

In today's political climate moral outrage about (alleged) propaganda and manipulation of public opinion dominate our discourse. Charges of manipulative information provision have arguably become the most widely used tool to discredit one's political opponent. Of course, one reason for why such charges have become so prominent is that the way we consume information through online media has made us more vulnerable than ever to such manipulation. Take a recent [story](#) published by The Guardian, which describes the strategy of information dissemination allegedly used by the British 'Leave Campaign':

"The strategy involved harvesting data from people's Facebook and other social media profiles and then using machine learning to 'spread' through their networks. Wigmore admitted the technology and the level of information it gathered from people was 'creepy'. He said the campaign used this information, combined with artificial intelligence, to decide who to target with highly individualised advertisements and had built a database of more than a million people."

This might not just strike you as "creepy" but as simply unethical just as it did one commentator cited in the article who called these tactics "extremely disturbing and quite sinister". Here, I want to investigate where this intuition comes from.

### **We can distinguish different ways one can go about changing someone else's mind:**

(1) by providing arguments,

(2) by appealing to epistemic trust (i.e. claims to expertise and epistemic authority),

and

(3) by exploiting how our cognitive mechanisms have evolved/developed to process information in general. For example, by directing attention in a certain way or by creating/triggering semantic associations between different contents (e.g. pairing an image of a happy family with a specific brand of car or by using terms such as 'flood' to describe migratory movements of people).

While in (1) and (2) one interacts with one's audience's cognitive machinery charged with evaluating critically what information to endorse (so-called epistemic vigilance mechanisms) in (3) one bypasses this machinery by exploiting the way that our cognitive system is designed to process information in general. For the purposes of this entry, I want to call (1) conviction, (2) persuasion and (3) manipulation. Put in these terms, the question I want to address here is why manipulation in contrast to conviction and persuasion often strikes us as ethically wrong.

One thing that bears mentioning right away is that the distinction between (1), (2) and (3) is not simply one between verbally communicated information and non-verbally communicated information. A political speech, say, can be manipulative while still employing arguments and appeals to epistemic trust explicitly. One distinction that is of importance here, however, is that of ostensively vs. non-ostensively communicated information. Manipulation, by definition, can never be ostensive since this would imply that the speaker intends her audience to know about her intention to manipulate. Manipulation is thus by default non-ostensive.

### **Why and how is manipulation ethically problematic?**

Nonetheless, when I ask the question why manipulation should be ethically problematic, I am talking of course about *intentional* manipulation. In interacting with others, we always rely on their general

cognitive capacities and the tendencies through which these capacities have evolved to process information. If we took 'manipulation' in this sense to be unethical, much of our interactions with others would be, too.

So why do we intuitively take intentional manipulation (i.e. non-ostensive forms of 'mind changing') to be ethically problematic? I want to suggest that this has to do with the kind of commitments that underlie communicative interaction. When I make an ostensive (i.e. explicit) claim, I thereby not only convey information but also take responsibility for its truth: I make myself *accountable*.

This analysis applies to conviction and persuasion in different ways. In conviction/argumentation, one provides an argument, which is intended to trigger specific inferences in one's audience based on what that audience already believes. This minimizes the extent to which the audience has to trust the speaker. Nonetheless argumentation relies on assertions to some extent, which always force the speaker to commit herself to the truth of whatever she is stating.

Persuasion, on the other hand, works only because of the commitments involved: whenever a speaker claims epistemic authority of some kind, this claim is only effective to the extent to which she makes herself accountable for the truth of her utterances. Crucially, it is rare that conviction and persuasion occur entirely apart from each other: we commonly rely on epistemic authority in our arguments.

Manipulation differs exactly in that there is neither an explicit claim to be evaluated nor does the manipulator incur any commitment towards her target. When I create an implicit association between immigration and terrorism, for example, by mentioning both words in the same sentence, I have thereby not committed to any such association indeed existing. In this sense, manipulation crucially differs from lying since a lie always consists in an explicit claim, which commits the speaker to its truth. Manipulation intentionally shortcuts those cognitive mechanisms designed to evaluate and keep track of such commitments. The reason we take manipulation to be unethical then is because it comes with some of the benefits of persuasion and conviction without their respective costs.