

# Contents

Acknowledgments	vii
Introduction	1
Chapter 1: Privacy	7
Chapter 2: The Attention Economy	37
Chapter 3: Nudging	57
Chapter 4: Echo Chambers and Polarization	83
Chapter 5: Misinformation	107
Chapter 6: Cancel Culture: Online Shaming and Caring	131
Chapter 7: Friendship	153
Chapter 8: Quitting	173
Conclusion	197
Bibliography	199
Index	209

# Introduction

Philosophers have developed the field of applied ethics to apply moral concepts to difficult ethical issues that arise in particular areas of our existence. There are environmental ethics, business ethics, medical ethics, and so on. The ethics of social media is one such part of applied ethics. Many of the ethical issues that this book discusses are also treated by philosophers working in the field of information ethics, which gets its name from the fact that it is concerned with information technology.

It is appropriate to talk about the ethics of social media, rather than information ethics, because doing so highlights more precisely what we are talking about—ethical issues that have arisen or become more prominent on account of social-media sites and apps, such as Facebook, Instagram, TikTok, X, LinkedIn, and many others. There is also the fact that many of the topics considered in this book, such as the value of online friendships, really do not have much to do with what people ordinarily mean when they talk about information technologies, such as search engines.

The largest obstacle to the ethics of social media is the set of thought-terminating clichés that surround the subject in popular media and discourse. When we say “thought-terminating cliché,” we have in mind those brief slogans that we usually roll out to end a conversation without engaging with much of the substance of what we are replying to. Some familiar thought-terminating clichés are “It is what it is,” “God works in mysterious ways,” and “Boys will be boys.” When we talk about the ethical issues generated by Facebook and Instagram, we might hear that such technologies are just tools or that they can be used for good or bad.

These are particularly pernicious clichés because there does seem to be something true about them. The latter, for instance, cannot possibly be *wrong*. After all, social-media apps *can* be used for good or for evil. When I teach the ethics of social media to students, I clarify on the first day of class that the course is not designed to relentlessly drag social-media apps through the mud. On the contrary, some of the chapters of this book, just like many of the lectures in my class, highlight the good things about social-media platforms: for instance, friendships on these platforms have real value, and they provide a community to people who are vulnerable and might not be able to meet like-minded people offline. All we

need to do to establish that social-media apps can do good is point to worthy crowdfunding opportunities. By the same token, this book's Table of Contents provides a list of many harmful things associated with social media, including echo chambers, misinformation, and violations of our privacy.

What makes it a thought-terminating cliché to say that social-media platforms can be used for good or for bad is that it does not really settle any ethical issues. It is equivalent to shrugging and saying "It is what it is." When people debate nuclear disarmament, someone might similarly say that nuclear weapons can be used for good or for bad, but even if we grant that point, we still want to know the answer to specific questions, such as how likely it is that they will be used well, what we should do about them in light of their potential for good and for bad, what using them well entails, and what risks we run by disarming or not disarming. Moreover, we should say *something* about the particular goods and evils that these things *do* cause. Similarly, if we are considering the use of fossil fuels, we would not be satisfied by knowing all the good things that come from having an easily available fuel source. We would want to know about all the bad things that arise from fossil fuels, and then whether the good outweighs the bad.

At this point, we run into the other thought-terminating cliché, which is the claim that social-media apps are merely tools. Since they can be used for good or for evil, how we use them is up to *us*. Of course, there is, yet again, something true about this. I can use TikTok to bully someone or I can use TikTok to advertise an important crowdfunding campaign. That is up to me.

The problem is that social-media technologies are *not* tools in the same way that hammers and saws are. This is an important philosophical point that will not really become clear until the end of Chapter 3, but I shall outline the basics now. Some tools, such as hammers, are inert and passive. They are waiting for us to use them as we see fit. In contrast, social-media platforms *act on us*. Specifically, they nudge us. This means that they are designed to promote the likelihood of some behaviors and choices on our part. Most often, this means that they are designed to get us to spend more time on their platforms and to get us to reveal more information about ourselves.

Social-media platforms simply *cannot* be neutral. They are built with design values in mind, and countless millions of dollars go into building every aspect of apps to promote the likelihood of certain behaviors without coercing you. Everything is designed with the values of the designers as guides, even the color of the

notification symbol. (Red, it turns out, is much more effective at getting you to click than blue is.)

Throughout this book, we will discuss more of these design values. For instance, one of them, called *consumer sovereignty*, values prioritizing content that the designers believe you would be most likely to click on and consume. (They infer from your clicking behaviors that it is content that you *want* to see.) This rather innocent-seeming value gives rise to communities in which users are shown only content that they already agree with. This creates a bias: users' beliefs grow more intense because they are shown only corroborating points of view and are never acquainted with the evidence against them.

If we keep our eyes on these considerations—which are mere previews of deeper discussions that unfold on the following pages of this book—then it is appropriate for us to ask how social-media technologies affect *us*. Once we ask this question, we can see that it is mistaken to say that social-media platforms are merely tools. In certain senses, they use us. (Chapter 2 will even bring to light the way that we, and our attention specifically, are the commodities being traded in a transaction between a platform and some advertiser.) Even putting aside philosophical discussions of nudges and design values, each of us can pause and reflect on the way that spending time on social-media platforms changes us: it might make us feel more anxious and lonely, or it might make us feel happy and more connected to other people. Either way, we are not describing tools as we might be when we are talking about hammers or saws.

There is an insight from ancient philosophers that we might profitably learn from today, namely, that the things that we use are active in a way that we downplay when we call them mere tools. Plato, the ancient Athenian philosopher, recognized this point when he talked about how pleasures affect our conception of what is real and valuable. They lure us into changing our beliefs.<sup>1</sup> Even more relevant to the present discussion is the ancient Confucian philosopher Xunzi, who argued that ritual and music can rework human nature itself. For Xunzi, humans are by nature bad, but we can change our nature through deliberate effort and hard work. The path to changing human nature is marked by sages, who invented practices and certain kinds of music to help bring about the desired change.<sup>2</sup> Yet, in the minds of many ordinary people today, who and what we are does not change when we pick up any so-called tool. According to a view that is common today, we might be using social-media apps badly on Monday, but we could just as easily be using them well on Tuesday if only we have the right intention and mindset.

This is overlooking the way that social-media apps are designed to get you in one mindset and keep you there indefinitely.

When I talk about this common view, I have in mind the popular discussions about social media in the news and, perhaps unsurprisingly, on social-media platforms themselves. I do not mean to attribute this common view to philosophers. The burgeoning field of the ethics of social media is rich and exciting, as I am sure that this book will confirm.

Each chapter of this book is designed to facilitate discussion around a different topic within the ethics of social media. I consider multiple points of view, and I have put together an annotated bibliography at the end of each chapter, which is supposed to simultaneously provide a list of further readings and substantiate the claims made in the chapter. I have also provided sets of cases to think through at the end of each chapter. I have chosen real cases because applied ethics is at its most valuable when thinkers work with actual cases. They give us an opportunity to see what philosophy can do when it turns its attention to the day-to-day problems that affect us, and I think that the result of wrestling with these problems will be a more acute sensitivity to what social-media technologies do, both for us and against us.

I shall conclude this introduction with a brief mention of my own experience with the ethics of social media. As a graduate student working on ancient philosophy at the University of Toronto, I became aware of discussions in the media about echo chambers and misinformation. I felt that these discussions were important but that they were not particularly well conducted. Part of the problem with the conversations was certainly due to the fact that the subject was (and still is) new: these technologies are new, and we have not yet developed the necessary social antibodies to protect against some of their harms. (“Social antibody” is a term, discussed in Chapter 2, that refers to the way that society spontaneously develops rules and norms to help it deal with threats, as our own immune systems do with the threats that they encounter.) The problem was also partly due to the fact that philosophers had not yet had their contributions to these discussions fully appreciated. My hope is that by introducing to students and other interested readers the ethics of social media, more people will benefit from the richness of philosophical analyses of these problems.

At the same time as I was noticing these discussions in the media, I attended a talk at a meeting of the American Philosophical Association by Professor Kathryn Norlock, who was discussing whether it is possible to consent to the terms and

conditions of use of what was then called Twitter. Her talk convinced me that philosophers have a lot to contribute to these discussions. Later that year, I designed and taught my course on the ethics of social media at the University of Toronto for the first time. Now I have moved to Alma College, where I still teach the course every year, and I find that each new cohort of students visits the various ethical problems in new ways. This book represents the culmination of years of work teaching students: presenting the material concisely and lucidly to them, choosing topics that speak to what they are experiencing online, and compiling cases that are both philosophically relevant and thought-provoking.

### Notes

1. See Plato's *Phaedo* 83b–c.
2. See *Xunzi*, chaps. 19, 20, and 23.