

Why It's OK to Speak Your Mind

Political protests, debates on college campuses, and social media tirades make it seem like everyone is speaking their minds today. Surveys, however, reveal that many people increasingly feel like they're walking on eggshells when communicating in public. Speaking your mind can risk relationships and professional opportunities. It can alienate friends and anger colleagues. Isn't it smarter to just put your head down and keep quiet about controversial topics?

In this book, Hrishikesh Joshi offers a novel defense of speaking your mind. He explains that because we are social creatures, we never truly think alone. What we know depends on what our community knows. And by bringing our unique perspectives to bear upon public discourse, we enhance our collective ability to reach the truth on a variety of important matters.

Speaking your mind is also important for your own sake. It is essential for developing your own thinking. And it's a core aspect of being intellectually courageous and independent. Joshi argues that such independence is a crucial part of a well-lived life.

The book draws from Aristotle, John Stuart Mill, Friedrich Nietzsche, Bertrand Russell, and a range of contemporary thinkers to argue that it's OK to speak your mind.

Hrishikesh Joshi is Assistant Professor at Bowling Green State University, and works on moral and political philosophy. He completed his Ph.D. at Princeton University.

First published 2021
by Routledge
52 Vanderbilt Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

*Routledge is an imprint of the Taylor & Francis Group, an informa
business*

© 2021 Taylor & Francis

The right of Hrishikesh Joshi to be identified as author of this work
has been asserted by him in accordance with sections 77 and 78 of the
Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or
reproduced or utilised in any form or by any electronic, mechanical,
or other means, now known or hereafter invented, including
photocopying and recording, or in any information storage or retrieval
system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks
or registered trademarks, and are used only for identification and
explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data
A catalog record for this title has been requested

ISBN: 978-0-367-14171-4 (hbk)

ISBN: 978-0-367-14172-1 (pbk)

ISBN: 978-1-003-14575-2 (ebk)

Typeset in Joanna MT Pro and DIN pro
by Apex CoVantage, LLC

Taylor & Francis

Not for distribution

Let us speak of this, you who are wisest, even if it be bad.
Silence is worse; all truths that are kept silent become
poisonous.

—Friedrich Nietzsche, *Thus Spoke Zarathustra*

Speaking your mind can have consequences. We constantly face pressures to conform to the opinions—or at least *perceived* opinions—of our peers, friends, and employers. Dissenting from these views can have real costs. Often, the best way to advance in one’s career, for example, will involve not “rocking the boat” too much. Saying what you really believe might come at the cost of lost job opportunities, promotions—and in more extreme, but nowadays not uncommon, cases, getting fired.

What’s more, these sorts of pressures exist in their strongest form precisely within those professions and institutions which are primarily responsible for producing new ideas, maintaining the stock of knowledge, and shaping culture. Dissent from the zeitgeist of the cultural elite is less likely to be the cause of lost employment for truckers, plumbers, or mining engineers. It’s much more of an issue for writers, journalists, academics, and artists. But if such pressures exist most strongly within these professions, and if such pressures can cause group-think and blind spots, then it seems particularly important to address them because of the power such professions have in determining public opinion.

Now, of course, legal sanctions for the expression of ideas are rare within modern democracies. The United States has protections built in to the Constitution itself, in the form of the First Amendment. You cannot go to jail for merely voicing

an opinion—no matter how heterodox or repulsive. But this protection is nowhere near enough to safeguard our collective ability to evade the perils of conformity. John Stuart Mill, perhaps the most preeminent historical defender of free expression, was acutely aware of this issue. He wrote:

In respect to all persons but those whose pecuniary circumstances make them independent of the goodwill of other people, opinion, on this subject, is as efficacious as law; men might as well be imprisoned, as excluded from the means of earning their bread.¹

Bertrand Russell, one of the premier analytic philosophers of the 20th century and a defender of liberalism, echoed this idea decades later:

Legal penalties are, however, in the modern world, the least of the obstacles to freedom of thoughts. The two great obstacles are economic penalties and distortion of evidence. It is clear that thought is not free if the profession of certain opinions makes it impossible to earn a living. It is clear also that thought is not free if all the arguments on one side of a controversy are perpetually presented as attractively as possible, while the arguments on the other side can only be discovered by diligent search.²

Unless you are independently wealthy, the costs of self-expression can be very real despite the legal protections. Even if you are wealthy, there are substantial risks associated with saying what you think about contentious topics. For one, there can be reputational costs—from friends, neighbors, or colleagues. And bad reputations, even if unfounded, have

tendencies to stick. In addition, if you stray too far from sanctioned opinion, you may eventually lose access to the very means of expressing your opinions to the broader public: op-eds or TV appearances, for example. That said, the democratization of information dissemination via the internet has in some ways reduced the power of gatekeeper institutions in this regard—though not entirely, because the internet has gatekeeper institutions of its own.

Given that speaking your mind has costs, what should you do?

This question is increasingly relevant for the many who feel that they often can't say what they really think. This is the question I want to explore. What this book is *not* trying to do is to defend, in the first instance, the ideal of free speech—that is, the notion that society and its institutions should be open to dissenting opinion. If you want a defense of that idea, you can do no better than to read John Stuart Mill's *On Liberty*, particularly Chapter 2.

Rather, what I want to do here is to defend the idea that we often have a duty to speak our minds, even in the face of the sorts of costs mentioned above. And furthermore, a good life involves the cultivation of intellectual independence, which we cannot achieve without the outward expression of our ideas. A life of intellectual conformity and status seeking, I will argue, leaves something important to be desired.

Below is a quick birds-eye view of the main arguments of this book.

SYNOPSIS

Speaking our minds, against social pressure not to do so, can often help improve the condition of what I will call the

“epistemic commons”—that is, the stock of evidence, ideas, and perspectives that are alive for a given community. This is because our knowledge is essentially *distributed*, due to the division of cognitive labor. We are deeply reliant on others for what we take ourselves to know. Few of us actually know how a zipper works, for example, but we often take ourselves to know this because we can easily access the relevant facts. In this way, we do not and cannot think alone.

This is in many ways an indispensable blessing: imagine having to figure out everything by yourself! Yet it is also a curse. As the ocean is vulnerable to overfishing and the atmosphere is vulnerable to pollution, the epistemic commons is vulnerable to *social pressure*. Social pressure can distort our picture of the world, often dangerously. And if we have a distorted view of what the world is like, then the actions we take can be counterproductive—even if our intentions are good.

The handling of the Chernobyl nuclear disaster in the former Soviet Union is a prime example of evidence distortion leading to catastrophe. Much of the disaster could have been averted had the authorities and the general public a better idea of what was going on. But because evidence was suppressed at crucial junctures, the response to the disaster allowed for many more deaths than it had to.

While the Chernobyl disaster occurred under an authoritarian government, democracies are not immune to the underlying phenomenon. Social pressure can cause evidence to pile up on one side of a debate or issue, while evidence on the other side is systematically screened out. A person who looks at the evidence as it is presented, then, will form a warped view of the matter, even if she rationally evaluates the evidence at hand. I will suggest that whenever there exists social pressure to conceal evidence on one side of a topic,

we should suspect that a (possibly dangerous) blind spot is lurking somewhere, due to this mechanism. The catch, and it is a big catch, is that the blind spot will not be recognizable by people who simply look at the evidence as it is presented. They will be unaware that their view of the world is distorted.

We should be especially mindful of social pressure in those institutions most closely associated with knowledge production and dissemination: particularly universities, but also thinktanks, newspapers, and so on. In contexts where there is social pressure within a community not to give evidence for certain conclusions, the output of that community cannot be taken seriously—or at least it must be taken with a generous serving of salt. Nonetheless, the danger is for outsiders to take this output at face value: through no fault of their own, they might be unaware of the social pressures within the community. This is our modern epistemic predicament, though it has received little recent philosophical attention.

Given these dynamics, I will suggest there is a duty to reveal our evidence against the social pressure, so long as the costs are not too high. In this sense, you have a duty to speak your mind against social pressure, at least some of the time. The duty is imperfect in the sense that we can pick our battles. We don't have to speak our minds in every single context; that would be too onerous and generally unwise. Nonetheless, there is a duty to speak up against social pressure because this constitutes doing your part to protect a common resource—namely, the epistemic commons. Someone who never does this is a free rider. He benefits from the work of others but never does his part to contribute—much like the roommate who never does dishes or the dog walker who never picks up after the dog.

If you work in one of the institutions of knowledge production mentioned earlier, one way you can significantly benefit humanity's epistemic position is by pursuing heterodox projects. The marginal value of this type of work is very high. The first article defending X or developing a new way of seeing things is vastly more valuable than the 1000th article defending not-X or making a small move in some arcane debate. Those with the protection of tenure are in an especially good position to do this and to help others who are pursuing such work.

Now, the structure of the free riding problem identified here creates a worry that many readers will share. Given the many actors involved in the maintenance of the epistemic commons (think of all the people who drive science or a particular culture), we might wonder if our individual actions make a real difference. Climate change seems to be a case like this. What difference will it make if you stop eating hamburgers? Probably not much, given the scope of the problem. But isn't speaking your mind like this? What difference can one person make?

Quite a lot, it turns out. Even *one* dissenter can have a huge psychological impact on others. This is borne out in several psychological studies, most notably the famous Solomon Asch experiments. One person calling it like they see it can relieve the fear of isolation that other potential dissenters might be experiencing. That person can be you. This is also a lesson we can draw from Hans Christian Andersen's fairy tale "The Emperor's New Clothes"—when the child points out that the emperor has no clothes, everyone suddenly musters the courage to say it themselves. In this way, a lone voice can burst the bubble of what social scientists call "pluralistic ignorance,"

which is a situation wherein most people think something but are unaware that they are in the majority.

Speaking your mind for the sake of the *common good* is the subject matter of Chapters 1–3. Chapters 4 and 5 contend that you should speak your mind *for your own sake*. Speaking your mind is an essential component of the good life. Or so I will argue.

The question of what makes for a good life has been a core preoccupation of philosophers for as long as there has been philosophy. When we look at a life as a whole, under what conditions might we say: “this here was a life well-lived”? Two natural answers present themselves immediately: pleasure and social status. According to the first option, a life goes well for the person who lives it to the extent that it contains lots of pleasure (which encompasses not only things like gustatory and sexual pleasure, but also feelings of contentment, satisfaction, and the like) and little pain (physical pain, but also frustration, depression, etc.) The second option is that a good life contains lots of social status: thus, kings, presidents, deans, Nobel Prize winners, musical celebrities, etc., are living the best lives.

The great ancient Greek philosophers considered these two answers and found them wanting. For Aristotle, what makes something—anything—good is whether it fulfills its characteristic work (*ergon*) well. Therefore, if we want to know what a good human life is we need to think about what is the characteristic work of humans. In other words, what is distinctive about humans *qua* humans? Aristotle thought it is our capacity to *reason*. Hence, he thought the good human life is one that exhibits the proper development and exercise of reason.

Recent work in philosophy and psychology suggests we simply cannot reason well in isolation. In order to reason well, we must find interlocutors with whom we can go back and forth. Reasoning is an *essentially social* activity. But if that's right then reasoning well involves speaking your mind, rather than keeping all your distinctive thoughts and ideas to yourself. Aristotle's teacher's teacher, Socrates, was a living embodiment of this ideal. He roamed Athens challenging the cherished assumptions of his fellow Athenians. For this, he was put to death by a jury. But even then, he had no regrets. From his perspective, the unexamined life was not worth living anyway. Now, this may well be a bit much, but if these ideas have something to them, then at the very least we should not sacrifice our integrity as thinkers willy-nilly for accolades and prestige and approval.

A somewhat different strand of thinking about the good life emphasizes *independence*. Great human lives do something unique, create something new, and refuse to follow the cultural zeitgeist everywhere it goes. This is a major theme for the 19th century philosophers John Stuart Mill and Friedrich Nietzsche. Their ideas are of deep relevance to thinking about the pressures to conform that we experience today, and might help to put these pressures into broader ethical perspective. Many of the trends they identified in their own time seem to exist in an even stronger and more potent form in the 21st century. Hence, revisiting their works, not merely as historical curiosities, but as offering perspectives worth engaging with and drawing from, can greatly benefit us with respect to our current predicaments.

Importantly, for our purposes, the independence they extol cannot be cultivated if we never speak our minds: because we are fundamentally *social* creatures, we must express and

exchange our ideas and values outwardly in order to be genuinely independent.

The book ends with an exploration of how the future of humanity is a condition of many things mattering to us here and now. Philosopher Samuel Scheffler has recently presented novel arguments for this idea. According to him, much of what we value here and now assumes a future for humanity; without such an assurance, we would be gripped by nihilism and despair. This is just part of the picture, however; I argue that we don't just want humanity to continue but also to flourish. But if social pressures can create dangerous blind spots, then given the speed at which modern life is evolving, it seems all the more important to speak our minds so as to combat these blind spots. There is far too much at stake, given what we care about.

Taylor & Francis
Not for distribution

Madness is rare in individuals—but in groups, political parties, nations, and eras it is the rule.

—Friedrich Nietzsche, *Beyond Good and Evil*

Every age has its peculiar folly; some scheme, project, or phantasy into which it plunges, spurred on either by the love of gain, the necessity of excitement, or the mere force of imitation. Failing in these, it has some madness, to which it is goaded by political or religious causes, or both combined.

—Charles Mackay, *Extraordinary Popular Delusions*

THE DIVISION OF COGNITIVE LABOR

Modern society is only possible because of the division of labor. Without division of labor, the most we could achieve is a very meager standard of living. Imagine you had to make everything you use, by yourself, from scratch—without tools created by others, without water and food provided by others, without medicines invented by others. Most of us would not survive for a month, if that. Division of labor makes modern standards of living possible because with individuals specializing in one area, society as a whole is able to be much more productive.

Adam Smith illustrated and developed this idea in his *Wealth of Nations* by using the example of a pin factory. Imagine ten

people tasked with making pins. If each person had to make a whole pin, perhaps each might make ten pins a day. Making a whole pin involves several distinct processes. Let's suppose it involves ten different tasks. Well, if one person had to do all these tasks we can expect that there would be time lost as that person transitioned from one task to another. Furthermore, it would be hard to become skilled at all these different tasks—that would require lots of training and effort. But what if each person in the factory focused on just one of the ten tasks instead? Time would be saved in a myriad of ways, and the factory would be able to produce a lot more pins—though, no person by himself would be making a whole pin. As a result, the factory might produce 10,000 pins total per day, whereas it would have produced only 100 without specialization. Modern society is like this pin factory writ large.¹

But division of labor in modern life is not limited to the production of physical goods. The other face of specialization is the division of *cognitive* labor. Our institutions of knowledge production (universities, thinktanks, private research labs) reflect this feature: researchers inevitably specialize in one tiny sub-sub-field or two in order to make new discoveries. Yet, the division of cognitive labor has deep implications. What we are able to know is inextricably tied to what I will call the *epistemic commons*—the stock of facts, ideas, and perspectives that are alive in society's discourse.

In their book, *The Knowledge Illusion*, cognitive scientists Steven Sloman and Philip Fernbach write: "Language, memory, attention—indeed, all mental functions—can be thought of as operating in a way that is distributed across a community according to a division of cognitive labor."² The authors argue that we know very little, but take ourselves to know a lot because the relevant facts are easily accessible to us. If Sloman and

Fernbach are right, then our epistemic health as individuals—i.e. the extent to which our beliefs accurately represent the world—is inextricably tied to the health of the epistemic commons.

Consider the following. Do you understand how a zipper works? How about a flush toilet? These objects seem basic enough. Knowing how they work isn't exactly rocket science. But people drastically overestimate their understanding of how these simple items function. In one study, Leon Rozenblit and Frank Keil asked people to rate from one to seven how well they understood the workings of such objects. They then asked participants to actually explain in detail how the objects worked. Many were simply unable to do so. And so when asked to revisit the question of how well they understood, subjects drastically lowered their ratings. Psychologist Rebecca Lawson performed a similar experiment where students were asked to explain, by sketching out the mechanism, how a bicycle works. The results were striking—most people were unable to complete the task, even though a bicycle is such a familiar object in our daily lives. This phenomenon, of people thinking they know much more than they actually do, has come to be known as the illusion of explanatory depth.³

Why might we fall prey to this illusion? Well for one, the relevant information is easily accessible. If you want to know how a zipper really works, a simple internet search will give you all the details you need. Though you may not actually as of this moment know the workings of a zipper, the knowledge is “at your fingertips,” as it were. What this suggests is our representation of the world is like a low-resolution map such that “zooming in” only gives a clear picture insofar as we are able to rely upon the knowledge others have. With respect to most areas of the map, we are unable to zoom in by ourselves—and

if we do, we'll just see large pixels that don't look like anything. The division of cognitive labor, then, renders our epistemic lives intricately tied with the efforts and contributions of others.

Furthermore, the very coarse-grained picture we have of the world will itself depend on which perspectives are "alive" in the discourse within our milieu. Consider for example a teenager within a deeply religious sect living in a small village. Suppose that this sect does not believe in Darwinian evolution. The arguments for evolution are not discussed, and when the topic is broached, people quickly dismiss it as an unsubstantiated theory. Some might raise what they take to be decisive counterarguments like: "How come we don't see monkeys turning into humans now?" or "Where are the missing links?" and so on. Now the teenager might be able, in principle, to discover the powerful arguments in favor of evolution by natural selection. There is a copy of the *Origin of Species* at the local library, and she could also spend time delving into encyclopedias and biology textbooks. But for all intents and purposes, her map of the world has a large hole in it. What's more, given that there are ample other constraints on her time, she might simply not find it worthwhile inquiring further.

In this way, there are lots of questions that we might lack the time or imagination to inquire about if the people we're surrounded by consider the issue settled. Division of cognitive labor means we simply cannot independently verify all the claims we take for granted. But that in turn means that if the view our community settles on is mistaken or impoverished, the distortion easily transfers to us. Our epistemic health thus depends on the epistemic health of our milieu.

The 19th century mathematician and philosopher W.K. Clifford underscored this social, interconnected nature of our

ability to understand and describe the world in his landmark essay on the ethics of belief:

Our lives are guided by that general conception of the course of things which has been created by society for social purposes. Our words, our phrases, our forms and processes and modes of thought, are common property, fashioned and perfected from age to age; an heirloom which every succeeding generation inherits as a precious deposit and a sacred trust to be handled on to the next one, not unchanged but enlarged and purified, with some clear marks of its proper handiwork. Into this, for good or ill, is woven every belief of every man who has speech of his fellows. An awful privilege, and an awful responsibility, that we should help to create the world in which posterity will live.⁴

For Clifford, this meant that each of us has an important ethical responsibility: namely, to believe only on the basis of proper evidence. As I will be arguing in the next chapter, if our epistemic situation is a common resource in this way, then we all have a duty to do what we can to preserve the integrity of this resource. However, believing on the basis of proper evidence, though important in its own right, is not enough—we also have a duty to speak our minds.

BLIND SPOTS AND SOCIAL PRESSURE

To set the stage for that argument, it is necessary to examine the way in which the epistemic commons is vulnerable. Tragedies of the commons arise because common resources are often susceptible to damage and degradation.⁵ For example,

industrial pollution can destroy river ecosystems. Analogously, I will argue below, *social pressure* can degrade the epistemic commons.

Consider again the village described above. Why might reasons to accept evolution be systematically repressed here? Presumably because publicly defending such reasons will come at some cost to one's social status, the maintenance of which is a strong motivation for most people.⁶ Somebody discussing evidence in favor of Darwinian evolution might be seen as deviant, and perhaps not a true believer of the religion. Furthermore, accusations of heresy or disbelief can invite severe repercussions in many deeply religious societies—even if such accusations end up being untrue. Thus, even if somebody were to encounter or think of a reason to believe in evolution, they might keep that thought to themselves, especially if they're unsure of the soundness of the reason. Why risk your reputation and social standing (or worse, in many places and times) just to voice some reason you're unsure of?

In this way, social pressure can systematically filter out reasons to believe a particular claim. The reasons that don't get filtered out will make it look like that claim ought to be rejected—even if had there been no such filtering, then people would be justified in believing the claim. In other words, filtering processes created by social pressure allow reasons to pile up on one side of an argument while those on the other side get discarded. Yet, the overall balance of reasons, had open discourse prevailed, might well have supported the other side. Any time we observe social pressure to avoid giving some kinds of reasons, then, we should suspect that a worrisome blind spot exists in some form or another.

Importantly, we can't dismiss the existence of such distortions simply by surveying the first-order evidence (i.e. evidence directly relevant to the issue at hand) presented to us. The problem is created precisely because evidence is filtered in such a way as to support one conclusion. It's then no good to simply look at the evidence that is presented and say: "but the conclusion is obviously right!" The conclusion looks obviously right because countervailing evidence is not allowed to surface and accumulate, due to the presence of social pressure. A collective blind spot can exist in this way even if the members of a community *respond rationally* to the first-order evidence they have.

A DETAILED EXAMPLE: THREE ENGINEERS AND A DAM

Consider the following example. Imagine a situation where three engineers are responsible for the construction and upkeep of a particular dam. Suppose that constructing the dam has been a project that has required enormous funding and mobilization of resources. Imagine also that the dam has an enormous positive impact on the livelihoods of the surrounding community—it provides essential power and irrigation. So, naturally, the community as a whole has a strong interest in the success of the project. Besides, the dam construction is a big feather in the cap for many local officials and politicians. People want to believe it will succeed, and opponents of the project as well as doubters are not looked upon favorably.

Now, a dam bursting, of course, can be really devastating. Suppose there are some good reasons to think that this particular dam will hold during this particular year. But there

are also reasons to think the dam will break. The reasons to think the dam will hold are common knowledge among the engineers, since there's no social pressure not to voice these reasons. However, the reasons to think the dam will break are distributed among the engineers. They keep these reasons to themselves because they don't want to be seen as naysayers.

Suppose that the reasons to think the dam will break outweigh the reasons to think the dam will hold. Basically, given the total evidence, the dam is going to break. The resulting case has the following feature: it would be rational given the evidence of the group as a whole to believe the dam will break. Yet it is not true of any particular individual that she should believe the dam will break given the evidence she has.

To fix ideas, suppose the following are the relevant considerations ('R' for 'reason').

Pro:

R_1 = the dam is constructed with good materials.

R_2 = the structural engineering is sound overall.

Con:

R_3 = the upstream rainfall has been unusually high this year.

R_4 = the spillway design has some defects.

R_5 = the outlet pipe maintenance has been suboptimal.

None of the con considerations is by itself sufficient evidence to think the dam will break, given the pro reasons. However, the considerations taken as a whole—i.e. R_1 through R_5 —support the conclusion that the dam will break. But now

suppose that each of the three engineers knows both R_1 and R_2 (since there's no social pressure to hide these) but only one of the remaining reasons.

So, suppose the first engineer knows R_1 , R_2 , and R_3 , the second knows R_1 , R_2 , and R_4 , and the third knows R_1 , R_2 , and R_5 . Each engineer believes the dam will not break—and given the evidence each has, this is indeed the rational conclusion for each to draw. But the group as a whole is irrational in a sense. For the group as a whole is in possession of evidence such that it would be rational to believe the dam will break, and therefore to take steps to fix it if possible or to evacuate the surrounding population.⁷

Why might each engineer not reveal her counterevidence? Well she might think something like this: "I have some evidence to think the dam will be in trouble, but the overall case for the integrity of the dam is really strong. If I raise concerns about the dam, all that will happen is I will invite social opprobrium. Nothing good will come out of it, because given what I know, the dam is not going to break anyway." So, in other words, there's a downside but there's no upside.

Notice that such silence need not be borne out of pure selfishness. It need not be the case that the engineers only care about themselves and not the people who would be affected were the dam to break. We can even suppose that if the dam breaks, all three of the engineers' houses will get destroyed. So, if any of these engineers knew the rest of the evidence, let's suppose they'd bring it to the attention of the others, despite the risk of social opprobrium. I'll risk some flak if it means saving thousands of lives, including my house, they might think. But the issue is that none of them is in a position to know that the dam will break because they act in a way that's rational given what they know. In

this way social pressure can blind us to what the right course of action is, given what our group as a whole knows.

LESSONS FROM THE 20TH CENTURY

Evidential situations like these can lead to catastrophe. If information is not freely shared within a group due to social pressure, deliberation on very important issues can be distorted. In the above case, the dam will break and ruin many lives.

Moreover, this is not simply an exercise of the imagination. Many avoidable disasters have occurred because there was pressure not to share certain kinds of information. The Chernobyl disaster, in which a nuclear powerplant malfunctioned and exploded in what is modern-day Ukraine, is perhaps a paradigm example. Due to the authoritarian, top-down government in place at the time, individuals had incentives not to raise alarms about radiation levels, the nature of the explosion, substandard materials, etc. The result was devastating for thousands of people, many of whom continue to feel the effects of radiation poisoning to this day. The HBO series *Chernobyl* offers a detailed look at the deliberations and actions of various individuals as they grappled with the situation in a way that brings out the incentives they had to distort or suppress information.⁸

Democracies typically do a better job of avoiding unnecessary disasters and missteps like this. The victory of the Allies in World War II can be partly attributed to the nature of information flow within democratic decision-making.⁹ In the democracies, members of the army were relatively more able and willing to offer information that would lead to course-correction by the upper chain of command. By contrast, within the German

army and air force, people were much more hesitant about displeasing their superior officers with news or information or strategic perspectives that might be seen to dampen the war effort.

Democracies are also able to allow the spread of key information through a more open media. Journalists are less prone to intimidation by the government, and thus can quickly disseminate crucial news to civilians and government officials alike. Luther Gulick, who served as a high-level American official during World War II, explained that in contrast, decisions within authoritarian governments are “hatched in secret by a small group of partially informed men and then enforced through dictatorial authority.”¹⁰ Democracies are thus able to avoid some of the epistemic pitfalls that beset authoritarian regimes because the channels of information are much freer.

This is no cause for complacency, however. Democracies are not immune to such problems. For example, the infamous Bay of Pigs Fiasco, a failed U.S.-backed landing attempt on Cuba in 1961, resulted in part because those who had doubts about the plan suppressed their reservations.¹¹ Moreover, social pressure need not always come from government authorities. Think of college students who feel pressure to binge drink, the many of us who feel pressure to dress in particular ways, teenagers who (used to) feel pressure to smoke cigarettes—or, what’s more relevant here, people who feel pressure not to publicly express certain social or political opinions. Such forms of social pressure do not come top-down, from some governmental chain of command. Rather, they are much more spontaneous and organic. These pressures emerge from the incentives, interactions, and choices of millions of people who shape a particular

culture. Democracy, then, does not solve all the informational problems systemic within authoritarian regimes.

THE IMPORTANCE OF REFERENCE NETWORKS

Which pressures to conform we experience depends crucially on our *reference networks*. Philosopher Christina Bicchieri, known for her work on social norms, uses the concept to illustrate that the set of people who matter when it comes to influencing your norm-guided behavior is not always the set physically closest to you. Thus, a religious, married Mormon individual's reference network might centrally involve her religious family and other Mormon friends rather than the atheist singles living in her city, who may be physically closer to her. An academic's reference network may largely consist of other academics even if he lives in a neighborhood comprised mostly of blue-collar families.

The cultures that have the greatest impact on us, then, are not necessarily those most physically close to us. And they may be cultures associated with a minority of the population, as in the Mormon example above. It's not always a matter of similar demographics either. For instance, according to Bicchieri, "A young woman in Philadelphia wearing very high heels will probably not care what other women do in India, or even New Orleans. Her reference network may be the 'fashionable' crowd in her town, those who she is likely to meet and give her a chance to 'show off,' or it may be a celebrity, magazine starlets, or TV series that girls in her reference network follow."¹²

Now the pattern I have been describing—where the evidence on one side is common knowledge but the evidence on the other side is distributed and isolated—has important epistemic

and ethical implications which have received little philosophical attention. The pattern may explain some of the phenomena we observe in public discourse and public opinion—in particular, how political polarization on a range of separate issues can be maintained. Crucial to the analysis is a characterization of the social pressures and information channels within people’s reference networks.

Consider a person, Alice, whose reference network consists mostly of supporters of one party. Now imagine that Alice agrees with her reference network, given her analysis of the information she has, on all but one issue. Let’s say that on this particular issue, she has some pretty strong evidence. Suppose this issue has to do with the causes of, and effective methods of reducing, violent crime. Alice has done a deep dive into the available data on crime and policing, and is statistically savvy enough to draw reasonable conclusions. Now, Alice disagrees with her group on this issue, but the evidence she has is strong enough that many others in her group, if they deliberate rationally with this evidence in hand, will come to change their views on the topic.

Nonetheless, Alice may not share this evidence with her group. For one, she might be keenly aware of the flak she will receive if she disagrees with her reference network on an issue of contentious partisan disagreement. Even if she doesn’t explicitly lose friends, people may look upon her with more suspicion. She might also lose out on future professional opportunities if she signals to others that she’s not a good member of the ideological group. So that’s the downside. What’s the upside? Well, by Alice’s own lights, her group gets things mostly right anyway. Sharing the evidence that she has will only empower the other group relative to her group, which will be, on balance, counterproductive from

her perspective. Thus, the rational thing to do is to keep that evidence to herself.

But now note that Alice may not be the only one facing such a situation. Imagine that Bob, another member of the group, disagrees on the issue of the minimum wage. He has good evidence that would suggest a position contrary to the group's accepted wisdom. But he agrees with the group on all the other issues, including crime and policing. Claire might disagree with the group's position on abortion, having thought a lot about the topic and delved into the arguments on both sides of the debate. However, she agrees with the group on all the other topics, including crime and policing as well as the minimum wage.

The dissenters don't share their evidence. However, were the evidence to be shared, the group's overall position on a variety of partisan issues may well be dramatically undermined. This fact, however, will not be transparent to the members of the group. Given what first-order evidence they have, it makes sense for them to think their group is right on the whole.

This dynamic may be a good explanation of the pattern of polarization we observe in modern life. We find public opinion divided strongly along partisan lines, but on issues that seem to be rationally disconnected. For example, particular positions on gun control, criminal justice, immigration, climate change, abortion, minimum wages, and a host of other issues travel together. In other words, if you pick a person at random and all you know about them is their view on gun control and climate change, you can probably predict their opinion with good accuracy on abortion and immigration. But why should this be so? It would seem that these issues have nothing to do with each other—a particular position on gun control shouldn't commit you to any view on abortion

or minimum wage. To put it another way: the set of considerations, statistics, and arguments relevant to each of these issues is quite distinct. So, what explains this pattern? Partisans of either side, of course, will be tempted to say the other side is simply wrong about every issue at hand. And they'll point to the first-order evidence they have on a variety of issues, which supports the views of their side, and thus implies that the other side's views are wrong.

But what could possibly explain why the other side would get everything wrong and one's own side would get everything right? There are a variety of flat-footed responses partisans will be tempted to give: the other side is dumb, brainwashed, evil, selfish, and so on. These responses, however satisfying from the point of view of the partisan, are challenging to sustain.¹³ It is also worth exploring non-cynical hypotheses that could show how generally well-meaning individuals come to form their beliefs on contentious issues. What is needed here is a way of explaining how people seem to come to accept one of two packages of disconnected beliefs. The model formulated above, of social pressure acting as a screen on contrary evidence, is a plausible candidate. The model can explain how rational people, doing the best they can with the evidence they have, can nonetheless form a group that is irrational.

THE DANGER TODAY

In her groundbreaking work on the dynamics of public opinion, political scientist Elisabeth Noelle-Neumann argued that fear of isolation can create a "spiral of silence," where only one side of an issue is publicly defended. The core mechanism she identifies is this: people don't want to say things that they believe might risk eliciting the disapproval of others; they

don't want to potentially lose friends and get pushed out of their social groups. There is a fear of isolation. So, instead of saying what they really think about a particular issue, such individuals keep mum. Once the process is set in motion, more and more people become silent about their true opinions.¹⁴

Spirals like these typically occur with regards to contentious, emotionally laden moral and political issues. A spiral of silence can drive even the majority opinion underground if the minority is sufficiently vocal, and especially if mass media repeatedly and concordantly come down on one side of the issue. Eventually, the spiral of silence causes the majority opinion to effectively disappear, while the previously minority opinion becomes the dominant societal assumption.¹⁵

What does this mean for us, now? Well for one, we shouldn't assume, for all the reasons explored so far, that such spirals of silence induced by social pressure (real or perceived) are going to line up with the truth all the time (or even most of the time). Spirals of silence are sensitive to social forces, not to the truth. Thus, they can cause society to settle on opinions that are quite misguided.

However, in order to know what policies to support or how to remedy various social problems, we need to have an accurate idea of what the social world is like. The very best of intentions can have terrible consequences if those intentions are not supplemented with an accurate picture of the world. (Indeed, under some description, more or less all of the worst actors and movements in history can be said to have "good intentions.") But social pressure can warp our collective picture of the world without individuals being in good positions to detect the distortion. So, the more we allow spirals of silence to occur, the more chance there is for the road to hell to be paved with good intentions.

The danger we face today is that many of us have quite confident views about lots of contentious issues, as well as lots of issues that have been “settled,” not via a process of institutionalized disconfirmation, but rather through spirals of silence. But this means that the steps we might take to mitigate economic and social problems could backfire, making things worse. The risk becomes greater the more radical, as opposed to piecemeal, solutions we embrace. We might also be misdiagnosing what the problems are in the first place. And we might be missing various forests for the trees. Our Chernobyl, so to speak, might not involve a nuclear powerplant, but might instead manifest itself in the way we conceive of and try to solve social and economic problems.

One way to respond to this predicament is to encourage epistemic humility.¹⁶ Perhaps we should all just check ourselves. This however, is far easier said than done. Knowing our epistemic limitations in abstract terms may not actually induce humility in us (especially the loudest among us) when the rubber meets the road. The only way to properly mitigate our dangerous blind spots is for courageous individuals to speak their minds, and refuse to buckle to social pressure. This is not to say that epistemic humility and other tools for critical thinking are not important or worth cultivating. But if knowledge is a collective enterprise, individual epistemic humility can only go so far. This humility, for instance, cannot prevent a Chernobyl—only people sharing their evidence can.

UPSHOTS FOR CONFORMITY AND COOPERATION

Whenever there is social pressure to support one particular conclusion and to refrain from giving reasons to doubt that conclusion, there will be a systematic filtering out of

important information. If, in the end, the group's conclusion is correct, it will be an accident, a stroke of luck. Even then, the group's picture of the world will likely be warped in some way or another. Depending on the context, such distortions can have terrible practical consequences as well.

Two qualifications are in order. First, there is often stigma attached to obviously irrational "giving of evidence." Imagine a person who says "Hey, here's some evidence for thinking the sun goes around the Earth: horses have four legs." People would probably think there's something off with this person, or that he's joking in some way; the argument is a *non sequitur*. Though, people presumably won't get mad at him. They'll probably just try to make sure he's not having some kind of breakdown. Second, there are pressures to be relevant in conversation. Thus, giving evidence about the employment effects of minimum wages, say, is not relevant in the context of an ongoing conversation about how the dinosaurs went extinct, and would rightly be frowned upon.

But society often applies pressure on us to conceal evidence in a way that is independent of the quality of the evidence *qua* support for a particular conclusion, or considerations of relevance. Thus, imagine a person giving evidence about the effects of minimum wages, which conflicts with the convictions of her social group, in the context of political discussion. Even if her arguments are good, there will be an inclination for people to act as if to say, "Which side are you on?" or "You're not the good person I thought you were." Or imagine giving reasons to believe in evolution by natural selection within the context of a deeply religious sect of a particular sort. There, people might suspect that the person making such arguments is actually a closet disbeliever, to be shunned.

These forms of social pressure, which come apart from the perceived badness of the argument or evidence on offer, are ubiquitous with regards to ideas that social groups have an affective investment in. The pressures are often tied to issues regarding which taking a particular stand is important to people's identity in some way—be it social, political, religious, national, or professional. What's more, people will often publicly display *anger* towards those who share evidence supporting disfavored conclusions of this sort.

It is these kinds of pressures, especially when they apply to topics of great epistemic importance or generality—like the theory of evolution or the effects of minimum wage—that can lead to worrisome blind spots. Therefore, providing evidence that challenges prevalent opinion, at personal cost, can be a useful service to society.

Of course, it's important not to overstate the point. A pure contrarian, that is, someone who disagrees with people for the sake of disagreement, is probably not doing much of a service to society. First of all, most people and groups, most of the time, get most things right. If that's correct, then pure contrarians will be wrong most of the time. Secondly, a pure contrarian's opinions will not contain much "signal"—since he disagrees simply because he wants to disagree, people won't be in a good position to take him seriously.

Legal scholar Cass Sunstein, in explaining the perils of conformity thus warns that "We should not lament social influences or wish them away. Much of the time, people do better when they take close account of what others do."¹⁷ Many of the conventions and social norms that societies adopt serve useful functions. Furthermore, even if the contentious issues that generate disagreement or social pressure seem pervasive,

they reflect the tip of an iceberg. In almost any society, there is a base of mutual agreement on facts and norms. Most members of most societies agree that theft is wrong, that plants need water to grow, that two and two make four, and that it makes sense to drive on the side of the road on which everybody else is driving. If there was constant disagreement about everything, society, with all the coordination and cooperation it requires, would be impossible.

The tendency to come to mutual agreement with others would have had enormous evolutionary import for our ancestors. Cooperating with others to hunt large game, for example, would have required a lot of agreement and coordination. The hunters would all have to agree about which kind of animal to hunt, what technique to use, what roles each individual must play, how to divide the bounty, and so on. Agreement, then, is indispensable for cooperation, which in turn is indispensable for human society.¹⁸

Yet, conformity also has a dark side. Individuals who speak their minds despite the pressure serve a crucial function. Sunstein goes on to say, "But social influences also diminish the total level of information within any group, and they threaten, much of the time, to lead individuals and institutions in the wrong directions. Dissent can be an important corrective; many groups and institutions have too little of it . . . conformists are free riders, whereas dissenters often confer benefits on others."¹⁹

INSTITUTIONS OF KNOWLEDGE PRODUCTION

Social pressure creates blind spots by making it costly to provide evidence on one side of an issue, while making it costless or even beneficial to provide evidence on the other side

of the issue. Whenever such incentives exist, we should suspect that our resulting view of the world is warped in some way. These incentives are particularly important to address within the institutions responsible for knowledge production and dissemination: research groups of various sorts and fields and academic departments within the university system.

Given modern division of labor, such institutions specialize in knowledge production; the rest of society thus relies upon them for providing an accurate picture of the world. Other individuals in society, however, do not have the time or resources to check all the work produced by such institutions, and so an element of trust is necessary. Analogously, you don't have the time or wherewithal to check all the work done by your lawyer, doctor, or accountant—when it comes to your interaction with such specialists, then, an element of trust is involved.

However, social pressures within institutions responsible for knowledge production can undermine their mission and distort their product. Science works well only in a context of institutionalized disconfirmation: that is, a situation wherein researchers are free and even incentivized to disconfirm any and all hypotheses that are in contention.

Over time, science has disconfirmed hypotheses that would seem exceedingly natural to humans observing their world. Many things that seem intuitive to us turned out to be false. The Earth, it turned out, is roughly spherical, though it looks flat from our vantage point. And while the sun looks like it goes around the Earth, the reverse is true. In the 17th century, Galileo Galilei suffered persecution at the hands of the Catholic Church for defending this idea. Science naturally works best when such costs are absent—so that it doesn't take a Galilean personality to seek the truth.²⁰

Modern physics has upended our intuitive picture of the world even further. The things that look 'solid' to us—tables, rocks, books, etc.—are actually made mostly of empty space.²¹ And the fundamental units of physical reality have both particle-like and wavelike properties. Albert Einstein famously showed that time is not absolute. Whether or not two spatially distant events are simultaneous depends on the observer's frame of reference. He further showed that space and time are intertwined in such a way that it's best to think of them as spacetime. According to the best models we currently have to explain the behavior of large objects, gravity is the result of spacetime "bending" around massive objects.²² Trippy stuff!

How has science made these remarkable discoveries that are so far from our intuitive sense of the world? Science is a collaborative effort, and no one person can do it all by themselves, even within a sub-sub-field. Science involves enormous division of labor. But for us to be able to trust the products of science, the incentives have to be right. The incentives that individual scientists face must be aligned with finding the truth, wherever it may lie. Generally, this is the case, and that is why science has been on the whole very successful. In physics or chemistry, if you are able to find experimental data that disconfirms an important and commonly accepted hypothesis, you will receive many professional goods—you'll likely get published in prestigious journals like *Nature* or *Science*, you might get big grants in the future, an endowed chair, maybe even the Nobel Prize.

Given these incentives, physics and chemistry are *self-correcting*. If a hypothesis is easily disconfirmed, it won't last for long. Researchers, incentivized to disconfirm it, will quickly design experiments to show why the hypothesis doesn't hold. Sloman

and Fernbach write: “Scientific claims can be checked. If scientists are not telling the truth about a result or if they make a mistake, eventually they are likely to be found out because, if the issue is important enough, someone will try and fail to replicate their result.”²³ Many scientists have echoed the importance of this feature of science over the years. Any time the accepted wisdom strays from the truth then, a course-correction will quickly follow.

Understanding knowledge production as a collective endeavor, which relies heavily on a well-maintained epistemic commons, helps us appreciate why John Stuart Mill defended his somewhat radical sounding account of justification for our scientific beliefs in *On Liberty*. He wrote:

If even the Newtonian philosophy were not permitted to be questioned, mankind could not feel as complete assurance of its truth as they now do. The beliefs which we have most warrant for, have no safeguard to rest on, but a standing invitation to the whole world to prove them unfounded.²⁴

Thus, imagine if critics of Newton’s physics found themselves unemployable or prone to receiving censure, threats, etc., as soon as they challenged part of the view. Could a person living in Mill’s time, circa the mid-19th century, be able to trust the science of physics? Could he have reasonably believed in Newton’s laws if people faced a very uphill battle in trying to disconfirm them and he knew about this situation? Plausibly not. For, especially if this person is not a physicist, he lacks the wherewithal to check the researchers’ work. For all he knows there may be good reasons to reject Newtonian physics that are just not allowed to surface.

Indeed, as it turns out, Newtonian physics was accurate only in approximation. For macroscopic objects traveling at relatively low speeds, i.e., well below the speed of light, Newton's laws allow us to make approximately true predictions. However, as Einstein later showed, some decades after Mill had passed away, Newtonian physics breaks down when it comes to objects moving close to the speed of light. Furthermore, while Newton assumed that space, time, and mass are absolute, Einstein showed that they are relative. Which events are simultaneous, how long an object is, how much mass it has, all depend on the observer's frame of reference. If you are traveling at, say, half the speed of light relative to where I stand, then the length of a particular table will be quite different for you as opposed to me. Hence, even Newtonian physics, which was by Mill's time well established and confirmed with countless experiments, turned out not to be sacrosanct.

The scientific process, then, must be structured in a certain way for it to merit our trust and reliance. If there were contrary evidence to be found, would it be discovered, published, and incorporated into the mainstream scientific consensus? The answer to this question must be yes.

In some sense, the scientific enterprise must be *objective*. What does such objectivity mean? Philosopher Helen Longino argues that it requires an openness to what she calls *transformative critique*. For Longino, science is fundamentally a social practice, and it is precisely due to this fact that its objectivity can be secured. Individual researchers are bound to have their idiosyncratic perspectives and biases. However, "science" is not simply the aggregation of the findings of individual scientists. Science is fundamentally practiced by social groups, not lone individuals. What gets counted as scientific knowledge results from social processes like peer review, attempts at replication,

citation patterns, and clashes between defenders of alternative hypotheses and paradigms. This is a feature, not a bug. “Only if the products of inquiry are understood to be formed by the kind of critical discussion that is possible among a plurality of individuals about a commonly accessible phenomenon,” says Longino, “can we see how they count as knowledge rather than opinion.”²⁵ Consequently, the more diverse points of view there are within a scientific community, the more objective the process is likely to be.

These lessons are not limited to science. Philosophy or literary criticism can be objective in this way too, according to Longino. However, the objectivity essentially depends on whether the social conditions within the field allow for robust critical discussion. A healthy field of inquiry, one whose product we have reason to take seriously, has to be one where people are incentivized to critique and disagree with ideas, such that no idea is sacred or beyond criticism.

To fix ideas, consider the philosophical field of metaethics. This subdiscipline asks foundational questions about the nature and epistemology of moral claims. These questions include, but are not limited to, the following. Are there any moral facts? If there are moral facts, are they subjective or objective? Would such facts be the sort of thing that can be discovered and investigated by the methods of natural science? How might we come to possess moral knowledge? When we say “murder is wrong,” are we expressing something more like a belief or something more like an emotion?²⁶

Now metaethics, given my own impression of it, is a good example of a field that is working reasonably well. People defending a wide range of positions—naturalism, non-naturalism, error theory, expressivism, constructivism, Humeanism—have climbed to the top of the profession, winning prestigious awards

and endowed chairs, working at elite universities, and so on. A variety of perspectives and styles of argument can thus exist and flourish within the discipline. There's no stigma, as far as I can tell, attached to working on either side of the various debates in metaethics. Consequently, younger members of the profession feel free to follow the argument where it leads. And so many different kinds of positions within the logical space have renowned and well-respected defenders.²⁷

When we look at the product of this discipline then, we can be fairly confident that few stones have been left unturned. If there was an easy argument to be made against some position it likely will have been made; the remaining fruits on the tree will probably be pretty high up. We don't have to worry about reasons piling up on one side of the debate but being filtered out and discarded on the other side. Part of why metaethics works as well as it does might have to do with the fact that its subject matter—though fascinating and stimulating—does not “excite the passions.” People just aren't going to get mad at you for defending non-naturalism or expressivism.

Due to the absence of such social pressure, we find each position having several defenders. This in turn reinforces the willingness of metaethicists to follow the argument where it leads. There's a kind of strength in numbers. Contrast this with a hypothetical scenario where there are 100 naturalists (i.e. those who believe that moral properties are natural properties, in principle investigable by natural science) for every non-naturalist (those who deny naturalism). In such a case, it is hard to imagine not feeling isolation or social pressure against defending non-naturalism. Such pressure, whether real or perceived, would especially impact early career researchers, such as graduate students, whose future careers are uncertain. A promising graduate student who is inclined to defend non-naturalism

might think twice. The fact that naturalists are in the overwhelming majority may be taken by such a student—whether consciously or subconsciously, rightly or wrongly—to suggest that defending non-naturalism is a bad career move.

Suppose now we add a stigma to this. Imagine that defenders of non-naturalism were publicly censured and ascribed bad character traits. We can see how this would cause reasons to pile up on one side of the debate. It would create perverse incentives that should undermine the trust we ought to have in the product of this community of research. Fortunately, as it stands, such pressures do not exist within metaethics. In fact, it would be considered grossly unprofessional to publicly ascribe bad character traits to one's intellectual opponents within the field. A person who engaged in ad hominem attacks would quickly lose standing in the profession.

I have been describing modern physics, chemistry, and metaethics as fields that model healthy atmospheres of research (though of course they may not be perfect). But is this true across the board with respect to our institutions of knowledge production? Along with others, economist Glenn Loury suggests there is reason to worry. In a provocative 1994 article called "Self-Censorship in Public Discourse," he writes:

Some areas of social science inquiry are so closely linked in the public mind to sensitive issues of policy that an objective, scholarly discussion of them is now impossible. Instead of open debate—where participants are prepared to be persuaded by arguments and evidence contrary to their initial presumptions, we have become accustomed to rhetorical contests—where competing camps fire volleys of data and tendentious analyses back and forth at each other.²⁸

In a later passage, Loury claims that perverse incentives within a community of research can reduce the degree to which we should take its output seriously:

The notion of objective research—on the employment effects of the minimum wage, say, or the influence of maternal employment on child development—can have no meaning if, when the results are reported, other ‘scientists’ are mainly concerned to pose the ad hominem query: ‘Just what kind of economist, sociologist, and so on would say this?’ Not only will investigators be induced to censor themselves, the very way in which research is evaluated and in which consensus about ‘the facts’ is formed will be altered. If when a study yields unpopular conclusions it is subjected to greater scrutiny, and more effort is expended toward its refutation, an obvious bias to ‘find what the community is looking for’ will have been introduced. Thus the very way in which knowledge of the world around us is constituted can be influenced by the phenomenon of strategic expression.²⁹

To the extent Loury is right, our epistemic condition with respect to the output of fields that are politicized in the way he describes above is shaky. Given the mountains and mountains of evidence relevant to all these policy-adjacent debates, though, none of us has the time, energy, or expertise required to dig through everything and properly make up our own minds. We inevitably have to rely on the journals, textbooks, and public lectures of the practitioners of these fields. But if the incentives within these fields are skewed in the way Loury describes, then such reliance will expose us to a lopsided selection and analysis of the facts out there. Depending on

the case, this may well put us in a worse position epistemically than either ignorance or suspension of judgment with respect to certain topics. It would be like a jury being made to hear hours of arguments from the prosecution, and zero from the defense. Likely, the jury would have been better off before, when they had no opinion on the case!

All this puts us in a serious predicament, especially because, unlike metaethics (sorry metaethicists), the kinds of research Loury alludes to are extremely important to get right from a practical, policy-making perspective. The proper maintenance of the epistemic commons, when it comes to such fields of knowledge, then, is all the more important.

CONCLUSION

Social pressure to conceal evidence can create blind spots that can often be dangerous. Any time there is social pressure of this kind, we should suspect that our view of the world is distorted in important ways. What's the ethical upshot of this? If this is right, what should we, as individuals, do? That is the topic of the next chapter.

Two

Let us not concur casually about the most important matters.

—Heraclitus (DK 47)

CONFORMISM AS FREE RIDING

Free riding is tempting. But we typically think it's unethical. At the very least, we don't see free riders as exemplars of moral virtue. Think of the roommate who shirks on the cleaning tasks because he knows the others will do them. Or think of the person who never picks up after their dog when walking through a park. Or imagine a person who drives a gas guzzling Hummer and never recycles. All these characters rely on others to do important work which they themselves benefit from—keeping the house clean, the park sanitary, the environment non-degraded—but aren't willing to put in the required effort themselves.

In general, whenever something is a common resource, people are tempted to free ride in this way. A clean sink, park, or atmosphere are common resources in the sense that it's not possible or feasible to exclude people from deriving the associated benefits. A clean atmosphere benefits everybody. Yet, for each person, the costs involved in keeping it clean—i.e. recycling, driving more efficient cars, eating less

meat, etc.—outweigh the benefits accrued to *that* person by her actions to keep it clean. So, it's nice to have a clean park, but perhaps for many, the effort involved in picking up after their dog is not worth the extra cleanliness that results from the vantage point of the individual. In other words, the costs are concentrated but the benefits are distributed. Any situation with this structure generates the temptation to free ride.

I have been arguing that there is an *epistemic commons* in an important sense as well. Our epistemic health depends on the epistemic health of others, due to the division of cognitive labor. For one, we can't verify everything ourselves. Furthermore, our patterns of thinking are deeply influenced by those in our surrounding culture. Pressures to suppress evidence can lead to avoidable and potentially dangerous distortions of our view of the world. In cases like the dam described in the previous chapter, or the Chernobyl nuclear plant accident, such distortions have far reaching, disastrous consequences. Just as it is important to preserve the health of the physical commons—our atmosphere, rivers, oceans, etc.—it is important to preserve the health of our epistemic commons, viz., the stock of evidence and perspectives alive within the communities we're part of.

The problem is a commons problem because typically everyone benefits (at least in the long run) from their community having a better picture of the world. Everyone involved with the dam case of the previous chapter is better off if information is freely shared. Likewise, presumably all parties would have been better off if the evidential bottlenecks leading to the Chernobyl disaster had been avoided. These cases involve practical considerations—the disaster has an enormous practical impact (uniformly negative) on all parties. But we can make the same point in purely epistemic terms as well. The

group of partisans described above is such that all individuals would benefit epistemically—i.e. their own picture of the world would become more accurate—if information were freely shared among members.

The issue, however, is that the perceived costs are concentrated. If Alice, in the partisan group described in the previous chapter, shares her evidence about crime and policing with others, then, as far as she can tell, she will lose standing within the group, but the (epistemic) benefits of her contribution will be spread out. Analogously, the lazy roommate's efforts to clean the dishes involve a concentrated cost to him but the benefits of a clean sink are enjoyed by everyone in the house. Thus, in both cases, there is a temptation to free ride.

In both these cases, it is *rational* in one sense to free ride. Assuming that the other roommates let him get away with it, the lazy roommate maximizes self-interest, construed narrowly, by allowing others to clean while he relaxes. Similarly, the person who doesn't pick up after their dog enjoys the benefits of a clean park while shirking the cost. Likewise, I want to add, conformists obtain benefits from a well-maintained epistemic commons while allowing others, i.e. dissenters, to do the hard work. This may not be irrational from the vantage point of self-interest, construed narrowly.

Now, I am not alone in making this observation. Sunstein emphasizes this aspect of conformity in the quote presented in the last chapter. Dan Kahan and colleagues, known for their work on cultural cognition and motivated reasoning, echo the idea:

It is perfectly rational, from an individual-welfare perspective, for individuals to engage decision-relevant science in a manner that promotes culturally or politically congenial

beliefs. What any individual member of the public thinks about the reality of climate change, the hazards of nuclear waste disposal, or the efficacy of gun control is too inconsequential to influence the risk that that person or anyone he or she cares about faces. Nevertheless, given what positions on these issues signify about a person's defining commitments, forming a belief at odds with the one that predominates on it within important affinity groups of which such a person is a member could expose him or her to an array of highly unpleasant consequences.¹

Others have pointed out that being an informed voter is in fact *irrational*, unless you simply enjoy being informed or think there are strong enough moral reasons to be informed. The idea is this. Being an informed voter takes a lot of time and effort. It involves looking closely at complicated bills and voting records. It also involves learning a lot of economics, reflecting carefully on public policy, being aware of various data, and so on. So, all this goes on the cost side of the ledger. What goes on the benefit side? Well, you might cast an informed vote. But individual votes almost never make a difference. The chance of your flipping a national or even statewide election is exceedingly small. So, there's not much accruing on the benefit side at all. Hence, most voters are *rationally ignorant* about the important issues.²

However, what I want to emphasize here is that we're often not in a position to know how the evidence we have will contribute to the epistemic health of the group. Our evidence may well be a crucial part of either eliminating the blind spot or changing the incentives of others so that they will reveal evidence that cures the collective blind spot. And this might be worth the cost for us given what we care about, or what we

should care about. While the notion of speaking your mind can be interpreted in different ways, this chapter will focus on the moral dimensions of individuals sharing their evidence.

DAM AND ENGINEERS REDUX

Consider again the case of the engineers and the dam, from the previous chapter. Suppose one of the engineers, the one with the evidence denoted by 'R₃,' comes forward and shares her information. Now, R₃ becomes common knowledge. Well, it then becomes clear to the other two engineers that they should come forward with their information—for now, their sharing information becomes manifestly consequential.

As I set up the case, each piece of counterevidence—i.e. evidence supporting the conclusion that the dam will break—is not enough, by itself, to defeat R₁ and R₂. It is only when you take R₃, R₄, and R₅ together that you can conclude that the dam is going to break, despite R₁ and R₂. Now, evidence functions in a myriad of different ways. But to make things simple, let's suppose that the evidence here works in a linear and aggregative fashion. Each piece of evidence, let's suppose, has a weight of one. And the evidence adds up as weights do on a scale. So, if R₁ and R₂ are on one side of the scale, and if only one of the other pieces of evidence is on the other side, then the scale tips in favor of thinking the dam will not break. If R₃, R₄, and R₅ are together on the other side, however, then the scale tips in favor of thinking the dam will break. As things stand, each engineer knows only one of the con reasons, and so from each engineer's vantage point the dam is safe, and there's no point in coming forward with worrisome evidence, only to invite social opprobrium. But the crucial thing is that no engineer knows that the others are in possession of

evidence such that were it to be revealed, the thing to conclude would be that the dam is going to break.

But when one of the engineers comes forward with R_3 , the reasons come into equipoise from the vantage point of either of the other engineers. Consider what their weight scales will look like. One of the remaining engineers will know R_1 and R_2 , on the pro side, and R_3 and R_4 , on the con side. The other engineer will know R_1 and R_2 , on the pro side, and R_3 and R_5 , on the con side. So, from both their vantage points, the dam may well really break but they can't be sure—the weight scale is in balance, so to speak.

But now, the reasoning that says “my evidence won't make a difference anyway” does not work. Each of the remaining engineers now has evidence that they definitively *should* reveal, even from a narrowly self-interested perspective. If it's up in the air whether the dam will break, given the evidence, then serious discussion and further testing should follow. This would thus incentivize the other engineers to reveal their evidence and seek discussion about the real possibility of breakage.

So then suppose one of the others comes forward, with R_4 . Well, it is now clear to the last holdout that the dam is going to break—for now, he has all the relevant evidence: R_1 through R_5 . Since the breaking of the dam is terrible from the perspective of what he cares about, we can assume, he will share the last piece of evidence if he's minimally rational. Catastrophe will thus be avoided.

What I want to show with the use of this example is that we're often not in a position to tell how our evidence impacts the overall case for a particular claim, and, moreover, how it might change the incentives of *others* with respect to revealing their own evidence. The key point is that the presence of social pressure can obscure the evidential landscape for the groups

we're part of. Thus, whenever there is social pressure to avoid revealing evidence, we should suspect—though we can't tell for sure—that an important blind spot exists. We may not know how exactly our evidence might bear upon epistemic commons. We don't know whether it will be the crucial piece that clears away a misunderstanding or shows some matter in an important new light. It may well be inconsequential, but then again, it may not. When social pressure obscures the evidential landscape, then, the problem becomes deeper than the one described by Kahan and colleagues.

A DUTY TO SPEAK YOUR MIND

Whenever there is social pressure to refrain from revealing some evidence we have, I contend, we should take ourselves to have a duty to reveal that evidence—it is in this sense that we have a duty to speak our minds. It is a *prima facie* duty: i.e., one that need not be decisive in all contexts. For example, we have a *prima facie* duty not to break promises. But if breaking a promise to get lunch with a friend is the only way you can save someone's life, then you should obviously break the promise.³

Indeed, sometimes the costs of revealing forbidden evidence will be prohibitive. Imagine a person eking out survival within Stalin's Soviet Union. Stalin conducted regular purges and orchestrated the killings of hundreds of thousands who showed the slightest signs of dissent. In many 20th century communist regimes, people could face execution for expressing evidence to the effect that particular policies were not working as intended or that the rations were too small for their families.⁴ In such gruesome contexts, it is too much to expect people to speak their minds—morality cannot be that demanding.

However, in many situations, the costs, though they exist, are nowhere near facing execution or death-by-labor in a Soviet gulag. The costs for us, living in modern democracies, can often be real. In the worst case, they might involve losing your job and having to find other employment. But often, the costs might just involve small losses in social or professional status. The partisans in the example mentioned earlier likely won't get fired for revealing dissenting information about crime or abortion. They may well, however, lose some standing amongst their social group. Nonetheless, this cost is not prohibitive. Morality cannot be too demanding, but it does make demands on us—it instructs us in many cases to forgo narrow self-interest for wider goods.

If morality never instructed us to sacrifice narrow self-interest, then no case of free riding would be objectionable (besides a whole host of other bad behavior). It would then not be wrong to litter, pollute heavily, let others do the dishes, and so on. But surely such behaviors are wrong because they rely on the cooperation of others while offering no cooperation from one's own end. Whenever social pressure exists to suppress one's evidence, then, I claim there is a *prima facie* duty to defy that pressure—i.e. share one's evidence despite the social costs, so long as those costs are not prohibitively high.

This is just a natural extension of the way we think of our ordinary moral obligations in a variety of commons contexts: we have a *prima facie* duty not to free ride, unless doing so would be prohibitive. Hence, we can understand a poor fisherman who overfishes a lake to feed his family. The costs of cooperation for him—namely, not being able to feed his family—are too high. But we rightly frown upon a healthy, able individual who does not pick up after their dog, or a

person who doesn't recycle even if the option is available and easily accessible to them.

IMPERFECT DUTIES AND KEEPING YOUR POWDER DRY

What exactly does the duty to speak your mind amount to? In other words, if there is such a duty, when does it recommend us to share our evidence? A natural model suggested by the discussion so far is the following:

Whenever there is a social pressure not to share piece of evidence E, and the cost of doing so for you does not meet some threshold T, you should share E.

As described earlier, the presence of social pressure not to reveal certain evidence suggests the likely presence of a blind spot. Revealing one's evidence may then be an important corrective and a service to others vis-à-vis the proper maintenance of the epistemic commons. On the other hand, it is not plausible that morality is so demanding that it requires individuals to share their evidence even in the face of execution or imprisonment. The principle above would thus seem to capture both these desiderata.

However, the principle as stated faces several problems. The first is the grandma's ugly sweater problem. Suppose your grandmother knits a sweater for you as a gift for Christmas. It turns out that you don't like the design and it's not really your style. Should you say this as you tear off the wrapping paper? Obviously not. But this case would seem to meet the conditions outlined in the principle above. There is social pressure to avoid saying you don't like it: your family will likely frown on you for saying this. And the costs are not prohibitive—you're

not getting sent to the gulag. What's going on in this case? Well, one key feature here is that the underlying matter is not of broad importance. In contrast, the cases we looked at earlier did involve matters of broad social importance—the stability of a dam, the proper policy with respect to crime, the ethics of abortion, and the like. Thus, plausibly, the duty to share evidence would seem to kick in only when the matter at hand is of sufficient importance.

The second problem is that even if the principle may not be too demanding in any one particular case, following the principle *all the time* would make morality too demanding. Consider: it is not too demanding for me to give \$100 to a well-run charity. But if there are one hundred such charities, the requirement to donate \$100 to *each* well-run charity would make morality too demanding given my financial situation and family responsibilities. Likewise, it may be too demanding to require of people that they share evidence in every single case where there is social pressure not to, even in cases where the matter at hand is of sufficient importance.

What this suggests is that we should understand the duty to share your evidence as an *imperfect* duty. Immanuel Kant famously made the distinction between perfect and imperfect duties.⁵ The former yield determinate prescriptions, and don't allow for any discretion or latitude. For instance, you should not steal your neighbor's car so as to sell it and go on a nice vacation. It's just not to be done, period. Kant also thought we have duties of beneficence to others and duties to ourselves for self-improvement. But a duty of beneficence does not yield a determinate prescription. It can be discharged in many ways: by volunteering at a local cleanup effort, by being a helping friend, and so on. A duty of self-improvement can also be discharged in many ways. For example, if such a duty involves

getting in better physical shape, then you can do so by taking up running, cycling, swimming, or weight training. There's a lot of room for latitude and discretion.

In a similar vein, philosopher Jennifer Lackey develops a characterization of a duty to object (roughly, a duty to correct falsehoods) in terms of an imperfect duty. She writes:

Just as I ought to generally contribute to the moral flourishing of others so, too, should I do my part in ensuring that false and unjustified beliefs aren't promulgated. But surely I need not step in every time. If there were no discretion allowed, I would quite literally need to spend all of the hours in the day objecting to what is said on the news, in my Facebook feed, at the dinner table, and so on.⁶

Like the duty of beneficence then, the duty to object, for Lackey, is an imperfect one. Lackey contends that how much we are required to do in order to fulfill such a duty depends on two further things: (i) our social standing and (ii) what the others in our group are already doing. The greater our social standing (status, wealth, etc.) the more we are required to do. For one, people with higher social standing will have more of an impact with their speech. Further, they will often have lower costs—a tenured professor risks less than a graduate student, for example. Regarding (ii), if others are already doing a lot to object, in Lackey's sense, then a smaller burden falls on our shoulders. Thus, we're required to do less if others are doing their fair share. Compare: if lots of other people are giving to charity, there's a lesser need for us to give as much. But if relatively few people are giving, then the duty of beneficence will demand more from us in terms of charitable donation.

The duty to share evidence in the face of social pressure will have similar structural features. Sharing one's evidence in the face of every single instance of social pressure not to do so (even in cases which do not run afoul of the grandma's sweater problem) would be too demanding. Note also that since the duty I am positing is a duty to share one's evidence in the face of social pressure, so long as the costs are not too great, it accommodates Lackey's qualifications automatically. If lots of other people are sharing a particular piece of evidence, then almost *ipso facto* the social pressure against doing so will be small. But costless sharing of evidence is not what the duty amounts to. I have no duty to say that the Earth revolves around the sun now: there's no cost for me in doing so. But one might have that duty circa the late 17th century. Second, since costs matter, tenured professors, for example, will often have more of a duty to share certain pieces of evidence than graduate students who might risk unemployment by doing so.

Finally, a third problem the proposal outlined above faces is that it's often wise to keep the powder dry. The phrase originates from the 17th century when English general and statesman Oliver Cromwell advised his soldiers to keep their gunpowder dry. If you got your gunpowder wet back then, you couldn't fire when needed—and so it was important strategically to mind your resources and be well prepared. Analogously, if you never take heed of social pressure and go against the grain all the time, you might develop a reputation as a contrarian even if you are not. Such a reputation can “poison the well,” and cause people not to take you seriously. Depending on the context, for example, such behavior might signal to others that you are part of the outgroup. And empirical evidence suggests that people are prone to quickly discount the

testimony of outgroup members. Sunstein explains: “If people seem to be from some group we distrust or dislike, or a kind of ‘out group,’ they are far less likely to influence us, even on the simplest questions. Indeed, we might say or do the very opposite (‘reactive devaluation’).”⁷

Thus, one must be somewhat forward-looking in order to fulfill one’s duty to properly maintain the epistemic commons. The latitude and discretion characteristic of Kantian imperfect duties include a *strategic* element here as well. Of course, what strategy is appropriate will be heavily context dependent, in a way that makes any general recommendations otiose.

SHARING EVIDENCE VS CONTRARIANISM AND TROLLING

If Sunstein is right, then conformists are free riders: hence, their behavior leaves something to be desired, morally speaking. But on the other end, a pure contrarian cannot be a paragon of virtue either. Recall that the pure contrarian disagrees with others simply for the sake of disagreement. Perhaps he gets a kick out of disagreeing with others, and thus disagrees simply for that sake. Such characters are, obviously, rare. Most people want to fit in with others. But nonetheless, the contrarian is a possibility we should keep in mind. The virtuous person in this regard, it would seem, lies in the mean of the two extremes—conformism on the one end and contrarianism on the other.

There is another related but perhaps more mischievous character: the troll. The troll is someone who enjoys getting people riled up. Trolls are common particularly within anonymous online discourse, and if they’re defending fashionable opinions, even non-anonymous online discourse. A troll will

say something he knows will provoke a strong emotional reaction from others—in particular anger, frustration, or disillusionment. Moreover, he says what he does in order to provoke that reaction. His goal is thus not to arrive at the truth on some matter or improve the discourse in some way.

It would be bad if the duty I have been positing so far entailed that such characters are doing something good. Fortunately, it does not. The duty is a duty to share one's evidence in the face of contrary social pressure. What is evidence? Philosopher Thomas Kelly explains: "Intuitively, one's evidence is what one has to go on in arriving at a view. Evidence is what Sherlock Holmes carefully collects and surveys, and that from which he ultimately infers the identity of the person who committed the crime."⁸ Rational thinkers arrive at their beliefs by properly basing them on the total evidence they have. This means that evidence can come in many shapes and sizes—it might involve direct perception, arguments, datasets and statistical analyses, or the testimony of others, among other things. These can all be proper bases for forming particular beliefs.

To fix ideas, suppose the issue at hand is crime and policing. For example, we might be interested in whether increased police presence within an area reduces violent crime. What's the relevant evidence here? Well it might include things like crime statistics in various neighborhoods along with data about police presence. It may also include first-hand accounts of individuals in various neighborhoods.

Now suppose there is social pressure within your reference network against giving evidence for a particular view about policing and crime. It is here that the *prima facie* imperfect duty to share evidence kicks in. Importantly, the evidence has to be actual evidence. Contrast this with what a troll might do: he

might anonymously assert an opinion that has no connection to his evidence. He may not have looked at any of the statistics or first-hand accounts; he'll simply say whatever he knows will rile others up. Similarly, a contrarian may just engage in fanciful hypotheticals without having any evidence to back them up. But crucially, these characters are not sharing their *evidence*—namely, whatever it is that would bear upon whether a view about crime and policing is justified. And that's what the duty described in this chapter recommends.

GOOD FAITH

We can share genuine evidence in order to mislead others. Hence, imagine that a detective finds out three pieces of evidence during her investigation. A particular suspect, Jones, was seen entering the house around the time the murder took place. This is a reason to think that Jones is guilty. However, it was Smith's DNA that was found on the crime scene and murder weapon. In addition, Smith has no alibi, while Jones does. So, the overall evidence suggests Smith is guilty, not Jones. But suppose our nefarious detective harbors a grudge against Jones and so wants him to be found guilty. She thus reveals only the first piece of evidence.

This sort of evidence sharing is, paradigmatically, not done in good faith. The detective in this case intentionally makes the epistemic position of the prosecutor and jury *worse*—they were better off before, when they would have suspended judgment about who the culprit was. Now, they have been exposed to evidence, but evidence carefully curated so as to mislead. Plausibly, our duty to share evidence in the face of contrary social pressure cannot be one to share evidence so as to deliberately mislead. Rather, we must do so in *good faith*—where the

intention is to improve rather than to further deteriorate the epistemic situations of others.⁹

THE IMPERFECT DUTY GENERALIZED

The duty to share evidence in face of social pressure against doing so is a special case of the more general imperfect duty to improve the condition of our epistemic commons. In analogous fashion, the duty to give to charity (assuming the possession of means do so) is a special case of the more general duty of benevolence. Hence, someone with little by way of pecuniary means can nonetheless fulfil their duty of benevolence in other ways. They may not give money to charity but may nonetheless help others by being a supportive friend or family member, for instance. These aren't mutually exclusive, of course. An excellent way of fulfilling the duty of beneficence for someone who is well off might include donating to charity in addition to being a helping friend or family member.

Similarly, some of us may not be in a position to offer evidence against social pressure. Thus, someone might not have looked too deeply into issues like the minimum wage or crime and policing or abortion. Indeed, this will be true for many of us. We have other things to do with our time: work a nine-to-five job, raise children, exercise, mow the lawn, chat with friends, and much else. However, even in this case, we can do our part in helping to improve the epistemic commons: namely, by doing what we can to alleviate social pressures and taboos against sharing certain types of evidence or perspective.

There are a variety of ways to do this. For example, an academic might stand up for a colleague who is under fire for

publishing research that is deemed controversial. As I will discuss in detail later, even one voice of support can often make a huge social and psychological impact on others. Typically, this will involve some cost. There will be a good chance that some might treat such a person as being on the wrong “side”—and thus to be shunned and ostracized. And so, standing up for a colleague in this way might well have adverse social and professional repercussions. But of course, as emphasized earlier, morality often instructs us to sacrifice our narrow self-interest, so long as the costs are not excessively high.

Further, those who work in professions responsible for knowledge production and dissemination—academics, journalists, etc.—might take steps to promote intellectual diversity within their professions: for example, by hiring researchers or writers who can offer evidence that might mitigate blind spots. Recent empirical work suggests that there are benefits to both ideological and cognitive diversity when it comes to the output of teams. Teams composed of ideologically/cognitively diverse individuals produce better output (as measured by objective metrics) than ideologically/cognitively homogenous groups.¹⁰

In general, we might try to defend and uphold social norms in which people can share evidence in a relatively free atmosphere. This is not to say there shouldn't be robust disagreement. You might offer some evidence against claim X while I offer evidence for X. But the key is to promote social norms and structures wherein disagreement doesn't lead to ostracism or shunning.

This is particularly important when those who might share certain types of evidence are in a minority within a given reference network or milieu. Here, the majority have an enormous upper hand and a lot of leeway to intimidate the minority. Thus,

suppose that within a particular context, the claim X is held by the majority. Moreover, imagine the majority has some affective investment in believing X. So, it's not simply a "dry" matter to them—like some abstract theorem or obscure polymer. Rather, the matter is close in some way to their social identity.

Such a scenario can likely lead to an atmosphere which can easily prevent the countervailing evidence from surfacing. The majority now will be able to intimidate the minority in a variety of ways. For example, when someone from the majority levels an ad hominem attack against a member of the minority, he might receive little if any pushback. The majority, by stipulation, doesn't want to believe the contrary proposition. Of course, the members of the majority may not explicitly think "I don't want countervailing evidence to surface because it will undermine my justification for this belief that I am psychologically attached to." The inner machinations of our minds are rarely transparent to us. Rather, the temptation will be to see the dissenter as evil or the evidence as fabricated, even when these things are not true. The opportunities, however, for being able to level ad hominem attacks will not be symmetric. If a member of the minority does this, she will be quickly called out by the majority—and rightly so. Ad hominem attacks are not how virtuous and mature people conduct inquiry. But the problem is that the ability to conduct unjustified "warfare" in this way will be profoundly asymmetric.

Another method of intimidation might involve the spread of unjustified rumors. These rumors will spread easily within the community because of what the majority wants to believe. If the dissenters are bad people, there's no reason to take them seriously. On the other hand, unjustified rumors will have a much harder time getting traction in the other direction.

Asymmetric intimidation can also result from uncharitable interpretations and “strawmanning.” The majority will be in a position to strawman the dissenters’ views, that is, present them as much less reasonable and well-argued-for than they really are. Any attempt by the minority to strawman the majority position, however, will be quickly (and rightly) met with opposition. The same goes for things like public censure, the imposition of professional costs, and so on. While the majority will be able to publicly censure or impose professional costs on the minority, the reverse will not hold.

Here, I am just highlighting a few of the ways in which the ability to impose costs will be asymmetric when there is an ideological majority within a social or professional context. Many of these things work on subtle and subconscious levels—on which, more later. The point of importance here is that one way in which individuals can take steps to improve the epistemic commons, even if they belong to the majority, is to try to curb the majority’s tendencies towards such behavior. Thus, for example, someone might resist the temptation to use ad hominem attacks, and further, discourage others from employing such attacks, even if they’re friendly to his “side” of the issue. Often, the urge is to do the exact opposite—to praise or applaud ad hominem attacks when they align with the majority opinion. The rewards are feelings of righteous indignation (if one shares the majority opinion) as well as social status and approval. But again, morality often requires us to forgo such rewards.

In a passage worth quoting at length, John Stuart Mill emphasized similar worries about our behavioral tendencies when it comes to lopsided opinion, regarding which there is widespread affective investment:

The worst offence of this kind which can be committed by a polemic, is to stigmatise those who hold the contrary opinion as bad and immoral men. To calumny of this sort, those who hold any unpopular opinion are peculiarly exposed, because they are in general few and uninfluential, and nobody but themselves feel much interest in seeing justice done them; but this weapon is, from the nature of the case, denied to those who attack a prevailing opinion: they can neither use it with safety to themselves, nor, if they could, would it do anything but recoil on their own cause. In general, opinions contrary to those commonly received can only obtain a hearing by studied moderation of language.¹¹

One way to help the epistemic commons, then, is to promote the “studied moderation of language,” in Mill’s words, as a social norm, regardless of one’s actual position vis-à-vis the issue at hand.

I want to emphasize here that ‘majority opinion,’ as it relates to social pressures, is not to be construed as relative to an entire country or the whole world. What ultimately matters for people’s behavior in this regard is their *reference network*. Something might be a minority opinion with respect to the country’s or world’s population but a majority (even overwhelming majority) position within a particular industry or social class or professional network. In fact, it can often take enormous courage to defend a view that has widespread acceptance in the general public, depending on one’s social and professional context.

The social costs for expressing opinions are heavily context dependent. What might take courage to say in a rural town may not take courage to say in a newsroom, and vice versa.

Within several religious communities across the world, it takes enormous courage to say that atheists should not be punished by law. But it takes no courage at all to say this within a typical university in the U.S.—it would be like saying the sky is blue or that grass is green.

Who counts as a dissenter and who counts as a conformist thus depends on the social context. Recall that this distinction is important if we think of conformism as free riding, in the way discussed earlier. If someone thinks and says the same things as everyone else in his reference network, he's not a dissenter even if his position is a minority position when it comes to the world at large. Rather, he's a conformist. For it's not the world at large that is able to impose social and other costs on him—instead, it is his reference network.

Taylor & Francis

EXPRESSING OPINIONS AND COURAGE

Sometimes, the ad hominem attacker on the side of majority opinion may even be lauded as courageous! But of course, an action that is lauded as courageous and one that is courageous can come apart—particularly where the courage in question is expressive as opposed to physical courage. Can a person who is widely lauded within his social circles for sharing an opinion be courageous in sharing that opinion?

He may be tempted to think so. What psychologists call moral self-enhancement is people's tendency to have a high moral opinion of themselves, especially in comparison to others. In particular, people often attribute to themselves high levels of moral virtues such as honesty and loyalty. There is a well-studied, high degree of irrationality involved here. Paradoxically, the average person judges themselves to be more morally virtuous (on a variety of dimensions) than average.¹²

Naturally, then, people will tend to think of themselves as courageous even as their expressions of opinion are lauded by their milieu.

But what is courage? The ancient Greek thinkers often framed the virtue of courage within the context of war. Aristotle famously analyzed the virtue of courage (or bravery) as he did all virtues: as a mean between two extremes. A cowardly soldier would be excessively afraid to be killed in war. A rash soldier would be excessively fearless, perhaps flying into battle when it's wise to wait. Such a person would be very rare, and "some sort of madman."¹³ A courageous soldier would be a proper mean between these two. He would be afraid of the right things at the right time—and he wouldn't be afraid *per se* of a noble death (according to Aristotle) in war.

But notice we can't apply the concept of courage in this way to a situation that presents no danger at all. One may well be courageous as a person and face a situation with no danger. A courageous soldier might be relaxing under a tree, in a meadow, during a time of total peace. But such a situation will not allow for the *manifestation* of the virtue of courage.¹⁴

Similarly, then, sharing an opinion that is widely lauded as courageous within one's relevant reference network paradoxically *cannot* involve courage—unless there is a penalty enforced from outside the reference network. For, being widely lauded as courageous is obviously an enormous social benefit. And what is the cost? In war, it may be death. But when it comes to expressing opinions in modern democracies, the costs to be borne are social—yet by stipulation, the person in question is being lauded, and hence receiving *benefits*. This of course, is not true in a case where most of one's reference network lauds us for saying X but doing so publicly might land the person in jail. Imagine, for example, a reference network of early

20th century independence activists working to overthrow a colonial power. Expressing dissent against the colonial power will elicit praise from the reference network, but doing so publicly will risk punishment from the colonial administrators. However, this is typically not the situation we find ourselves in within modern democracies.

Manifesting courage as it relates to sharing our evidence, then, by its nature implies taking risks of being frowned upon, or worse, ostracized, by one's social network. Yet of course, courage can't be had cheaply. It is the very essence of the virtue of courage that it doesn't come painlessly. As Aristotle says, "standing firm against what is painful makes us call people brave; that is why bravery is both painful and justly praised."¹⁵ Now, courage may well be praised by one's reference network when it comes to war—the enemy there is external to that network. But the "pain" or cost involved in sharing our evidence, within the context of modern democracy, is typically imposed by the network itself. So we can't be widely lauded by our social network as courageous while at the same time actually displaying courage. Can't have your cake and eat it too.

MARGINAL VALUE AND HETERODOX RESEARCH

Economists and decision theorists model rational decisions as being made at the margins.¹⁶ When consumers buy goods, they base their decisions on the marginal cost and the marginal benefit of the n^{th} unit. Consider, for example, the purchase of cups of coffee. The first might give you more benefit as compared to the cost (let's say it's \$3). The second might still benefit you more than \$3. But the benefit of the third may be below \$3—or, the alternate use to which you could put that \$3.

You've already had two cups and gotten a good dose of caffeine. You might have some decaf if it was being sold for 50¢, but at its current price, it's just not worth it.

In this way, the marginal benefit of a typical good decreases for us, while the cost remains the same. The coffee is going to cost \$3 regardless of how many you buy (typically) but the benefits of the third cup of coffee are far less than the benefits of the first cup. You stop buying the good when the marginal cost exceeds the marginal benefit—in this hypothetical case, you stop at two cups. The decision, if made rationally, is made at the margins. The question is not “how much are three cups of coffee worth to me, rather than zero?” Three cups of coffee may well be worth \$25 to you if the alternative is zero cups. You've just got to have that daily caffeine fix, let's say. But this doesn't mean that you will buy three cups since they cost \$9, and 25 is bigger than 9. Rather, in making the decision as to whether to buy the third cup of coffee, you ask yourself “how much is the third cup worth to me?” And here, the answer may well be: less than the \$3 that it costs.

Similarly, the third bag of oranges is way less valuable to an individual consumer than the first. Same goes for the third car, the third computer, the third house, the third TV, etc. Of course, there's nothing special about the number three—rather the point is that typically, the value of the marginal unit of a good decreases as the amount of that good we have increases. Thus, the tenth TV is even less valuable than the third for us. Unless one is a filthy rich person with a mansion, she typically won't buy ten TVs or ten cars.

Marginal analysis applies to scenarios beyond just individual purchasing decisions. If you're studying for an exam, the first hour of reviewing the material is much more valuable than the 13th. It might be wise overall to spend that first hour studying

rather than hanging out with friends, but by the time you've already studied for 12 hours, it could be a much better use of time to relax and talk to people. Likewise, the first half hour of exercise per day is much more valuable than the next half hour.

The same sort of analysis applies to decisions involving what economists call *public goods*. Public goods are goods that are non-excludable and non-rivalrous. Clean air is a paradigm public good. We can't exclude others from enjoying clean air. And clean air is non-rivalrous in the sense that my enjoying the clean air doesn't affect your ability to enjoy the clean air. Contrast this with a cup of coffee, which is a classic case of a private good. There, you can exclude me from drinking from your cup. And, my drinking your coffee means there's less remaining for you—so it's rivalrous.¹⁷

The notorious problem with public goods is that they throw our incentives way out of whack. Consider roads. A road is a public good (unless it's a toll road). Everybody benefits from having roads. However, nobody has an incentive to pay for them voluntarily. Let's say that for the road to be built, each person in the town has to contribute \$100. Now, the road may be worth way more than \$100 to each individual. But this is not the relevant question. The relevant question each individual will consider is: "how much benefit will I receive from contributing \$100 to the road project?" Here, the answer will be: very little. To see this, note that adding that money to the project might make the road a quarter of an inch wider. But it's not worth it to you to spend \$100 to make one road in your town a quarter inch wider. You might rather put that money towards a comfy chair or a nice dinner—those things would add more to your life-satisfaction. The problem is that everyone will think this way, and so nobody will contribute to the road fund.

But the issue is that *everyone* would be better off if they all contributed to the road fund. After all, the road is worth more than \$100 to each of them, as stipulated. Thus, there are gains to be made, from everyone's perspective, left on the table. Enter: the government. The government can make use of coercion to create public goods like roads. It coerces by means of taxation. People don't have a choice as to whether to pay taxes—they can't thus voluntarily decide not to contribute the \$100. But, paradoxically, by making the choice to pay taxes involuntary, the government makes them better off.

The provision of public goods in this way has been considered to be one of the main justifications for government. Adam Smith thought the three jobs of government were: (i) provision of national defense; (ii) provision of a justice system; and (iii) "erecting and maintaining those public institutions and those public works, which though they may be in the highest degree advantageous to a great society, are, however, of such a nature, that the profit could never repay the expense to any individual, or small number of individuals; and which it, therefore, cannot be expected that any individual, or small number of individuals, should erect or maintain."¹⁸ Notice that all these functions can be subsumed under the idea of public good provision.

However, the presence of the government as a tool for public good creation doesn't make all the problems associated with public goods go away. Should the next road be built? Rational decision making from the standpoint of the community would consider the total cost of the road and compare it with the benefit the whole community reaps from having the road built. The thing to do then is to keep building until the former quantity is greater than the latter.

But the issue is that communities don't decide to build particular roads. Administrators within the government do. And the rational thing to do for an administrator may come apart from doing what is best for the community. This can cause the under-provision of some public goods, and the over-provision of others. Ideally, of course, a government administrator would only take actions that are to the benefit of the community. But people are not ideal—sometimes they favor themselves or their friends at the expense of the average community member. Within democracies, voters have some control over their administrators, but such control is not perfect—it goes via the channel of representative government, rather than direct democracy on every particular decision to be made. Such is the *principal-agent* problem in government administration.¹⁹

Now, academic research is a public good—non-excludable and non-rivalrous—and thus presents people with similar kinds of incentives as roads. Without some kind of external funding, it would almost never be worth it for any researcher to do their work. Imagine a medical researcher who is looking to study the potential benefits of a particular therapy for a kind of cancer. To work on this, she will need millions of dollars—to support herself, her research assistants, purchase lab equipment, etc. But her research will never really “pay for itself,” as far as her direct interests are concerned. The chance of her getting that particular kind of cancer is rare. Maybe she might write a book for popular audiences and recuperate some money that way, but it's likely not going to cover all the costs. So, if she were left to her own devices, it would never be worth it for her to spend years and years working on the problem. Nevertheless, her research may well be tremendously

valuable for the world as a whole—if she succeeds, she may save many lives. Even if her chance of finding a cure is 50/50, say, it might be worth it for society to fund her work, depending on how prevalent the cancer is. Thus, to get out of this predicament, researchers need external funding—which typically comes from governments, private donors, and tuition paying students.

But now, notice that academic research is only beneficial for society if its marginal benefits exceed marginal costs. I want to emphasize that ‘marginal benefits’ here is to be construed broadly—the benefits can include both the instrumental value of research (building new technologies, finding cures, etc.) as well as its intrinsic value (knowledge itself being the good in question). Presumably, research within my own field of philosophy is largely justified in the latter way: philosophers don’t typically discover how to cure cancer or make smartphones or spaceships.

Clearly, not all possible research is worthwhile in the sense that the marginal benefits (broadly construed) exceed marginal costs. Here’s a possible research program. The aim is to count exactly how many blades of grass there are within a large pasture in Montana. Let’s say this will take a total of 1000 people working for one day, for a total expense of \$500,000. Should society fund this? Of course not, it would be absurd. It’s not worth spending that amount of money to figure out how many blades of grass there are in this pasture. Resources are scarce, and we should use them in better ways.

Imagine now that Bill, who is a researcher, has received a check from the government and private donors for \$500,000. He could have refused the check, in which case the money would have gone to the representative sorts of things that governments and private donors fund: poverty relief, public works,

art galleries, etc. Now that Bill has gotten the money, there are no strings attached, however. He can spend this money however he likes, provided it's on research. So of course, he can't just spend it on a Rolex and a yacht.

Suppose Bill can choose one of two programs. One program involves the grass counting in Montana. The other program involves looking into the environmental effects of large-scale farming. Stipulate, for simplicity, that Bill has the right training and other resources to successfully conduct either research program. It would be wrong for him to choose the former project. He'd be wasting money that could be much better used elsewhere. It seems he has a duty to pursue the latter program, even if it takes more effort, say. Part of research ethics then involves seeking projects that have a sufficiently high marginal value for society—ideally, higher than the marginal cost that society bears to fund said research.

This hypothetical is relevant to modern researchers, because we often find ourselves in exactly such a situation. For instance, consider my incentives. I receive a salary that is not conditioned upon the content of what research I pursue. I'm free to do whatever I like. Of course, I don't have tenure yet, and in order to get tenure, I will have to publish work with reputable journals and presses. But other than that, there are no content-based constraints. Now, this is the model of the humanities. The social and physical sciences work differently—the research there is often grant funded, which largely involves government agencies like the National Science Foundation awarding money for specific research projects. However, researcher professors themselves decide which projects get funded through these agencies—not on their own behalf, of course, but on behalf of others in their fields. I contend that just like Bill, individuals like me, as well as people deciding

which grant applications to fund, have a duty to see to it that the funded research is expected to have the greatest (or at least, sufficiently great) benefit at the margins relative to the cost.

Sometimes this might involve personal costs, though. Working on the highest value research might involve more effort or it might mean the work is harder to get published. From a narrowly self-interested perspective, of course, the thing to do for researchers is to work on the projects which yield the most publishable papers in top journals, most opportunities for better jobs and grants etc., while requiring the least effort. If that involves counting grass, then count grass! Nonetheless, as I have been emphasizing, morality often instructs us to act contrary to our narrow self-interests, so long as the costs are not too high.

Now, in a field of inquiry that is working well, individual researchers will be incentivized to pursue those projects which in fact have the highest marginal value. Indeed, proposals to count blades of grass won't stand much chance of being funded by the National Science Foundation.

Within philosophy, I mentioned metaethics as field that is paradigmatically working well. I asked you to imagine a hypothetical situation where there were 100 naturalists and one non-naturalist in the field. Such a hypothetical situation would be unstable. Presumably there would be lots of good arguments to make for non-naturalism, which that one person hasn't formulated. This would incentivize philosophers to publish articles defending non-naturalism—there will be “low hanging fruits” to pick. And picking these fruits will be rewarded with publication in the best journals in the field. This in turn will lead to higher status and better job opportunities, and so on. A situation where only one person defends

non-naturalism, then, won't last for long—in game theoretic terms, it would not be an equilibrium.²⁰

Indeed, as noted previously, a wide variety of positions exist within the modern profession of metaethics, with each position having several high-profile defenders. Insofar as the total amount of money spent by society to support metaethics is warranted, then, the incentives of researchers align with the good of society. Individual researchers are incentivized to seek out projects with high marginal value. Similar points presumably apply to physics and chemistry. Researchers work on a wide range of topics and are constantly testing various hypotheses. A hypothesis that can easily be refuted won't stand unrefuted for long. A new, useful polymer that can easily be studied won't remain unstudied for long. This is not to say these three fields are working perfectly, but rather just to say they approximate the ideal.

It's helpful in this context to think of fields of inquiry as consisting of individuals harvesting a tree bearing fruit. A healthy field of inquiry will involve individuals being free to pick whichever fruit they can get their hands on. Thus, the lowest hanging fruits will be gone quickly. Eventually, the only fruits left will be high up in the tree. In contrast, an unhealthy field of inquiry might be akin to a tree where there are obstacles to picking fruit on a particular side of the tree. On such a tree, individuals will be picking fruit at ever-higher locations on one side, but on the other side, there will be fruits hanging very low, which individuals are disincentivized from picking. Yet if fruits are the goal here—as truth is presumably the goal of inquiry—then the marginal benefit to society of trying to pick the lowest fruits will be very high. Doing so won't require expensive ladders, to stretch the metaphor a bit further.

What should individual researchers do within fields where low hanging fruit are “forbidden” in the sense that picking them invites social and professional costs? Are there any such fields today? In the passage quoted in the previous chapter, Glenn Loury expresses his worry that some areas of social science are structured in this way. The worry is that there are some conclusions that most members within some fields want to reach, and thus there are social costs associated with conducting research whose output undermines, or provides evidence against, those conclusions.

However, if the analysis above is correct, there will be low hanging fruit to pick here *precisely because* people have been disincentivized from conducting such research. But if the work of researchers is ultimately justified in terms of what they add to the knowledge stock of humanity, as opposed to merely the status and prestige they gain for themselves, then it would seem that individual researchers have a duty to pursue such “heterodox” projects. For, the marginal benefit they provide to society is likely to be the highest if they do so.

Of course, as I have been emphasizing, the duty is a *prima facie* duty—if the costs of pursuing such research are prohibitive, then plausibly the duty fails to apply. Furthermore, the costs will often depend on the context, and in particular, the professional position of the researcher. Thus, imagine a scenario where pursuing some research projects, however well, makes Ph.D. students un-employable. Now consider a young Ph.D. student deciding what project to pursue. Morality would be too demanding if it required her to pursue research that would destroy her career.

But note that typically such costs decrease as one advances in one’s career. If and when this individual gets tenure within a research university, for example, the costs of pursuing the

research program in question will drop dramatically. Now that her living is secure, the brunt of the costs will involve things like social sanctions from some of her colleagues, poor luck in some of the field's top journals, etc. These costs are real—they can be unpleasant and might decrease her status. But morality sometimes demands us to do things at some cost to ourselves.

A classic historical example of someone who fulfilled this duty—or rather, went above and beyond—is Galileo Galilei. Galileo argued for the heliocentric model of the world, according to which the Earth revolves around the sun, rather than vice versa. But this was at a time in the 17th century when the Catholic Church was heavily invested in defending the opposite. Thus, the issue was not a “dry” one which people could discuss and debate openly, without cost. Ultimately, the Church found grounds for ‘vehement suspicion’ that he defended the Copernican heliocentric model, thereby committing heresy. His *Dialogue Concerning the Two Chief World Systems* was banned and Galileo was condemned to house arrest for life. In this way, Galileo dramatically improved our understanding of the world at cost to himself.²¹

Which research invites social and professional sanction has changed over the years. We don't nowadays get worked up about what physicists find. The work of modern physicists, though fascinating, is dry in this sense. Nobody is going to get fired or put under house arrest for discovering and publishing something about black holes or neutrons. But the worry expressed by Glenn Loury, among others, is that this “dryness” is absent in some fields within social science. Defending some hypotheses, however well, invites censure, ad hominem attacks, and professional costs, according to this worry. The idea would be that as the Catholic Church was invested in

particular hypotheses being true with respect to the physical world, so are the majorities within today's institutions of knowledge production invested in certain hypotheses regarding the social world.²² Insofar as Loury is right, it is here that modern-day Galileos are to be found. Let me illustrate with a hypothetical example.

In the 20th century, there was a lively macroeconomic debate between the Keynesians and the Monetarists. The Keynesians believed, among other things, that during recessions, governments ought to pursue expansionary fiscal policies to stimulate the economy. Monetarists thought that such policies had no positive impact on long term growth. Keynesians favored more flexibility for central banks, while Monetarists favored more constraints and rules on central bank actions. Each side gave various arguments and evidence for their conclusions. Today, macroeconomists draw from the insights of both sides to construct their models.

But imagine the following scenario. Imagine that in a distant possible world, the economics community was heavily invested in Monetarism being true. Journal article after journal article defended Monetarism. Monetarists received prestigious awards and university posts. But Keynesianism was verboten. A Keynesian could get an article published here and there, by stroke of luck, but it was an uphill battle. Graduate students defending Keynesianism found themselves with few job prospects. Defending Keynesianism provoked ad hominem attacks of the kind Loury describes.

What are the marginal benefits of pursuing research in such a scenario? Well, note that a good typically gets less and less valuable the more of it we have. The first article defending Monetarism is extremely beneficial. But, assuming similar

quality, the 1000th article defending Monetarism will not add nearly as much. The low hanging fruit are to be found in defending Keynesianism. Moreover, given such incentive structures, our general epistemic position with respect to monetary policy will be impoverished. We just won't know if the Keynesian insights have some truth to them—because people are strongly disincentivized from voicing them!

In such a scenario, then, it is all the more important to pursue work within the Keynesian research program. By publishing the second article defending Keynesianism, you would do much more good as compared to writing the 1000th article defending the Monetarist view. In addition, even if you are a Monetarist in such a situation, you should help reduce the social and professional costs borne by the few who are defending Keynesianism, in whatever ways you can. This might involve doing your best to seek out unbiased referees for Keynesian papers as a journal editor, attempting to reduce ideological bias in hiring decisions, and so on.

The point doesn't only apply to social science, of course. Similar conclusions may be drawn with respect to the humanities. Within philosophy for instance, many careers, books, and journals are devoted to inquiry about the basic principles of justice and their application to contemporary issues. Much work in the humanities more broadly touches on social issues that are contentious and alive today, regarding which people have affective investment. But insofar as a community of academics wants to reach a particular conclusion, the worries Loury raises will be relevant. In such cases, one provides much more by way of marginal benefit by defending the heterodox side rather than piling on with the n^{th} paper or book defending the popular conclusion.

CONCLUSION

We have a duty to preserve the health of our epistemic commons. One chief way to do so is to speak our minds despite social pressure. Social pressure distorts the evidential landscape, and so threatens the health of the epistemic commons. Hence, when we reveal our evidence to our community against countervailing pressure, we might be curing a dangerous blind spot. The duty is particularly relevant to researchers and intellectuals, given the influence they have and the social role they play. They can fulfill this duty by pursuing and encouraging others to pursue heterodox research and ideas, where there is pressure from the mainstream (relative to the context) against doing so.

An important worry this discussion raises is the question of whether individuals can make a difference. Why risk social status if all we say is going to be a proverbial “drop in the ocean” anyway? This is the topic of the next chapter. I will argue that you can often make a huge difference.