into account the option list used to arrive at the judgment. Once we take into account the option list utilized, the above judgment can be seen as being elliptical for the following:

- 2 Moe ought to have a sandwich *rather than* burn down a hospital.⁷

But if the person’s judgment that “Moe ought to have a sandwich” is interpreted as meaning that “Moe ought to have a sandwich rather than burn down a hospital,” then there seems to be nothing erroneous about the person’s judgment after all. Given that having a sandwich is arguably a much better option than burning down a hospital, the judgment seems entirely correct.

Of course, if our deliberator had started out with a different option list—e.g., {have a sandwich, burn down a hospital, keep promise to Susan}—then she would likely have reached a different result, namely that:

- 3 Moe ought to keep his promise to Susan.

But given the expanded option list used to arrive at this result, the judgment would actually be elliptical for (and would thus mean) the following:

- 4 Moe ought to keep his promise to Susan *rather than* burn down a hospital or have a sandwich.

And notice that (4) is not inconsistent with (2). In fact, once we understand the meanings of these two different judgments, both seem to be true.

We call this view about the semantics of ‘ought’ claims the *contrastivist view*. The view is contrastivist in that it interprets the meaning of judgments about what we ought to do as tacitly referencing a contrast class of alternatives, alternatives that comprise the option list that is evaluated to arrive at the judgment.

We can develop a similar contrastivist semantics for political discourse. Such an account would claim that the meaning of any judgment of the form “We ought to elect X” is determined by the option list used to arrive at that judgment. For instance, if we are given an option list such as {Hillary Clinton, Donald Trump}, then we might decide that:

5 We ought to elect Hillary Clinton.

But such a judgment would really just be elliptical for:

6 We ought to elect Hillary Clinton *rather than* Donald Trump.

If, on the other hand, we are given an option list such as {Hillary Clinton, Donald Trump, Bernie Sanders}, then we might judge that:

7 We ought to elect Bernie Sanders.

But by this we would only be saying that:

8 We ought to elect Bernie Sanders *rather than* Hillary Clinton and Donald Trump.

And our support might shift again if Jill Stein is included in the option list—and might shift yet again if Russ Shafer-Landau is included in the option list. But, according to the contrastivist account, these shifts in support do not represent any fundamental change of mind. Instead, they each represent fundamentally compatible judgments, each of which is equally correct. To think otherwise is to fail to take account of the semantic influence that one’s chosen option list has on the meaning of the claims one makes in political discourse.

This contrastivist account is compatible with (but does not require) a further view, according to which there just isn’t any correct or incorrect way to assemble the option list. Once augmented with this view, the contrastivist account provides a kind of deflationary solution to the demarcation problem that motivates this paper. If we suppose, in accordance with the contrastivist account, that every judgment with the surface form “S ought to x” is merely elliptical for a judgment of the form “S ought to do x *rather than* y” (where y is a set of alternatives determined by the option list used to arrive at that judgment), then it seems to follow that our initial selection of an option list does not affect our ability to reach true judgments at all. Thus, any option list we might choose has the potential to yield true judgments; and thus, one might think, there is really no way to go wrong in selecting the option list at the start. Of course, on

this account, it is still possible to err *after* we've selected our option list (e.g., if we mistakenly judge that one option in the option list is better than another, when in fact the first option is worse than the second one). But contrastivism seems to imply that we cannot make any error *in the selection of the option list itself*; and so it is simply unnecessary to search for a criterion with which to justify the selection of one option list rather than another.

This solution to our demarcation problem is tantalizing, but far too quick. To begin, note that production of true 'ought' judgments is not the *only* thing that political discourse aims to achieve. (If it were, people would just go around speaking tautologies in the context of political discourse!) To see why this matters, return to the case of Moe. Even if the contrastivist analysis of our judgments about Moe is correct, it still seems clear that there are better and worse ways to construct an option list with regard to Moe. For instance, we seem clearly to be missing something if our option list includes just one option, as in {have a sandwich}. If we begin with this option list and then reach the judgment that *of course* Moe ought to have a sandwich, then we seem to have gone wrong somehow—regardless of whether we accept a contrastivist interpretation that implies that this judgment is somehow true. Likewise, if we consider an option list like {burn down a hospital, detonate an atomic bomb in a major city}, and then reach the judgment that Moe ought to burn down a hospital (because clearly this is better than detonating an atomic bomb in a major city), then we seem to have made a mistake at some point along the way. In particular, we have mistakenly left out some options that ought to have been considered.⁸ Simply put, even on the contrastivist analysis, it seems eminently plausible to suppose that there are better and worse ways to construct an option list for Moe.⁹

And much the same will be true of political discourse. Imagine a misogynist who is unwilling to even consider candidates who are not men. If this misogynist refuses to consider Hillary Clinton or Jill Stein for President and thereby reaches the judgment that we ought to elect, say, Donald Trump (or Bernie Sanders, or whomever), then it seems quite clear that an error was made in the selection of the option list—even if we are prepared to grant that the misogynist's judgment about whom to support is perfectly sound given the option list with which he started.

It is important to underscore that we are not objecting to a contrastivist semantics for 'ought.' We grant for the sake of argument that contrastivism might provide the most plausible semantics for that word. What we deny is the further claim (which is neither required nor entailed by the contrastivist semantics) that there are no wrong ways to choose an option list. We'll still be in need of a principled criterion to use in selecting an option list regardless of whether we accept a contrastivist account of political discourse.

The Practical Possibility Criterion

In this section, we're going to try to motivate a strong inclusivism about political discourse. According to this inclusivism, participants in political

discourse should assume an option list that includes every candidate who *can* be elected. This view shares practical difficulties with the Legality Criterion (considered on pages 182–184). As we noted before, the Legality Criterion would imply that participants in political discourse must consider option lists that are unmanageably large. The inclusivism that we'll suggest here has precisely the same problem. Despite this problem, we think that this inclusivism has some advantages that cannot be easily dismissed.

Return to the case of Moe. When we are trying to decide what Moe ought to do, what are the options that we should be prepared to consider? Here's a simple answer: we should be prepared to consider any and all of the various actions that Moe *can* perform. Thus, for instance, no matter how desirable it might be for Moe to end world poverty instantly with a snap of his fingers, this is not an action that we should consider, simply because it is not something that Moe can do. But anything that Moe can do is fair game for consideration—even if he is very unlikely to do it. Thus, we should be prepared to consider the possibility that Moe ought to keep his promise, even if he is not likely to do so. And, by the same token, we should also be prepared to consider various outlandish options—e.g., the option of Moe dumping the entire contents of his bank account into GiveDirectly's coffers.¹⁰ If we are going to set such outlandish options aside, we'll need to do so at the evaluation stage, by showing that they are in some sense inferior to the alternatives; we cannot simply discard such options at the outset without consideration.

We think that this view of how to reason about what Moe ought to do is highly plausible. And we want to suggest that an analogous view of how to approach political discourse is also plausible.

At the beginning of this paper, we suggested that the central question of political discourse surrounding the upcoming American presidential election is this: *Whom should we elect?* This question is structurally similar to the question about Moe: *What should Moe do?* Thus, we might want to approach political discourse in a way that is analogous to the approach we've suggested for the case of Moe. The idea here is straightforward. Our option list should include all (and only) the different candidates we *can* elect, just as Moe's option list should include all (and only) the actions he *can* perform. Or, if we want to introduce a bit of jargon, we might say that the option list in the case of Moe should include all of the actions that are *practically possible* for him, and that the option list in the case of political discourse should include all of the *practically possible* candidates. If it is practically possible for us to elect a given candidate, then she belongs in our option list; otherwise not. Here's a criterion based on this idea:

The Practical Possibility Criterion: A candidate X is to be included in the option list if it is practically possible to elect X.

Notice that the idea of practical possibility that we are working with here is not directly related to probability. The idea is *not* that if a given candidate is unlikely to win, then her election is practically impossible. True, people do sometimes

use the language of practical possibility in this way. For instance, it has been said that “Bernie Sanders cannot win,” by which it is meant that he almost certainly will not win. And that is true—he almost certainly will not win. But there is also an important sense in which he *can* win even though he almost certainly will not. He can win because we can elect him. Indeed, we can elect him rather easily. All that we need to do, in order to elect him, is to show up to support him at the polls in sufficient numbers on election day, and then he would win.

To be clear, the claim here is that *we*—as a collective, i.e., as the mass of voters—can easily elect Bernie Sanders. The claim is not that *any given individual* can bring it about that Bernie Sanders is elected. Indeed, no one is able to ensure Sanders’s election. But of course, it is in the nature of the democratic process that no one voter is able to bring about any electoral outcome. An individual voter is unable to bring it about that Hillary Clinton is elected, just as surely as she is unable to bring it about that Bernie Sanders is elected. In other words, it is difficult (indeed, it is not even possible) for an individual voter to elect Hillary Clinton, but it is easy for *us*—the collective mass of voters—to elect her. The same seems true of Bernie Sanders: no one voter can elect him, but all of us, collectively, can do so easily. And so, according to the Practical Possibility Criterion, Sanders needs to be included in our option list when we engage in political discourse.

What can be said for the Practical Possibility Criterion? Thus far, we have tried to motivate that criterion by analogy with the case of an individual agent like Moe. But how strong is this analogy?

Here is one way to develop the analogy. We might argue that the collective mass of voters, the electorate, comprises a *group agent*—an agent made of individual agents. This group agent has various options open to it, just as an individual agent such as Moe has various options open to him. That is, there are various things that the electorate can do, just as there are various things that Moe can do. Now, none of the participants in political discourse is identical to the electorate—each of us, in political discourse, stands outside of the electorate (despite being partly constitutive of it). In political discourse, we are engaged in a conversation about what it—the electorate—ought to do. Analogously, we can discuss what Moe ought to do without being identical to him—that is, we can discuss what he, as an individual agent, ought to do. And when we are talking about what an individual agent (such as Moe) ought to do, our task is straightforward: we should consider the list of all of the things he can do; we should try to find the best option in that list; and we should judge that he ought to choose that best option. Likewise, according to the present analogy, when we are talking about what a group agent (such as the electorate) ought to do, our task is equally straightforward: we should consider the list of all of the things that it can do; we should try to find the best option in that list; and we should judge that the group agent ought to choose that best option.

There are two primary points where the analogy can be disputed. First, one might argue that the electorate is not a group agent, and indeed isn’t an agent

at all. At a minimum, one might claim, an agent is capable of acting (as against merely exhibiting behavior) and this requires some kind of mental life, or at least a capacity to grasp and respond to reasons. But the electorate as a whole has no mental life and cannot grasp and respond to reasons (even though the individual voters who constitute the electorate do have a mental life and can grasp and respond to reasons). Second, even if it is granted that the electorate is a group agent in some important sense, one might still argue that the point of political discourse is not to reach a judgment about what that agent ought to do. Rather, political discourse is concerned with judgments about how we, as individual voters, ought to direct our support. On this account, the guiding question in political discourse is not *Whom should we elect?* but is instead *Whom should I support?*

The first of these points raises an issue that we'll unfortunately be unable to resolve here. The question of whether a collective, such as an electorate, can count as a genuine group agent is exceptionally difficult, and there are good arguments on both sides.¹¹ We are of the view that group agents exist and that an electorate can count as one, and we have defended a view along these lines elsewhere.¹² In defense of that sort of view, we might initially point out that people often talk as if these sorts of collective entities can be responsible for what they do. For instance, many observers of the Republican Party have argued that when the Republican Party chose Donald Trump as its presumptive nominee, it failed in some culpable way. It is not immediately clear how to make sense of this kind of talk unless we assume that these sorts of collective entities can count as agents (since, after all, responsibility seems to require agency; we do not hold non-agents such as comets or viruses responsible for their behavior). And so ordinary talk about collective entities in politics seems, at least at first blush, to presuppose that these entities can be agents of some sort.¹³ But this is only an initial argument for the possibility of group agents; it is not a conclusive argument for that view.

In response to the second point, we are prepared to grant that *Whom should I support?* is a guiding question political discourse, but we don't think that this causes serious problems for the analogy that we have presented. After all, supporting a candidate in political discourse paradigmatically involves defending the judgment that we (the electorate) ought to elect that candidate. To support a candidate is, as we observed in the introduction, to provide arguments in favor of that candidate; and it is hard to see what an argument in favor of a candidate would be if it were not an argument for the view that we ought to elect that candidate. And thus, we suggest, even if political discourse is most directly concerned with questions about individual support of candidates, it is still very difficult to get around the conclusion that political discourse is centrally concerned with the question *Whom should we elect?*

The major advantage of the Practical Possibility Criterion is that it can finally provide a *principled* way to decide which candidates to consider, and which candidates not to consider, in political discourse. And if one finds the analogy with Moe that we developed above compelling, then the Practical Possibility

Criterion may begin to look very tempting. As we acknowledged at the start of this section, however, the Practical Possibility Criterion faces a serious practical problem, because it would entail a vast expansion in the range of candidates that need to be considered in political discourse.

Unlike the Legality Criterion, however, the Practical Possibility Criterion—as we’ve argued—has strong theoretical virtues on its side that the alternatives lack. It may be, then, that while the view has practical difficulties, it is nonetheless the most plausible criterion available.

Conclusion

Our central question in this paper has been: When a voter is engaged in political discourse surrounding an election like the American presidential election of 2016, what is the right way to assemble an option list—a list of candidates to be considered? We’ve argued that voters need to have a principled criterion by which to separate candidates to be considered and evaluated from candidates to be set aside and ignored. Throughout the paper, we’ve considered a number of such criteria. We’ve made a preliminary case in favor of the Practical Possibility Criterion, according to which voters should be prepared to consider each and every candidate whose election is practically possible. As we’ve observed, the Practical Possibility Criterion seems to imply an extreme inclusivism about political discourse—because there seem to be a vast range of candidates whose election is practically possible.

If participants in political discourse were to abide by this sort of inclusivism, there would perhaps be some good effects. The range of practically possible candidates includes a great many individuals who could do great good if given the power of the presidency, as we pointed out in the second section. It is arguable that an inclusivism that requires us to consider those individuals would be beneficial. But this sort of inclusivism would also seem to have a number of unwelcome effects. After all, it appears that voters simply do not have the resources to consider all practically possible candidates, given that there seem to be many millions of such candidates. This seems like a powerful objection to our style of inclusivism.

There seem to be three ways that we could conceivably respond to this objection. First, we could argue that it is actually not as difficult as it may initially seem for participants in political discourse to consider millions of different candidates. (To pursue this strategy, we would need to offer a method that voters can use to meaningfully compare millions of individuals. It does not seem inconceivable that such a method could be developed; but that task lies well beyond the scope of this paper.) Second, we could argue that in fact the number of candidates who can be elected as president is in fact not very large—perhaps numbering in the dozens rather than in the millions. (To pursue this strategy, we might try to develop a fairly restrictive view about what an electorate “can” do. If it turns out, for instance, that in fact the electorate literally *cannot* elect anyone other than a small handful of candidates—say, a dozen or

so—then our inclusivism would suddenly be much more manageable.) Third, we can bite the bullet and simply accept that, on our view, a truly principled approach to political discourse will require voters to get involved in an extremely difficult and time-consuming comparison of millions of different individuals. Voters would then be faced with an unfortunate choice between either approaching political discourse in an unprincipled way, by arbitrarily selecting a manageable number of candidates to consider, or approaching political discourse in a principled but extraordinarily burdensome way.

We are not sure which, if any, of these responses to the present objection against the Practical Possibility Criterion is viable. Given this, we do not wholeheartedly endorse that criterion. We think the Practical Possibility Criterion *may* be the best such criterion available, but we are also hopeful that some superior criterion can be developed and defended.

Notes

- 1 According to the Tyndall Report (Tyndall 2015), an analysis of news coverage by ABC, CBS, and NBC by media analyst Andrew Tyndall, in the first four months of 2016 Donald Trump received approximately three times as many minutes of news coverage on the three major networks than as Hillary Clinton and Bernie Sanders.
- 2 As one illustrative example, consider that *ABC World News Tonight* offered 81 minutes of Trump coverage over the course of 2015 while only about 20 seconds was afforded to Sanders coverage. (See Boehlert, 2015 and Nichols, 2015 for this and other disparity details.)
- 3 For example, see (Sanders Campaign, 2015).
- 4 For example, see Scott (2015) and CBS News (2015).
- 5 Alternatively, Rule-Consequentialism is sometimes defined in terms of the set of rules that would have the best consequences if adopted by *most of the people* in one's society. (For a defense of this version of Rule-Consequentialism, see Brad Hooker (2013), "Rule-Consequentialism," in Hugh LaFollette and Ingmar Persson (eds.), *The Blackwell Guide to Ethical Theory* (Second Edition), Wiley-Blackwell, 238–260, esp. pp. 247–249.) But note that the Rule-Consequentialist Criterion could be easily modified to accommodate this version of Rule-Consequentialism without impacting our comments here.
- 6 For a classic rebuttal, see J.J.C. Smart (1956), "Extreme and Restricted Utilitarianism," *The Philosophical Quarterly* 6.25: 344–354.
- 7 Jonathan Schaffer (2004; 2006) has advocated for an interpretation of verbs like 'knows' and 'prefers' that is similar to what we are describing here for 'ought.' According to Schaffer, a sentence such as "Ann prefers chocolate" contains a covert, contextually determined variable that makes it elliptical for a sentence of longer form that references a contrast class (e.g., "Ann prefers chocolate *rather than strawberry*").
- 8 In discussing a contrastivist account of moral judgments, Baumann (2008) has argued that such judgments are special in that some options must always be present in the contrast classes that we use in moral reasoning. Specifically, in comparison with other normative judgments, moral judgments "...are certainly less 'pragmatic' and more 'absolute'. There is a moral value at stake here...This value 'forces' us to include certain options in the contrast class" (p. 465). We think that this point applies with just as much force to the judgments of political discourse.

- 9 Admittedly, not everyone would agree with us on this point. For instance, Sinnott-Armstrong (2006, pp. 434–452) has argued that there can be no account of what it means for a contrast class in moral epistemology to be more relevant than another. Sinnott-Armstrong writes: “I suspend belief about which contrast class is relevant and about whether any contrast class is relevant, even in a given context.”
- 10 What exactly does it mean to say that an agent such as Moe *can* do X? If a definition is needed, we think the one offered by Peter Vranas (2007) is satisfactory; Vranas writes: “I understand the claim that an agent *can* do something as the claim that the agent has both the ability and the opportunity to do the thing. The agent has the *ability* to do the thing in the sense of having the requisite skills, physical capacities, and knowledge—even if *psychologically* she is “unable” to do it..., and even if it would be *unreasonable* to expect her to do it.... The agent has the *opportunity* to do the thing in the sense of being in a situation which allows her to exercise her ability...” (170–171).
- 11 For a defense of the view that *at least some* collectives are agents, see Christian List and Philip Pettit (2011), *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford University Press).
- 12 David Killoren and Bekka Williams (2013), “Group Agency and Overdetermination” (2013), *Ethical Theory and Moral Practice* 16.2: 296–307.
- 13 For views of collective obligations according to which collectives have obligations only if they are agents, see (for example) Holly Lawford-Smith (2012), “The Feasibility of Collectives’ Actions,” *Australasian Journal of Philosophy* 90. 3: 453–467; and Stephanie Collins (2013) “Collectives’ Duties and Collectivization Duties,” *Australasian Journal of Philosophy* 91.2: 231–248. Alternatively, one could accept that groups can have obligations (and thus culpably fail to meet them) yet reject the agency criterion for obligation. For arguments that collective obligation does not require collective agency, see: Bill Wringe, “Global Obligations and the Agency Objection” (2010), *Ratio* 23.2: 217–230, and “From Global to Institutional Obligations” (2014), in Peter French and Howard Wettstein (eds.), *Midwest Studies in Philosophy* 38: 171–186; and Felix Pinkert, “What We Together Can (Be Required to) Do” (2014), also in *Midwest Studies* 38: 187–202.

References

- Baumann, Peter. (2008), “Problems for Sinnott-Armstrong’s Moral Contrastivism,” *The Philosophical Quarterly* 58: 463–470.
- Boehlert, Eric. (2015, December 11), “ABC World News Tonight has devoted less than one minute to Bernie Sanders’ Campaign this year,” *mediamatters.org*, available at <http://mediamatters.org/blog/2015/12/11/abc-world-news-tonight-has-devoted-less-than-on/207428> (accessed May 20, 2016).
- CBS News. (2015, December 24), “Bernie Sanders blames Donald Trump’s rise on the media,” *cbsnews.com*, available at www.cbsnews.com/news/bernie-sanders-says-donald-trumps-rise-is-an-indictment-of-the-media/ (accessed May 20, 2016).
- Collins, Stephanie. (2013), “Collectives’ Duties and Collectivization Duties,” *Australasian Journal of Philosophy* 91.2: 231–248.
- Hooker, Brad. (2013), “Rule-Consequentialism,” in Hugh LaFollette and Ingmar Persson (eds.), *The Blackwell Guide to Ethical Theory* (Second Edition), Malden: Wiley-Blackwell, 238–260.
- Killoren, David and Williams, Bekka. (2013), “Group Agency and Overdetermination,” *Ethical Theory and Moral Practice* 16.2: 296–307.

- Lawford-Smith, Holly. (2012), "The Feasibility of Collectives' Actions," *Australasian Journal of Philosophy* 90. 3: 453–467.
- List, Christian and Pettit, Philip. (2011), *Group Agency: The Possibility, Design, and Status of Corporate Agents*, Oxford University Press.
- Nichols, John. (2015, December 10), "The Discourse Suffers When Trump Gets 23 Times As Much Coverage as Sanders," *The Nation*, available at www.thenation.com/article/the-discourse-suffers-when-trump-gets-23-times-as-much-coverage-as-sanders/ (accessed May 20, 2016).
- Pinkert, Felix. (2014), "What We Together Can (Be Required to) Do," in Peter French and Howard Wettstein (eds.), *Midwest Studies in Philosophy* 38, 187–202.
- Sanders Campaign. (2015, December 11), "Why the Bernie Blackout on Corporate Network News?" Press release on www.berniesanders.com, available at <https://berniesanders.com/press-release/why-the-bernie-blackout-on-corporate-network-news/?version=meter+at+1&module=meter-Links&pgtype=article&contentId=&mediald=&referrer=https%3A%2F%2Fwww.google.com%2F&priority=true&action=click&contentCollection=meter-links-click> (accessed May 20, 2016).
- Schaffer, Jonathan. (2004), "From contextualism to contrastivism," *Philosophical Studies* 119.1: 73–103.
- Schaffer, Jonathan. (2006), "Contrastive Knowledge," in T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology* Volume 1, pp. 235–271.
- Scott, Eugene. (2015, December 24), "Sanders: Trump is 'bombastic' so he can get media coverage," *cnn.com*, available at www.cnn.com/2015/12/24/politics/bernie-sanders-donald-trump-new-day/ (accessed May 20, 2016).
- Sinnott-Armstrong, W. (2006), *Moral Skepticisms*, pp. 434–452, Oxford: Oxford University Press.
- Smart, J.J.C. (1956), "Extreme and Restricted Utilitarianism," *The Philosophical Quarterly* 6.25: 344–354.
- Tyndall, Andrew. (2015), "Let the Trump Circus Continue," blog post - tyndallreport.com, 2016, available at <http://tyndallreport.com/> (accessed May 20, 2015).
- Vranas, Peter. (2007), "I Ought, Therefore I Can," *Philosophical Studies* 137: 167–216.
- Wringe, Bill. (2010), "Global Obligations and the Agency Objection," *Ratio* 23.2: 217–230.
- Wringe, Bill. (2014), "From Global to Institutional Obligations," in Peter French and Howard Wettstein (eds.), *Midwest Studies in Philosophy* 38, 171–186.

Taylor & Francis
Not for Distribution