

A Too Intimate Internet: What is Wrong with Precise Audience Selection?

*Thomas Mitchell**

Oxford Uehiro Centre for Practical Ethics, University of Oxford.

*Corresponding author. Email: thomas.mitchell@philosophy.ox.ac.uk

Keywords: politics; social media; influence; internet; audience selection

Abstract

It is commonly recognized that the modern capacity for mass online communication carries various dangers: fake news, rampant conspiracy theories, trolling, and so forth. It is less commonly realized that moral problems remain when the contents of online communications are completely innocuous. This article discusses one of the noteworthy features of modern digital technology, the fact that it is possible to precisely target specific audiences, and argues that this can make mass communications such as advertising and political campaigns morally problematic. What is more, this holds even if the communicator is using only rational persuasion. In being selective about who sees which arguments, one becomes liable to mislead the audience despite sticking to honest, evidence-based, rational argumentation.

The Internet allows communicators to reach a far larger audience than ever before, and also to select with far greater precision who will be in that audience. This article argues that this raises moral problems, even if the messages being communicated are not themselves problematic. It uses the example of an election campaign based purely on rational argument to demonstrate that precise audience selection can wrongfully mislead the electorate, despite the candidate sticking to honest, evidence-based, rational argumentation.

There are two features of mass online communication that make it both very useful and potentially dangerous. The first is the sheer immensity of its reach; it is possible to get an idea to more people than ever before through Twitter, YouTube, Facebook, and other popular platforms. The second is that savvy communicators

can be precise about who sees those ideas. Using the data gathered about potential customers, for instance, adverts are shown to those most likely to become actual customers. Social media sites, too, will show their users the content to which they are most likely to respond.

Some of the dangers that come with these features are well known. Fake news, conspiracy theories, and harmful ideologies spread rapidly in a system with such an immense reach combined with a propensity to pass their messages to those most susceptible to them. The Internet is a powerful tool that can lead to serious harm when dangerous views are fed into it.

However, I would like to draw attention to how the ability to precisely select your audience can be morally problematic even if your message is not itself dangerous. This is because it violates one of the conversational maxims proposed by



Paul Grice, a philosopher who worked on the philosophy of language in Oxford between 1939 and 1967. These are not themselves moral norms, but violating them can mislead your audience, which is morally problematic.

Consider the following case. Peter is a politician seeking election to public office, so he needs to persuade enough people to vote for him. He likes to think of himself as a decent, honest person, so he resolves to use no underhand tactics. He will not trick, lie to, or make false promises to the public; his campaign will be based solely on rational argument. He therefore marshals a series of well-reasoned, evidence-based arguments supporting the various policies that he proposes. Naturally, he realizes that not everyone will vote for him, no matter how well he puts his case. Furthermore, even for those who can be convinced, some of his arguments will be more persuasive than others. Some voters may have reasonable objections; others may have biases

that discline them from being convinced; and still others may simply fail to understand his reasoning. Given these conditions, what should he do to maximize his chances of getting elected, while sticking strictly to rational persuasion?

An obvious thought might be to give all his arguments to everybody. That way, all those who would be convinced by any (combination) of his arguments would hear what is required to make them vote for him. Some would remain unconvinced, but these people would never have voted for him anyway. In this way, he would convince everyone he possibly could. However, this is unlikely to work. The strategy requires that his audience pay close attention to all his arguments and carefully weigh them up against what his opponents say. But unfortunately, people are not perfectly rational, often have short attention spans, and generally have limited cognitive resources. Peter wants to

make things as clear and as easy as possible for his potential voters.

‘... precise audience selection can wrongfully mislead the electorate, despite the candidate sticking to honest, evidence-based, rational argumentation’

A better strategy would therefore be for Peter to tailor his audience to his various arguments. He identifies those who might be convinced and presents them with only the arguments that are most likely to convince them, omitting any arguments that are unlikely to have the desired effect. Note that he is not changing his arguments, which remain as they were, based on reason and evidence – he is merely being selective about who is being shown which arguments. Furthermore, he does not present his case at all to those who will never be persuaded; that would be a waste of time and resources and would risk galvanizing the opposition, who would start formulating objections to the arguments they are shown. This strategy would be virtually impossible to implement using traditional campaigning methods. But using the immense reach and precise targeting of the Internet, it would be feasible for Peter to deploy his arguments just where they would be most effective. So, he launches an online campaign, first gathering data on as much of the electorate as possible, then using algorithms to ensure that individuals see posts, adverts, and articles that support those arguments that each person is most likely to find persuasive.

This may be an effective method for getting elected – and, in general, getting a population to do what you want. Furthermore, it may seem

morally innocuous. Peter has relied solely on rational persuasion; there have been no damaging false rumours spread about opponents, no lies posted on the sides of buses, no endorsement of ‘alternative facts’, and no appeals to voters’ prejudices. But there is nonetheless a problem with the strategy of precisely selecting an audience even for his reasonable and harmless messages.

Paul Grice, in his essay ‘Logic and Conversation’ (H. P. Grice, ‘Logic and Conversation’, in *Syntax and Semantics*, vol. 3: *Speech Acts* (New York: Academic Press, 1975), 41–58; henceforth ‘Grice 1975’), writes of what he calls the Cooperative Principle: ‘Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged’ (Grice 1975: 45). He draws this from the observation that talk exchanges, even when they are one-sided, are to some extent cooperative practices, with the participants sharing some goal to which they are expected to contribute. The goal may be only vaguely defined and might change over time, but all participants should have some sense of it. For instance, if we were to have a conversation about the relative merits of two rival political parties, it would not be appropriate for me suddenly to start talking about the weather, or to make responses that ignore what you have just said, or deliberately not to mention some relevant aspect of one party because I think it will not support whatever point I am trying to make. Such behaviour would be failing to make a contribution required by the accepted purpose of our conversation. As a cooperative practice, each participant is expected to play their part.

Following this general principle, Grice goes on to suggest a series of more specific maxims that govern conversation, which he divides into four categories: Quantity, Quality, Relation, and Manner. For example, under the category of Quality fall maxims requiring that you should only assert that for which you have adequate evidence and that you should not assert that which you believe false. Under the category of Relation falls the single maxim, ‘Be relevant’. Avoiding unnecessary ambiguities and being brief are both included under the category of Manner

(Grice 1975: 45–6). This is not a full list of the maxims Grice proposes, only a representative sample. Grice himself allows that there may be further maxims beyond those he highlights (1975: 46). The point is that there are certain, typically unspoken, conversational norms which help us to understand one another. Infringement can result in the audience being misled or feeling confused. There is normally a background assumption on the part of the listener or reader that these maxims will be observed unless otherwise indicated. Suppose, in our political discussion, I were to start commenting on local agricultural practices without signalling a change in topic. You might reasonably assume that I will shortly show how it is related and that the tangent was necessary for making a certain relevant point, for instance that one political party's policies would make a positive impact on farming communities. That is, you would assume that I was adhering to the maxim of relevance. If I was not, and I really was just talking about something completely unrelated for my own amusement, you would probably be left feeling rather nonplussed.

Grice distinguishes four ways in which someone might fail to fulfil a maxim. First, they might *violate* it, which is done surreptitiously; the audience is generally left believing that the maxim is still being followed. Second, they might *opt out* of a maxim, by making it clear that they are not following it on this occasion. For instance, beginning a sentence 'Anyway, the other thing I wanted to mention ...' indicates that what you are about to say will not be relevant to what has been said previously. Third, there may be a *clash*, which occurs when the speaker cannot fully satisfy one maxim without failing to satisfy another – you must then make a judgement as to which should take priority in the circumstances (some maxims might generally be more important than others). Fourth, they might *flout* a maxim, which is the unmarked but very blatant failure to fulfil it. For instance, when you are being ironic, you flout the maxim of not asserting what you believe to be false, but the audience will usually understand what you are doing (Grice 1975: 49). For our purposes, the most relevant

kind of failure is violation, which is the most apt to mislead.

The maxim relevant here is the first maxim of Quantity, which is as follows: 'Make your contribution as informative as is required (for the current purposes of the exchange)' (Grice 1975: 45). Essentially, you should not leave out bits of information relevant to the purpose of the communication. For example, if, driving to the station, someone stops to ask me for directions, and I tell them that it is at the end of the next road but fail to mention that that road is a one-way street and so does not provide access to the station for vehicles, then I have not been as informative as the exchange requires, though I have been truthful. The purpose of the communication was for them to find out how to get to the station and they would assume that I am adhering to the maxim of informativeness. They would therefore be misled into thinking that they could access the station that way since, they assume, I would have mentioned it if there had been anything preventing them from doing so.

I believe that, in my election campaign example above, Peter fails to fulfil this maxim. Specifically, he *violates* it, leaving him liable to mislead. When he displays his messages to those whom he hopes to convince, he fails to present them with factors relevant to each voter's decision. Since all his arguments are evidence-based, rational and on the relevant topic, they constitute information that the communication requires. So, by withholding some of his arguments on the basis that they are unlikely to be effective, he renders their natural assumption that he has been as informative as is required false. Suppose that a certain group of voters is dead set against one of Peter's views, but is susceptible to being convinced of another, and so they are presented with arguments for the latter and not the former. They assume that Peter would have mentioned the former if it was an issue he cared about. The purpose of the communication (in this case, an online post or advertisement) is to provide the recipient with reasons to vote for him, but he has left out some of what he thinks of as important reasons.

‘There is ... no shortage of harmful content online. What makes cases like this different is that the content is itself quite harmless and even seems an exemplar of innocuous influence.’

It may be objected that the maxim of informativeness is here clashing with another Gricean maxim mentioned above: that of brevity. In order to be brief, a speaker must restrict the amount of information given. Grice himself acknowledges that sometimes his maxims will clash with one another (Grice 1975: 49) and so it is not always clear what they prescribe. In this case, it seems that the Gricean maxims do not straightforwardly preclude Peter leaving out many of his arguments when addressing each part of the electorate, contrary to what has been argued. He has to leave some out, or the messages would get far too long. However, although maxims can come into tension, this is not occurring here. This is because Peter’s communicative activities are no briefer than if he had decided to lay out all his arguments. As part of an online campaign, a message is not merely shown to a user once, but is repeated, often in various forms. The purpose of adopting this strategy, after all, is to make things very clear to the audience and repetition is a simple means of doing this. If the targeting is successful, posts favouring the arguments Peter wants an individual to see will appear frequently on their newsfeeds, articles supporting those arguments will often be recommended to them, adverts preferring those arguments will be shown to them, and so forth. If he were to give all his arguments to everybody, his audience would be treated to

just as many campaign messages. The difference would merely be that those messages would not be specifically tailored to them; he would be communicating with no greater brevity.

So, precise targeting violates a maxim of conversation and thus fails to adhere to the Cooperative Principle. But what is so *immoral* about doing so? Put simply, the audience is misled into thinking that Peter’s online campaign has presented a more or less complete picture of the case he is arguing, when of course it has not. They are led to believe that he is adhering to the maxim of informativeness, when in fact he has violated it. This is more serious than it may initially seem. In misleading them into thinking that they have all his arguments, Peter also misleads them about the kind of candidate he is. He has shown only a part of what he stands for, while leading the electorate to believe that they have seen the whole. It is a subtle kind of duplicity; he has denied them the chance to consider all that he would do and be, were he elected to office, without ever lying or even departing from honest, evidence-based, rational argumentation. There may be worse ways of trying to gain power, but surely the ideal in a democracy is for candidates to be evaluated on the basis of all their merits and flaws as judged by the electorate. This Peter prevents without any active concealment and without anyone realizing. It is one thing for an electorate to consider a candidate who is believed to be hiding something and to weigh up whether it is worth voting for them given that there is an unknown element to them. It is another for them to consider a candidate who is apparently open about what they stand for, but in fact is only partially so; who is not merely two-faced, but has a multitude of faces, showing a different part of themselves to different sets of voters.

We have here considered the example of getting voted into office, but what has been argued above will apply to any scenario involving trying to get a population to do or believe something. Consider companies advertising their products, charities finding potential donors, corporations seeking investors, institutions spreading ideas, ... The list goes on. Many kinds of agents have an interest in getting groups and individuals to act and think in certain ways. It is unlikely that

Thomas Mitchell

all will use rational persuasion as their favoured method – and almost certainly not all the time. But where they do rationally persuade, it may commonly be thought that their actions are innocuous. This idea ought to be challenged.

There is, as mentioned at the outset, no shortage of harmful content online. What makes cases like this different is that the content is itself quite harmless and even seems an exemplar of innocuous influence. It is rational argument,

the apparent antithesis of problematic online communication in the modern day. But I hope to have shown that we do not need hateful messages, virtual shouting matches, and pernicious misinformation to make online communication morally problematic. The nature of the Internet itself, its immense reach and facility for precision, can render even the most seemingly unproblematic communications unethical.

Thomas Mitchell

Thomas Mitchell is Research Fellow in the Ethics of Persuasive Digital Technologies at the Oxford Uehiro Centre for Practical Ethics, University of Oxford. He thanks the European Research Council (grant number 819757) for their funding.

Cite this article: Mitchell T (2024). A Too Intimate Internet: What is Wrong with Precise Audience Selection? *Think* 23, 37–42. <https://doi.org/10.1017/S1477175623000337>

© The Author(s), 2023. Published by Cambridge University Press on behalf of The Royal Institute of Philosophy. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is unaltered and is properly cited. The written permission of Cambridge University Press must be obtained for commercial re-use or in order to create a derivative work.