



EPISTEMIC DISAGREEMENT, INSTITUTIONAL ANALYSIS, AND THE FIRST
AMENDMENT STATUS OF LIES

*Mark Tushnet**

(from the Knight Institute’s Lies, Free Speech, and the Law symposium**)

Introduction 652

I. Preliminaries: The Institutions of Interest 653

II. What Is a False Statement of Fact? Herein of Facts, Opinions, and the
Difficulty of Determining the Content of Factual Assertions 655

A. Facts and Opinions 656

B. What Does a Factual Assertion Mean? How Meaning and
Normative Assertions Are Intertwined..... 660

C. Some Generalizations 662

III. Is There a Constitutional Interest in the Dissemination of False
Statements of Fact and/or Lies? The First-Order Analysis..... 664

A. The Core Analysis 664

B. Conclusion on First-Order Reasons 669

IV. The Second-Order Analysis: Can Legal Institutions Reliably Distinguish
Between Mere Falsehoods and Lies? 669

A. The Institutional Analysis of Regulation of Mere Falsehoods and
How It Can Be Extended to Deal With Lies 670

* Mark Tushnet is William Nelson Cromwell Professor of Law emeritus, Harvard Law School.

** This Article is being reprinted essentially as it was published on the Knight Institute site, and adheres to the Institute’s stylistic preferences. More information about the Lies, Free Speech, and the Law symposium is available at <https://knightcolumbia.org/events/symposium-lies-free-speech-and-the-law>.

B.	A Sketch of an Institutional Analysis of When the Risk of Mistaken Inferences About Belief Will Arise	676
1.	Legislation.....	677
C.	Administrative Agencies	680
D.	Forms of Liability	681
	Conclusion.....	682

INTRODUCTION

According to Justice Powell’s opinion in *Gertz v. Robert E. Welch, Inc.*, “[T]here is no constitutional value in false statements of fact.”¹ This is a claim about what we can call first-order free expression interests, the values both individual and social of the dissemination of statements.² The first step in my argument is that the first-order claim requires substantial analysis, and that, though there might be no social value in the dissemination of a false statement of fact *with respect to its content*, a Millian argument of a certain sort shows that the first-order claim might be mistaken when *other* individual and social interests are taken into account. But, I argue, a different analysis is required when we come to lies, defined as false statements of fact known or believed by the speaker to be false.³ Like mere falsehoods, lies might not have social value with respect to their content, but the Millian argument that supports the conclusion that there might be social value in the dissemination of falsehood doesn’t support the conclusion that there might be such value in the dissemination of lies.

The next step in the argument turns to second-order concerns, mostly about lies but with implications for the analysis of mere falsehoods. Second-order analysis deals with the institutions we have for implementing the rules regarding first-order individual and social interests. It asks whether those institutions have characteristics that allow them to generate results that are reasonably reliable in determining when the first-order interests will be promoted or impaired by regulation. Second-

¹ *Gertz v. Robert E. Welch, Inc.*, 418 U.S. 323, 340 (1974).

² For important recent discussions of falsity and lies in connection with free expression, see Frederick Schauer, *Facts and the First Amendment*, 57 UCLA L. REV. 897 (2010); SEANA VALENTINE SHIFFRIN, *SPEECH MATTERS: ON LYING, MORALITY, AND THE LAW* (2014).

³ I quite briefly address as well whether the definition should encompass statements made with reckless disregard of their truth or falsity.

order concerns, I argue, support the conclusion that broad bans on the dissemination of lies should be viewed with great suspicion but that bans targeted at well-defined, quite specific lies shouldn't be seen as violating free expression principles. The principal second-order concern is the possibility that juries in particular (but other decision-makers as well) will wrongly infer from a statement's evident falsity that it must have been made with knowledge that it was false.

This argument has significant implications for First Amendment doctrine. For example, it suggests that *United States v. Alvarez* was wrongly decided because it failed to recognize that the second-order concerns it properly identified in connection with a "Ministry of Truth" were inapposite with respect to a statute prohibiting someone from lying about having received a military honor.⁴ The argument suggests that a statute creating a Ministry of Truth charged with identifying specific falsehoods that, if disseminated with knowledge of their falsity, would be constitutionally problematic because of the bureaucratic incentives the ministry would have to find something to do.⁵

I. PRELIMINARIES: THE INSTITUTIONS OF INTEREST

This essay deals with several institutions that might be charged with regulating falsehoods or lies. I assume throughout that these institutions are located in systems of governance that are reasonably well-functioning but flawed democracies, where the flaws are shortfalls from the system's own understanding of democracy's core characteristics.⁶ These institutions include: (1) *legislatures*, which can enact statutes targeting (a) specific falsehoods or lies, like the Stolen Valor Act, which made it a

⁴ *United States v. Alvarez*, 567 U.S. 709 (2012). The decision may have been correct given the ways in which the government defended the statute's constitutionality. The Supreme Court assumed that upholding the constitutionality of the Stolen Valor Act would imply that a statute penalizing the dissemination of any lie would be constitutional. That assumption might have flowed, arguably correctly, from the government's position that liability for falsehoods could be imposed only when those falsehoods worked something akin to material harm *and* that the impairment of the incentives created by the system of military honors counted as a sufficient harm because any lie could have effects that would count as "material" under the government's analysis.

⁵ A problem made familiar in the First Amendment literature by Thomas I. Emerson's observation that the job of a censor is to censor, implying that censorship boards have incentives to find things to censor that legislators and judges, with much more on their plates, do not.

⁶ Whether the United States is (or remains) such a democracy is of course contested, and I take no position on that question. I do note that after a certain point, democratic decay may have gone so far that it would be impossible to enact new regulations of falsehoods and lies.

crime falsely to claim having received a military honor; (b) general classes of falsehoods or lies, such as statutes creating commercial fraud liability for disseminating false and misleading information in connection with consumer products; or (c) all falsehoods; (2) *administrative agencies*, which can be charged with (a) identifying falsehoods or lies and imposing liability for their dissemination, as with the Ministry of Truth that so concerned Justice Kennedy in *Alvarez*; or (b) imposing liability for harms that aren't defined by the presence of falsehoods or lies but can sometimes arise because of their presence (Consider for example the Federal Trade Commission, which can impose liability for harms arising from creating a monopoly *or* for harms arising from false and misleading information about the risks associated with a product which causes consumers to purchase goods or pay prices that they otherwise wouldn't. I will sometimes refer to this as a "special purpose" agency.); and (3) *ultimate fact-finders*, such as judges and juries charged with determining that a specific utterance was a falsehood or lie.

Each of these institutions has different incentives with respect to the choices available to it. Legislators respond to electoral incentives, including the prospect of gaining votes from constituents or raising campaign funds from donors. In addition, they have many things on their plates. Whether we should be concerned about the possibility that legislators will enact any of the three types of statutes I've identified will depend upon our evaluation of their incentives and workloads.

A Ministry of Truth is a permanent institution with a single charge, and the usual analysis is that such an institution will go out of its way to find work to do—that is, it will seek to identify "enough" falsehoods or lies to justify its continuing existence. In contrast, a special purpose agency focuses on the harms it is charged with averting, and only incidentally upon the subclass of cases in which those harms are associated with the dissemination of falsehoods or lies. At the comparative level we might expect special purpose agencies to identify fewer falsehoods or lies than a Ministry of Truth would. And to the extent that constitutional concerns about imposing liability for such dissemination turn, as they sometimes will, upon the sheer mass of regulated falsehoods or lies, we might be less concerned about liability imposed by a special purpose agency than by a Ministry of Truth.

What of ultimate decision-makers? Juries come together to make a single decision and, our system assumes, follow their instructions closely enough that their incentives are simply to get the correct answer to the question, Is this statement a falsehood or lie? They are, that is, mostly responsive to "the law," and so draw on

their personal experiences to decide whether the statement is a falsehood or a lie, though perhaps their view of what the law is might be influenced by an unarticulated evaluation of the person making the statement or by the statement's content.⁷

Judges are a bit different. Our system assumes that, like jurors, judges are largely responsive to "the law's" requirements, with the same qualification I mentioned as to jurors. Like legislators and special purpose administrative agencies, judges have a wide range of tasks to perform. Only occasionally will a judge be tasked with determining whether a statement was true or false; in bench trials, they will of course decide who's telling the truth, but they won't see the entire universe of falsehoods and lies. Finally, unlike legislators and administrative agencies, judges must wait for someone to come to them with a case in which determining whether a statement was a falsehood or lie is legally relevant; they do not have a "roving charge" to seek out falsehoods or lies. These institutional features might reduce the number of falsehoods or lies that become subject to judicial scrutiny.

The takeaway point here is that we shouldn't talk about the "regulation of lies" in general but rather should focus on the characteristics of the specific institutions charged with such regulation (and as I discuss below, on the types of sanctions the institutions are authorized to impose).

II. WHAT IS A FALSE STATEMENT OF FACT? HEREIN OF FACTS, OPINIONS, AND THE DIFFICULTY OF DETERMINING THE CONTENT OF FACTUAL ASSERTIONS

Lies are false statements of fact disseminated by a person who knows (or believes) the statements to be false.⁸ But what is a false statement of fact? Philosophers

⁷ This qualification is relevant, for example, in connection with the imposition of liability for intentional infliction of emotional distress when one of the elements of the tort is a statement's outrageousness. The Supreme Court's decision in *Snyder v. Phelps*, 562 U.S. 443 (2011), appears to have been influenced by this concern. *See id.* at 458 ("'Outrageousness,' however, is a highly malleable standard with 'an inherent subjectiveness about it which would allow a jury to impose liability on the basis of the jurors' tastes or views, or perhaps on the basis of their dislike of a particular expression.'") (citation omitted).

⁸ For present purposes, "knows" simply means that the speaker has a quite high degree of confidence in the proposition that the statement is false, and "believes" means that the speaker's confidence level isn't that high. For completeness, I add the following: Suppose a person believes something to be true but insincerely disseminates a statement that it is false (the person believes that Hillary Clinton went to Harvard Law School but insincerely disseminates the statement that she went to Yale Law School) or believes something to be false but insincerely disseminates a statement

have discussed the distinction between truth and falsity for centuries, developing extremely complex accounts, some inconsistent with others. The legal system can't "rely on" some well-established philosophical account of what makes a statement false because there isn't one.

We need an analysis of what a falsehood is for purposes of legal regulation. I put the point that way because there are many purposes for which we might want to distinguish between truth and falsity—for assessing the character of a person making a statement, for deciding how to invest our money, and of course, many more. The legal system has institutional characteristics, of the sort described in Section II, that might provide usable boundaries around the distinction between truth and falsity—or at least that's the premise of what follows.

Two problems help frame my discussion of how we can identify false statements of fact: the distinction between facts and opinions and the treatment of so-called memory laws such as laws banning the dissemination of assertions that the Holocaust didn't occur. Each problem exposes real difficulties in developing a legal regime for regulating falsehoods—and perhaps for regulating lies.

A. Facts and Opinions

An important distinction in First Amendment law—including the law governing the regulation of libel—is the distinction between false statements of fact and "false" opinions. A central assumption of free speech law is that liability can't be imposed for disseminating false opinions.

Identifying what are statements of fact and what are statements of opinion is not always a simple matter, however. Simply labeling a statement an opinion can't immunize it from liability. If the statement "Donald Trump cheated on his taxes" is libelous, so is the statement, "In my opinion, Donald Trump cheated on his taxes." Whether a statement is a factual one instead depends upon a host of circumstances, including the statement's words and its context.

that it is true (the person believes that ivermectin is a useful treatment for COVID-19 but disseminates the statement that ivermectin isn't effective). We can argue about whether the person should be said to have lied, but for free expression purposes, the answer to that question is, basically, "Who cares?" For in both cases, accurate information has been let loose on the public, and there's no free expression concern (though there might be other concerns going to the person's character and the like).

Facts and opinions, it is commonly said, lie on a continuum.⁹ Identifying the metric, so to speak, for that continuum is notoriously difficult. A seat-of-the-pants definition would be that factual statements are those that are capable of being shown to be true or false. But as I've indicated, we can't look for criteria for determining whether a statement is "really" true or false in some transcendental sense; all we can do is come up with criteria for determining whether it is true or false for purposes of legal regulation. And translating "capable of being shown to be true or false" into a legal doctrine is quite difficult.

Consider several statements that present themselves as factual. (1) "There's some root beer in my refrigerator." (2) "The paper on which this essay is written is made up of atoms, which themselves are made up of electrons, protons, and neutrons, which are in turn made up of other kinds of subatomic particles." (3) "Our new product provides more effective and longer-lasting relief than our major competitor's product." (4) "Steph Curry is the best player in the NBA today." (5) "Time and time again the Republican program of cutting taxes on the wealthy has proven to be a motor for economic growth and the improvement of well-being for everyone in the United States." (6) "Corn dealers are robbers of the poor."

Though the first two "look" purely factual, and the rest blend factual assertions with words that look more opinion-like, I argue next that all the statements are capable of being proven true or false, and that the way in which that capacity manifests itself shows how difficult it is to come up with a legally tractable definition of factual statements.

(1) "There's some root beer in my refrigerator." We can prove this true or false by going to the refrigerator, opening it, and seeing whether there's some root beer there. Or can we? Suppose we do that and find no root beer there. Does that mean that the statement was false when made? Maybe not. Maybe somebody came into the kitchen and took the root beer out in the time we spent getting from the place where the statement was made to the kitchen. We can try to rule out this possibility by looking for clues indicating the presence of someone else (fingerprints on the refrigerator door, perhaps). In the end, though, we'll say that the statement was false when made when we think about how serious the factual claim is, come up

⁹ See, e.g., Michaele Sanders, *The Fact/Opinion Distinction: An Analysis of the Subjectivity of Language and Law*, 70 MARQ. L. REV. 673, 680 (1987).

with a list of techniques for verifying the claim, and use those techniques that seem appropriate given the claim's significance.

Suppose that we do find root beer in the refrigerator. Does that mean that the statement was true when made? Again, maybe not, because a parallel scenario might reconcile the root beer's presence with falsity when made. Of course, we're unlikely to investigate the possibility that someone sneaked in and placed root beer in the refrigerator because our usual experience suggests that the possibility is quite slim. Putting the two cases together, we can see that here we understand "capable of being proven true or false" to refer to the use of techniques that are pragmatically useful in helping us make decisions in daily life.

This point can be driven home by considering a slight variant: "When I looked in the refrigerator five minutes ago there was some root beer there." We assess this statement by considering first the speaker's general veracity insofar as we know it, then the possibility that the speaker isn't telling the truth this time, which we would usually rule out by assuming that speakers with this one's history of veracity conform to their history—in short, by trusting the speaker. Again, there's nothing hard and fast about this, just a bunch of judgments about what it's worth worrying about in connection with the statement.

(2) "The paper on which this essay is written is made up of atoms, which themselves are made up of electrons, protons, and neutrons, which are in turn made up of other kinds of subatomic particles." There are fancy and simple versions of how we can prove this true or false, both of which end up having the same structure.

(a) The fancy version is that we take the piece of paper to a physics lab, put it in some expensive atom-scanning equipment, look at the screen or other form of output, and see dots that we're told are images of the atoms (and similarly, but with even more equipment involved, for the claims about subatomic particles). It's pretty clear, though, that we aren't doing the same thing here as looking into the refrigerator. We are relying on the physicists who run the equipment and tell us what the dots mean. Brian Leiter refers to this as reliance upon the physicists' epistemic authority.¹⁰ We simply take their word for it because we think that they know

¹⁰ Brian Leiter, *The Epistemology of the Internet and the Regulation of Speech in America*, 20 GEO. J. L. & PUB. POL'Y 903 (2022). See also LINDA TRINKAUS ZAGZEBSKI, EPISTEMIC AUTHORITY: A THEORY OF TRUST, AUTHORITY, AND AUTONOMY IN BELIEF (2013) (offering a foundational philosophical argument, the quality of which I'm not competent to evaluate).

what they're talking about (and have no reasons to misrepresent what they believe their training allows to say about what their equipment shows).

(b) The simple version is that we've read a lot of articles about science in newspapers, magazines, and science classes, all of which present us with this picture of how the physical world is made up. Here the epistemic authorities are the authors and publishers of those articles.

I offer a somewhat more extensive analysis of epistemic authority in Section V, but for present purposes it's enough to say that we accept epistemic authorities (when we do) because doing so makes it easier for us to go on with our ordinary activities. Once again, pragmatic considerations dictate what we understand the practice of determining truth or falsity to be.

(3) "Our new product provides more effective and longer-lasting relief than our major competitor's product." We can prove this true or false by coming up with a list of criteria we associate with effectiveness and length of relief and asking users how much of each the two products provided. We aggregate the answers and see whether they support the statement.

Such surveys are common in cases dealing with claims like these, and the problems with them are well-known. Consider a survey of people each of whom uses one but not the other product. Suppose one respondent says, "The new product provided relief at level five for a full day." The other says, "The old product provided relief at level four for eight hours." Maybe the first respondent thinks that level five relief is decent and the second that level four relief is really spectacular and getting that level of relief even for eight hours is a blessing. We can come up with a slew of examples of individual interpretations of the survey questions such that the survey can't tell us whether the statement is true or false.¹¹

(4) "Steph Curry is the best player in the NBA today." Generations of heated disputes in neighborhood bars and restaurants confirm that this is a statement of an opinion if anything is. And yet, it would be easy enough to characterize it as a

¹¹ For completeness, I'll assert without elaboration that other problems attend surveys of people who have used both the new and the old products.

statement of fact: Design the basketball equivalent of sabermetrics,¹² rank all current NBA players, and find out where Steph Curry is on the list. Of course, the hitch here occurs at the first step, where one would have to gain agreement about the components of the ranking of the sort we have about basic scientific and physical facts. The “opinion” component lies in the choice among competing ranking systems.

Though the example is mundane, it offers a version of a quite important consideration. We believe what scientists say (when we do) because they have achieved a consensus for the moment on what the evidence shows; they are not choosing among alternative systems for evaluating the facts. As Judge Lynch put it in *ONY, Inc. v. Cornerstone Therapeutics, Inc.*, “[W]hile statements about *contested* and *contestable* scientific hypotheses constitute assertions about the world that are in principle matters of verifiable ‘fact,’ for purposes of the First Amendment . . . they are more closely akin to matters of opinion, *and are so understood by the relevant scientific communities.*”¹³ The words “contested” and “so understood” alert us to the importance of consensus in generating the confidence we have in assertions by scientists (within their domain of expertise).

B. What Does a Factual Assertion Mean? How Meaning and Normative Assertions Are Intertwined

The two final examples involve assertions whose factual content is contestable and disagreements about meaning can’t be disentangled from normative assertions.

(5) “Time and time again, the Republican program of cutting taxes on the wealthy has proven to be a motor for economic growth and the improvement of well-being for everyone in the United States.” As we will see, this is the kind of political statement that John Stuart Mill characterized as an opinion in defending the proposition that dissemination of false opinions had first-order social value. Yet in a way similar to the Steph Curry statement, it certainly looks as if it is a statement

¹² For a description of several such efforts, see *Advanced Statistics in Basketball*, WIKIPEDIA, https://en.wikipedia.org/wiki/Advanced_statistics_in_basketball [https://perma.cc/Y9K9-YJYZ] (Apr. 21, 2022, 3:03 AM).

¹³ *ONY, Inc. v. Cornerstone Therapeutics, Inc.*, 720 F.3d 490, 497 (2nd Cir. 2013) (emphasis added). I thank Rebecca Tushnet for directing me to this case.

about facts revealed by historical inquiry: Look at the economic statistics for periods following the enactment of Republican tax cuts to see how much economic growth occurred and how whatever growth occurred was distributed (and rule out other explanations for growth or its absence).

It should be apparent, though, that imposing liability for disseminating the statement, should the historical inquiry turn out to show its falsity, would be blatantly inconsistent with principles of free expression—without our having to do any fancy analytic maneuvers to explain why disseminating false statements has social value. Put another way, the Millian defense of affording protection to the dissemination of false opinions works too hard to reach a conclusion that should have been obvious from the outset. The reason, I suggest, is that key terms, including at least “economic growth” and “well-being,” are normatively freighted: It’s not that we agree on what we are pointing to when we use the terms but find it difficult to measure whether growth or well-being has occurred but rather that we have different normative views about what *counts* as growth or well-being. That makes the statement a normative one rather than one about what I’ve been calling basic facts about the physical world.

(6) “Corn dealers are starvers of the poor.” This is a modification of an example Mill uses in a different context. Like the statement about tax policy, this one presents itself as factual. Unpacking it: The speaker has a theory about how the market economy works that generates an account of how wealth is distributed. (The speaker also has a theory about the just distribution of wealth, but the truth or falsity of that theory isn’t central to my point here.) The statement’s truth or falsity depends upon the theory’s truth or falsity. And we can test the theory through ordinary empirical inquiries, so it’s capable of being proven true or false. But as with the statement about atoms, whether we accept or reject the theory rests on our assessment of the epistemic authorities brought forth in its support.

To summarize: If we think that we shouldn’t impose liability for disseminating false opinions, our inquiry into the permissibility of doing so should be based upon an account, pragmatic in John Dewey’s sense, of how we make decisions about truth and falsity in our daily lives. I develop later the idea that such an approach leads to the conclusion that liability for the dissemination of false factual statements should be limited to cases where decision-makers impose liability in connection with basic scientific, physical, and similar facts. And as already noted, those cases are likely to be rare outside the context of commercial fraud—though when they

occur, they can be quite important, which is why we are today concerned about the dissemination of falsehoods and lies about basic facts.¹⁴

C. *Some Generalizations*

The preceding, perhaps overelaborate discussion, has several payoffs. First, the foregoing arguments suggest that we should distinguish among three types of factual statements. (1) *Statements about basic historical events, basic scientific facts, and basic descriptions of phenomena in the real world.* Here most of the work is done by the term “basic.” It means something like “used by people as they go about their daily lives, whether or not they’re conscious that they’re using the facts.”¹⁵ As we’ll see, this sort of pragmatic definition has to play a rather large role in the law relevant to regulating the dissemination of falsehoods and lies. (2) *Statements founded in substantial part upon theories about how the physical or social world works.* These statements will be true if the theories are true (and the statements follow from or are compatible with the theories). (3) *Statements that almost necessarily employ normatively inflected terms to describe aspects of the historical, physical, or social world.* Here the term “normatively inflected” should immediately raise red flags about the constitutional permissibility of regulating the dissemination of such statements.

Section IV argues that some second-order considerations strongly suggest that legal regulation of the dissemination of “merely” false statements of all three types should be disfavored, with regulation of the second and third types especially problematic. As we will see, some of those considerations aren’t applicable to the dissemination of lies, but other second-order considerations might be applicable, with the consequences that only some forms of regulating the dissemination of lies, principally lies of the first type (about “basic” facts), should be viewed as consistent with general principles of freedom of expression.

¹⁴ Discussions of “memory laws,” which identify some historical events and make it an offense to deny (or sometimes, to affirm) that they occurred, are prominent in the literature on regulating lies. The foregoing analysis provides the tools for assessing such laws: In short, attempting to apply such laws to specific assertions demonstrates how difficult it can be to figure out precisely what counts as a factual assertion and to disentangle the factual from the normative in the words used to describe past events.

¹⁵ By “going about their daily lives,” I mean things like driving their cars and picking things off the shelves at grocery stores, not things like deciding what job to take or whom to marry.

If so, we can't do without some legally tenable distinction between basic facts and "nonbasic" ones. And if pragmatic considerations drive our understanding of that distinction, similar pragmatic considerations should shape the distinction's legal version.

I've argued that we're likely to get into a morass if the legal version requires us to decide whether a statement is capable of being proven true or false. But I suggest, we don't have to come up with a legal test aimed at guiding that decision (for example, a test that lists some criteria for determining whether a statement is capable of such proof). We can get by with a rule that instructs decision-makers to attach the label "statement of basic fact" only when doing so is appropriate, with no further analysis—with one exception—of what constitutes appropriateness.¹⁶ The exception is that the decision-maker's conclusion that a statement is about a basic fact must not be wholly unreasonable.¹⁷

Finally, many of the concerns I've raised about using the criterion "capable of being proven true or false" disappear when we're dealing with liability for *lies*. The reason is that the difficulties are associated with the ability of listeners and other "outsiders" to determine whether a statement has the relevant characteristics, but the liar knows (or believes) the statement false.¹⁸ The speaker's knowledge (or belief) makes irrelevant the listeners' assessment of whether the statement is factual.

¹⁶ An analogy here is to instructing juries that they can convict only if they are convinced beyond a reasonable doubt that the defendant committed the elements of the crime charged, with no substantial elaboration of what "beyond a reasonable doubt" means. *See* *Victor v. Nebraska*, 511 U.S. 1 (1994) (holding it constitutionally permissible to give juries a traditional "beyond a reasonable doubt" instruction). To be clear, this is an analogy, not a suggestion that decision-makers can impose liability only if their conclusion that a statement is about a basic fact is supportable beyond a reasonable doubt.

¹⁷ For me, the requirement of "not wholly unreasonable" is applicable to all exercises of public power, and in particular isn't a special rule for legislative or administrative determinations that a statement is about a basic fact.

¹⁸ SHIFFRIN, *supra* note 2, describes certain contexts in which obligations of sincerity are suspended or less stringent. So for example, high school physics teachers might know that the "solar" model of atoms isn't accurate but teach it to their students nonetheless as a way of getting the students closer to the truth. (I thank L. Michael Seidman for the example.) My arguments, like Shiffrin's, don't apply in these contexts. But I emphasize, none of the central problems I discuss do involve such contexts.

Consider here some statements about the 2020 presidential election: that a sixth-degree equation can show that some official electoral tallies couldn't have been honestly reported, that some Italians using military technology remotely altered the results on many U.S.-based voting machines, and that a large number of fraudulent ballots were "dumped" late on election night in several key states. The next section explains why decision-makers shouldn't impose liability on people who disseminate those statements, mistakenly believing them to be true: The statements are ideologically and normatively inflected, and they are located towards the "opinion" end of the fact-opinion continuum. The picture changes dramatically, in my view, if the speaker believes the statements to be false—if, that is, the speaker is lying. Focus on the proposition that you can't impose liability on the opinions a person holds because doing so would be inconsistent with many of the first-order values protected by the law of free expression. It's hard to see how imposing liability on a person for false factual statements they put forth but actually don't believe is inconsistent with those values. Second-order consideration might alter that conclusion, but first we need to explore its foundations.

III. IS THERE A CONSTITUTIONAL INTEREST IN THE DISSEMINATION OF FALSE STATEMENTS OF FACT AND/OR LIES? THE FIRST-ORDER ANALYSIS

A. *The Core Analysis*

Frederick Schauer pointed out that the discussion of the dissemination of false factual statements is underemphasized in the free expression literature.¹⁹ The reasons are probably manifold: Outside the context of commercial fraud, itself outside the free expression tradition until recently, reasonably well-functioning but flawed democratic governments rarely target false factual statements for regulation, with libel regarding government officials (specifically, seditious libel) and more recently memory laws being the largest (and problematic) exceptions. Most regulations of seemingly factual assertions involve assertions that sit close to the "opinion" end of the fact/opinion continuum.

Yet the dissemination of false factual statements—my focus here—often will undermine rather than promote the values the law of free expression seeks to pro-

¹⁹ Schauer, *supra* note 2.

mote. This last point is clear in connection with the insertion of false factual statements into political discourse.²⁰ People who are told that there's a very high probability that Iraq has weapons of mass destruction available for use may support policies that they wouldn't support were they to be told, more accurately, that the probability is rather low. Similarly with people who are told that an infectious disease can easily be passed from one person to another by a handshake. What free expression interests are served by allowing such false information to circulate freely?

Mill argued that allowing false statements of this kind to circulate promotes free speech values by training listeners in the ability to distinguish truth from falsity. Confronted with what we initially believe to be falsity, he argued, we have to think about the grounds for that belief—that is, the grounds we have for holding the view that we believe to be true. As Christopher Macleod puts the point, “Lack of discussion of false beliefs . . . can lead to the loss of our ability to connect our true beliefs with a network of related beliefs and actions—in these circumstances, a belief is ‘held as a dead dogma, not a living truth.’”²¹ Or as Mill put it, engaging with false beliefs leads us to “a clear apprehension and deep feeling of [the] truth.”²²

Schauer observes that most of Mill's examples of false beliefs involve what I've called opinions rather than facts, and that Mill noted that “on a subject like mathematics . . . there is nothing at all to be said on the wrong side of the question.”²³ In my view, Mill's argument carries through for many basic facts, even if it doesn't for mathematics.

Suppose someone tells you that Donald Trump won a majority of the lawfully cast ballots in the 2020 presidential election. You believe—know?—that's untrue.

²⁰ I address the “personal autonomy” values associated with free expression below, in analyzing the permissibility of regulating social lies.

²¹ Christopher Macleod, *Mill on the Liberty of Thought and Discussion*, in *THE OXFORD HANDBOOK OF FREEDOM OF SPEECH* 3, 9 (Adrienne Stone & Frederick Schauer eds., 2021) (quoting JOHN STUART MILL, *On Liberty*, in 18 *THE COLLECTED WORKS OF JOHN STUART MILL* 228, 243 (J.M. Robson ed., 1996)).

²² MILL, *supra* note 21, at 252.

²³ Schauer, *supra* note 2, at 905 (“[Mill] talks of the wrongness of suppressing advocacy of ‘Tyrrannicide,’ of the importance of being able to discuss ‘open questions of morals,’ of the value of ‘professing and discussing, as a matter of ethical conviction, any doctrine, however immoral it may be considered,’ of the freedom to challenge ‘belief in a God,’ and of ‘religious opinions’ in general.”).

Mill's argument asks you to produce the reasons you have for *your* belief. As I've argued, those reasons are rooted in the epistemic authorities on which you rely: the mainstream media in the first instance, and ultimately, the experts on ballot counting on whom the media rely.

You then ask yourself, "Why should I rely on those authorities?" Leiter uses the term "epistemic authorities" to evoke Joseph Raz's notion of authority. For Raz, authorities are institutions whose judgments *displace* each individual's assessment of their first-order reasons for action or, in the present context, their first-order reasons for belief. Relying on an authority means that you accept its assessment without yourself looking at the bases for the authority's conclusion—without, as today's conspiracy theorists put it, doing your own research. And why shouldn't you do your own research? Because, Raz argues, when you do, you're more likely than the authorities to get the wrong answer. That's not guaranteed: Sometimes the authorities have biases that lead them to make systematic errors that you wouldn't, and once in a while, you might actually do better research than even unbiased authorities. But overall, the system works better—our lives run more smoothly—if we accept the judgments of authorities without doing our own research.

So confronted with a factual statement inconsistent with your antecedent belief, you examine the authorities on which you've relied. You don't do your own research (even if you could), but you might well ask, "Is there some reason that on this question, the authorities on whom I'm relying are biased?" If you end up thinking that they aren't likely to be biased, you end your inquiry, now with your belief strengthened (or perhaps better, with more confidence that you had already arrived at the right answer).

Raz developed his argument in connection with the authority of the legal system. The case for allowing the dissemination of false factual statements because it leads us to think more seriously about the epistemic authorities on which we rely might be strengthened by noting a difference between legal authority and epistemic authorities. For each of us, there is only one legal authority whereas we have available to us many epistemic authorities. Raz argues that life would be quite bumpy if people routinely challenged law's authority. Not so, perhaps, if people occasionally

or even routinely pit one or a few of the available epistemic authorities against another.²⁴ More so, again perhaps, if one questions one or a few epistemic authorities only when the stakes are quite high.²⁵

The Millian argument, then, supplies first-order reasons for allowing the dissemination of false statements of fact. Doing so enhances our understanding of the truth by leading us to question and then gain confidence in our reliance upon epistemic authorities—or perhaps, leads us to question that reliance in some circumstances, thereby enhancing our ability to make decisions for ourselves.

That's not the end of the inquiry, of course. These first-order reasons might be offset by countervailing first-order reasons. For example, some people may mistakenly accept a falsehood as true without going through the inquiry into epistemic authority. We then do a first-order analysis of the situation. We might end up thinking that we get a bit more confidence in our understanding of the facts when we think through issues about epistemic authority, but that increment is overwhelmed by the distortions of judgment induced by widespread dissemination of falsehoods. In that event, we would have first-order reasons for regulating the dissemination of falsehoods.

What can we say about what has been called epistemic disagreement, that is, disagreement about which institutions should be treated as having epistemic authority?²⁶ Epistemic disagreement manifests itself today in the wholesale rejection

²⁴ To put the point a bit snarkily, even people who believe that Donald Trump won more votes in 2020 get on airplanes: They challenge the epistemic authority of the mainstream media but not that of aeronautic engineers.

²⁵ This is a phenomenon akin to or perhaps the same as what I understand philosophers refer to in describing “pragmatic encroachment.” For a brief discussion, see Dorit Ganson, *Pragmatic Encroachment*, ROUTLEDGE ENCYCLOPEDIA OF PHIL. (2019), <https://www.rep.routledge.com/articles/thematic/pragmatic-encroachment/v-1> [<https://perma.cc/JYA8-367X>]. Other occasions on which one might question an epistemic authority might arise when the authority reports facts that seem strongly counterintuitive to you or when the reported facts are novel and knowledge in the field is developing rapidly. For a discussion of the limits on but also the possibilities of doing your own research, which offers warnings but doesn't say you never should, see Nathan Ballantyne & David Dunning, Opinion, *Skeptics Say, 'Do Your Own Research.' It's Not That Simple*, N.Y. TIMES (Jan. 3, 2022), <https://www.nytimes.com/2022/01/03/opinion/dyor-do-your-own-research.html> [<https://perma.cc/SP6J-EAHE>].

²⁶ Michael Patrick Lynch, *Truth as a Democratic Value*, 64 NOMOS: TRUTH AND EVIDENCE 15, 22–25 (2021).

of the mainstream media as epistemic authorities. And I argue below, epistemic disagreement (or its absence) is a central condition for determining whether regulation of lies is consistent with free exercise principles. For the moment, I simply note my conclusion that epistemic disagreement is no different than disagreement about whether the Democratic or Republican parties are better at governance. All disagreements of this sort have to be handled by means other than content-based regulation.

I forgo analyzing the issues associated with the overall balance of first-order reasons because, as I've suggested before, the case presented by lies is different. What happens when the liar puts before us factual claims that he knows or believes to be false? Good Millians, we start to examine the bases for our beliefs. The benefits of doing that work, though, are almost certainly outweighed by its opportunity costs, which are deliberately imposed and might indeed be an important reason for lying in the first place. The liar has diverted our attention from the facts themselves to something else and thereby deprives us of the opportunity to devote our attention to other matters (we start worrying about why we should believe scientists' assertions about COVID-19 and can't use the time devoted to exploring that issue to working to support the expansion of paid family leave).²⁷ Antitrust law has a concept "raising rivals' costs" that describes actions that allow a potential monopolist to gain market share not by making a better product but by making it harder for competitors to make their own products.²⁸ As in antitrust law, raising rivals' costs has no first-order free expression benefits.²⁹

²⁷ Cf. SHIFFRIN, *supra* note 2, at 137 ("[W]e are adversely affected when lies introduce an epistemic need to investigate and confirm the particular reliability of individual speakers or their reliability about specific topics. . . ."), 141 ("Our time and attention are limited. If we aim to identify and appreciate the truth, it is beyond foolhardy to devote those limited resources by launching off from random starting points.").

²⁸ Thomas G. Krattenmaker & Steven C. Salop, *Anticompetitive Exclusion: Raising Rivals' Costs To Achieve Power Over Price*, 96 YALE L. J. 209 (1986).

²⁹ I briefly address in this note two topics that regularly arise in discussion of the regulation of lies. (1) *Autobiographical lies*, such as that told by Xavier Alvarez. See David S. Han, *Autobiographical Lies and the First Amendment's Protection of Self-Defining Speech*, 87 N.Y.U. L. REV. 70 (2012). As Shiffrin has argued, speech sustains networks of relationships among people. Autobiographical lies destabilize those networks because you quite literally don't know whom you're dealing with. Knowing that the person in front of you might be someone else, you can't trust anything they say.

B. Conclusion on First-Order Reasons

To sum up: The Millian argument shows that dissemination of false statements even of (some or many) basic facts can have first-order value by provoking serious reflection about epistemic authorities but not that dissemination of lies about those facts has first-order value. Perhaps the dissemination of autobiographical and social lies has first-order value; that value might be outweighed by other first-order reasons and regulation of such lies would be permissible; or that value is great enough to support a categorical ban on regulating those lies even if regulation of other types of lies is permissible.

IV. THE SECOND-ORDER ANALYSIS: CAN LEGAL INSTITUTIONS RELIABLY DISTINGUISH BETWEEN MERE FALSEHOODS AND LIES?

Analysis of free expression legal issues can't stop after identifying and evaluating the values associated with various forms of expression. It has to continue to an institutional level by asking whether or when which of our various legal institutions can reliably identify circumstances under which regulating some form of expression will promote or at least not undermine the values served by the system of free expression. Put another way: The system of free expression includes regulatory institutions as well as speakers and listeners, and understanding how institutions

And Shiffrin argues, interpersonal trust enables us to be autonomous individuals in a world of relationships. SHIFFRIN, *supra* note 2. Autobiographical lies undermine the value of autonomy they are said to advance; banning them would promote that value. (2) *Social lies*. According to Justice Breyer, "False factual statements can serve useful human objectives, for example: in social contexts, where they may prevent embarrassment, protect privacy, shield a person from prejudice, provide the sick with comfort, or preserve a child's innocence. . . ." *United States v. Alvarez*, 567 U.S. 709, 733 (2012). Some social lies simultaneously serve and undermine "useful human objectives." The lie, "I was working late at the office" might prevent embarrassment and protect privacy, but when the liar was actually out having a beer with other friends or worse, it places the relationship between the speaker and the listener on a shaky footing. And there are frequently (always?) ways of achieving the desired goals without telling a social lie ("I expect to be working really hard this week and was planning to veg out on Saturday night" rather than "Sorry, we can't go to dinner with you on Saturday because we have another engagement."). Shiffrin generalizes this last concern in a way that suggests that social lies *can't* serve "useful human objectives." She argues that all social lies are inconsistent with the trust that lies at the bottom of human relationships and that human objectives are constituted by such relationships. They "ambiguat[e] signals that function well only when fairly clear, signals whose preservation and use are crucial for sustaining a functioning moral and political culture." SHIFFRIN, *supra* note 2, at 136.

work is necessary for understanding what regulations should be allowed or prohibited.

So even if one believes that disseminating lies about basic facts lacks first-order free expression value, regulating dissemination of such lies might be inconsistent with free expression values *if* (and *when*) we have good reason to believe that the legal institutions tasked with regulating lies can't reliably distinguish between lies about basic facts and "mere" falsehoods about such facts. Understanding this second-order analysis in the present context requires us to begin by understanding the second-order analysis of the dissemination of mere falsehoods about basic facts.

A. *The Institutional Analysis of Regulation of Mere Falsehoods and How It Can Be Extended to Deal With Lies*

New York Times Co. v. Sullivan offers the canonical—and correct—institutional analysis of the regulation of mere falsehoods (outside the context of commercial speech).³⁰ Focusing on jurors and judges as ultimate decision-makers (to deploy the distinction developed in Section I), the Court began by noting that imposing liability for disseminating a false factual statement solely on the ground that the statement was false raised a substantial concern about "chilling effect." That effect arises because ultimate decision-makers acting in entire good faith might sometimes make a mistake and label as false a statement that's actually true. Concerned to avoid liability for publishing something, publishers will steer clear of the forbidden zone and refrain from publishing statements that might mistakenly be found to be false. That results in a reduction in the availability of true statements for the public to think about and predicate decisions upon.³¹

This insight is pretty clearly correct.³² In the libel context, it leads to efforts to structure liability rules to achieve a socially desirable balance between protection of

³⁰ *New York Times Co. v. Sullivan*, 376 U.S. 254 (1964).

³¹ Ultimate decision-makers might err in the opposite direction, identifying as true statements that are actually false. Unfortunate as that might be in some cases (for example of an exonerated defendant who fails to obtain relief from a civil jury because the jury believes him to have committed the crime), the availability of false statements isn't a cost to free expression values. It might be a cost to the good functioning of a democratic system, which is why we try to devise methods of reducing the distribution of false factual statements.

³² A commenter on an earlier version of this article observed that the possibility of liability might also induce the speaker to reflect on her epistemic premises—that is, why they believe the statement to be true. That's probably a good thing, so the measure of the chilling effect should offset

reputation and dissemination of information to the public. The details of those rules don't matter here,³³ though I emphasize that *Times v. Sullivan*'s analysis of the way institutions in the libel system operate doesn't necessarily extend to the analysis of other institutions in other contexts.³⁴

The concern about regulating "mere" falsehoods is that institutions will misidentify true statements as false ones. What's the parallel concern about regulating lies? That institutions will misidentify mere falsehoods as lies. Liability for disseminating a lie might be imposed on someone who actually believes the false statement to be true. A person who says that Donald Trump received more votes than Joe Biden in 2020 is making a false statement of basic fact but isn't necessarily lying if they honestly believe the assertion. We might worry, though, that some relevant decision-maker will conclude that they're lying or that anyone who makes such a statement must be lying.³⁵ The general version of this difficulty is straight forward: A decision-maker might infer from a statement's evident or obvious falsity that the person making it must have done so knowing it was false.³⁶

When is this risk likely to arise? When, I suggest, the false statement of basic fact's truth or false is tested by referring to epistemic authorities *in situations of*

the concern about being mistakenly found liable by taking into account that reflecting upon epistemic premises might lead the speaker to conclude that the statement is actually false (and refrain from disseminating it for that reason).

³³ When courts develop those details, as they must, another set of institutional question arises about the capacity of courts to do a good job (with "good" specified appropriately). Theories of judicial review address those institutional questions, but taking them into account here would take the present inquiry far too far afield.

³⁴ That is, *Times v. Sullivan* doesn't imply that we can regulate the dissemination of any kind of falsehood only when the falsehood is made maliciously or with deliberate indifference to its truth or falsity. In particular, what kind of falsehood is involved and who the regulating institution is might matter a lot. For present purposes, the relevant kinds of falsehoods are those dealing with basic facts. (Under the *Times v. Sullivan* doctrine, Barack Obama could recover damages for reputational injury itself established by evidence for the dissemination of the falsehood that he wasn't born in Hawaii by a person who believed the statement to be false, or published the falsehood with reckless disregard to its truth or falsity.)

³⁵ See AL FRANKEN, LIES (AND THE LYING LIARS WHO TELL THEM): A FAIR AND BALANCED VIEW OF THE RIGHT (2003), for a comic version of this proposition.

³⁶ The difficulty would be exacerbated were ultimate decision-makers allowed to impose liability if they concluded that the speaker acted with reckless disregard of the statement's truth or falsity.

epistemic disagreement, that is, when there's a real possibility that the person making the statement doesn't regard those on whom institutional decision-makers rely as epistemic authorities.

The concept of epistemic disagreement is central to my argument, so it's important to be clear that epistemic disagreement is different from what we might call "ordinary" disagreement about what the facts are. Consider a recent example of ordinary disagreement. During the oral argument about staying the effect of an Occupational Safety and Health Administration regulation, Justice Gorsuch asked a question that incorporated a reference to the number of deaths caused by flu each year. Those who listened in real time disagreed about whether he said "hundreds of thousands" or "hundreds, thousands." How do they deal with that disagreement? They "go to the tape" and listen again;³⁷ they apply a principle of charity in interpretation ("hundreds of thousands" is so wildly wrong that it's implausible to think that Justice Gorsuch said that); they observe that the force of the justice's argument depended on the number of deaths caused by flu to be roughly comparable to that caused by COVID-19; they might someday have access to the notes the justice took in preparing for the oral argument; and more.

At the end of the inquiry, some might still think that he said "hundreds of thousands" and will conclude that those who disseminate the statement that he said "hundreds, thousands" are disseminating a falsehood (and of course reciprocally for those who think he said "hundreds, thousands"). In cases of ordinary disagreement, participants in the discussion agree that certain data constitute the set of facts from which further factual inferences are to be drawn.³⁸

Cases of epistemic disagreement are different. Suppose that, on listening to the tape again, both sides agree that it clearly shows the justices saying "hundreds of thousands." But those who initially heard him say "hundreds, thousands" contend that the tape was altered after it was made in real time. The other side says, "Well, let's get some experts in audio reproduction technology to examine the tape and tell us whether it's been altered." A task force of 10 experts is convened, and the

³⁷ They can't go to the transcript because the initial transcript had "hundreds of thousands" and then was modified to "hundreds, thousands."

³⁸ Typically, not much will turn on that conclusion; one side might become more skeptical in the future about the other sides' factual assertions. Or the statement might be drawn into some authorized forum in which an ultimate decision-maker is authorized to take some action when they conclude, after applying a specified burden of proof, that one of the statements was false.

members unanimously conclude that the tape wasn't altered. Epistemic disagreement occurs when the "hundreds, thousands" side responds by (perhaps) finding a lightly credentialed student of audio reproduction technology who says that the tape was altered, by casting aspersions upon the professional credentialing process that treats "their" expert as less qualified than the task force members, and the like.³⁹

The distinction between ordinary and epistemic disagreement isn't inscribed in nature. It arises because people sometimes disagree not about what the balance of evidence is (no matter what we require that balance to be—that is, no matter whether we're looking for the preponderance of the evidence or for some more substantial outweighing of the evidence against the asserted basic facts) but in the special case where the disputants disagree about the epistemic authority of some institution or institutions that provide one (significant?) component of the balance of evidence. Almost any ordinary disagreement can become an epistemic disagreement if one of the disputants thinks the stakes are high enough: The stakes lead the disputant to look for some *new* epistemic authority supporting their position.⁴⁰

Without suggesting that the following provides a structure for allocating burdens of proof at trial (although it might), we can identify when epistemic disagreement *doesn't* lie at the base of the assertions of the false statement of fact when the speaker can't or won't direct our attention to any epistemic authority on which they

³⁹ The possibility that speakers disagree about who has epistemic authority is one reason that imposing liability for lies on the basis of deliberate indifference to the statement's falsity is problematic. Using terms introduced by Harry Frankfurt, an ordinary liar and a truth-teller orient themselves to the truth as identified by an epistemic authority they both accept, but in different directions. HARRY FRANKFURT, *ON BULLSHIT* (2005). I don't think it's unreasonable to impose liability when a speaker is deliberately indifferent to whether he's oriented toward truth or falsity. In contrast, where speakers disagree about who is an epistemic authority they are both orienting themselves toward the truth. But if one epistemic authority is widely accepted and the other not, the risk seems to me nontrivial of allowing liability for deliberate indifference to lead to an inference of deliberate indifference from the acceptance of an eccentric epistemic authority. For additional discussion, see Mark Tushnet, *Trust the Science but Do Your Research: A Comment on the Unfortunate Revival of the Progressive Case for the Administrative State*, 98 *IND. L.J.* 335 (2023).

⁴⁰ I've inserted the qualified "almost" to preserve the possibility, mentioned by Mill, that disagreements that within one or two steps become disagreements about agreed-upon measurements ("Is Steph Curry taller than LeBron James?") can't be converted into epistemic disagreements.

rely.⁴¹ An alternative equally informal “test” might be this: We ask ourselves why we believe the statement to be false and identify the epistemic authorities we’re relying on (the mainstream media, well-respected scientists, and the like). Then we ask why the speaker *might* believe the statement to be true. For ordinary people, quite often the answer will be that they are relying on a different set of epistemic authorities (Fox News, a scientist who disagrees with their colleagues). In such cases, the speaker isn’t lying, and we’re back to the “mere falsehood” case. We can continue our inquiries, though, and ask why the alternative epistemic authorities might believe the statement to be true. Sometimes, the answer will be that they are relying upon something like one or two un-peer-reviewed scientific studies. In these cases, the epistemic authorities aren’t lying either. But—and this is crucial—sometimes the answer to our inquiry about the bases for the epistemic authority’s assertion will be, “They got nothing” (or in Donald Trump’s words, “A lot of people are saying”—which is a statement of fact, just not a reference to someone supplying evidence *about* the underlying fact).⁴² At that point, we are indeed in the land of lies, not by the person making the statement but by the authorities on whom they rely.

Here are three examples. (1) An easy example is a false statement about where a polling station is located (in ordinary circumstances, that is, when there haven’t been recent changes in the station’s location). No decision-maker is likely to mistakenly infer from the statement’s obvious falsity that the speaker believed it to be true. (2) A slightly more difficult example is a false statement that COVID-19 vaccines contain microchips that allow the government to track your location.⁴³ Without some reason to believe that *some* epistemic authority supports that assertion, institutional decision-makers are unlikely to make the mistaken inference with which we are concerned.

⁴¹ QAnon-type falsehoods *might* be good examples of such assertions, except to the extent that they rest on the denial of epistemic authority to, for example, the Federal Bureau of Investigation with respect to claims about a pedophilic conspiracy among Democratic Party leaders. Cf. RUSSELL MUIRHEAD & NANCY ROSENBLUM, A LOT OF PEOPLE ARE SAYING: THE NEW CONSPIRACISM AND THE ASSAULT ON DEMOCRACY (2020) (discussing inter alia evidence-free conspiracy theories).

⁴² Occasionally, the speaker might be relying upon someone else’s “research,” but doesn’t believe that that person has epistemic authority, that is, has access to sources of factual knowledge that the speaker themselves is incompetent to evaluate.

⁴³ I regard this as slightly more difficult because I haven’t seen any evidence that people who do their own research actually have identified epistemic authorities that support the assertion.

And (3) an easy example in the other direction—that is, an assertion about a basic fact that is founded upon epistemic disagreement—is, alas, the false statement that Barack Obama wasn't born in Hawaii.⁴⁴ The epistemic disagreement here is over whether the Hawaiian authorities that generated Obama's long-form birth record can be trusted not to have faked it coupled with the undeniable fact that Obama's father was from Kenya and the common-sense (though false quite often) proposition that most children are born in their father's home nation. It's fairly easy to see how institutional decision-makers might infer from the statement's falsity and the inaccuracy of the common-sense observation that the speaker knew the assertion to be false when, we know as a matter of regrettable fact, that many people actually do believe the statement to be true.

Why though should we worry about imposing liability in situations of epistemic disagreement? Because, I suggest, epistemic disagreement is a form of political disagreement—disagreement with “the powers that be” with respect to what are reliable sources of knowledge.⁴⁵ This is clear enough when epistemic disagreement is presented as a challenge to the “lamestream media,” or to the “deep state,” or the professional ideologies of doctors in the pocket of “Big Pharma” (or scientists employed by the deep state).⁴⁶ I'm reasonably confident that, when analyzed carefully, all forms of epistemic disagreement will turn out to be challenges to the powers that be.⁴⁷

⁴⁴ I use the example because it is the first one offered in Schauer, *supra* note 2.

⁴⁵ As I read his article, this is the thrust of Leiter's concern about epistemic disagreement—that is, that in the current conditions of the United States, the epistemic disagreement we're experiencing is a pernicious form of political disagreement, one in which some people take positions on policy matters that have empirical foundations while rejecting the proposition that specialized knowledge is sometimes a necessary precondition of the policies' defense. I observe that Leiter appears to believe, as I do, that imposing liability in cases of political disagreement isn't categorically prohibited by free speech principles.

⁴⁶ If my argument is right, epistemic disagreement as I've defined it is ultimately resolved by the ordinary methods of political contestation—but precisely for that reason shouldn't be resolved by regulating those who rely upon “minority” epistemic authorities.

⁴⁷ Without claiming to have mastered the relevant literature, I have a reasonably strong sense that this is one of the core claims of recent (that is, over the past generation or two) work in the field known as “science and technology studies.” Cf. Harald Rohrer, *Science and Technology Studies, History of*, in *INTERNATIONAL ENCYCLOPEDIA OF THE SOCIAL & BEHAVIORAL SCIENCES* 200, 201 (James D. Wright ed., 2nd ed. 2015) (“Facts and artifacts are but temporarily stable outcomes of

B. A Sketch of an Institutional Analysis of When the Risk of Mistaken Inferences About Belief Will Arise

In a reasonably well-functioning but flawed democracy, what's the political economy generating regulations prohibiting the dissemination of lies? The analysis has two steps: (1) What's the political economy generating regulations targeting the dissemination of false statements of fact? (2) With respect to the false factual statements identified at the first step, what's the political economy of generating regulations targeting the dissemination of such falsehoods knowing (or believing) them to be false?

Notably, the legal institutions of reasonably well-functioning but flawed democracies rarely have incentives to impose liability for the dissemination of falsehoods about basic facts.⁴⁸ Rarely, but not never: Consider an election-protection provision making it an offense to disseminate a false statement about the location of polling places (“If you live on Holly Street, you should go to Shepherd Library to cast your vote,” when actually, the polling place is located at the Shepherd School). Or a public-health regulation targeting the dissemination of false statements about the conditions under which an infectious disease is transmissible or about the efficacy of public health measures aimed at reducing the transmission of such diseases (although drafting a regulation precise enough to avoid serious overbreadth concerns will be quite difficult).⁴⁹

I believe we can put to one side the possibility that such a democracy would produce a Ministry of Truth with a general charge of identifying false statements of basic facts and regulating the dissemination of such facts by those who know or believe them to be false. There's no obvious interest group other than a “good government” lobby that might support creating such an institution, and many—indeed, I suspect almost all—interest groups that would oppose doing so: Legacy and new media and universities would oppose doing so on principle, and substantive

heterogeneous activities of scientists and engineers and *their entanglement in wider social and political relations.*” (emphasis added).

⁴⁸ Orwell's example of the party's ability to get people to believe that two plus two equals five suggests that totalitarian governments might have incentives to penalize statements about basic facts, though it isn't directly about penalizing true statements about basic facts.

⁴⁹ I note that both examples, the latter more than the former, would probably be best implemented by legislation delegating the power to develop such regulations to a general-purpose administrative agency.

interest groups might reasonably fear that someday the Ministry of Truth would decide that their members were producing false statements of basic facts.⁵⁰

1. Legislation

Memory laws show that there is a political economy that generates legislation targeting *some* falsehoods. Analysis is complicated for the United States by the fact that we have legislatures operating at many levels: national, state, local, and school boards most prominently. The political economy of each institution is at least a bit different, and analysis is made even more complex because we have to consider not only the political economy generating regulation at each level but also the law and political economy associated with a higher level's power and practical ability to override regulations adopted at lower levels.⁵¹ For that reason, I focus on the political economy of national and state-level legislation, with the hope that the form of the argument will suggest how analysis of lower-level regulation would go and then turn to the political economy of legislation delegating authority to regulate lies to administrative agencies with substantive charges, such as regulating food and drugs, communications technologies, and the competitive economy.

(1) *Legislation against lies.* I believe that we can fairly assume that every interest group at one or another point gets annoyed enough at falsehoods associated with whatever it is they're interested in to think that laws against such falsehoods would be a good thing. "Big Pharma" wants to regulate lies about the risks associated with vaccines; "Big Ag" wants to regulate lies about the risks associated with genetically modified organisms; examples could be proliferated ad infinitum. These efforts will of course face opposition, from the media and from universities concerned about deterring the distribution of research findings, among others.

Sometimes the groups favoring regulation will simply have more power than the opposition. I suggest that another way of putting that point—that sheer power

⁵⁰ In my view, concerns about Ministries of Truth in reasonably well-functioning but flawed democracies are substantively trivial and wouldn't deserve attention in the text, but Justice Kennedy's prominent mention of the Ministry of Truth in *Alvarez* requires textual treatment.

⁵¹ I note my personal view that Section 5 of the 14th Amendment (and cognate enforcement provisions elsewhere in the Constitution) gives Congress the *power* to override all state and local regulations aimed at actions that Congress believes (with minimal rationality) to threaten constitutional values. That view is not current law, and in any event, saying that Congress has such power doesn't tell us when it might choose to exercise it.

is involved—is this: Interest groups can sometimes obtain regulation of false statements of basic fact by exercising their greater *epistemic* power, that is, their ability to determine the epistemic authorities upon which the legislature relies. Scientific consensus, as I noted earlier, is one form in which scientists’ epistemic power manifests itself.

If that’s a fair statement about the role of sheer power, it shows why these regulations of lies about basic facts are impermissible: The targeted false statements of basic fact are within the class as to which the risk of mistaken attributions of knowledge of falsity is worryingly high (because those making the statements do have their own epistemic authorities on which they rely).

Sometimes, though, the proponents of regulation of lies about basic facts don’t face epistemic disagreement (recall the example of lies about a polling station’s location). I suspect that this occurs only in connection with what would traditionally have been called pure “good government” reforms, which might be quite a small category today. But should a legislature enact a statute targeting lies about a specific class of lies about basic facts where there’s no epistemic disagreement, such a statute should not be understood as inconsistent with free expression principles. The Stolen Valor Act is a good example here.

I conclude this discussion with comments about two additional issues.

(a) *Particularity*. Why “a specific class of lies”? Because what’s permissible depends upon the absence of epistemic disagreement, and it’s exceptionally difficult to draft a statute of even modest generality that wouldn’t overbroadly sweep in cases where there’s epistemic disagreement. Trust me on this: I’ve tried. Think about drafting a statute that would somehow generalize “Barack Obama wasn’t born in Hawaii” into a ban on—what? Lies about whether a person certified by Congress as president satisfies the constitutional qualifications for the presidency? Similarly for a statute, generalize the polling station example: lies about the procedures for casting lawful ballots? I leave it as an exercise for the reader to explain how such statutes would be overbroad.⁵²

⁵² In principle, concerns about overbreadth can be addressed by excising overbroad provisions from statutes, sometimes by interpreting them narrowly, sometimes by holding the overbroad provisions unconstitutional; overbreadth doctrine arises because sometimes those techniques are unavailable (in the United States, with respect to state legislation and some federal statutes) or their use would raise questions about fair notice.

(b) *Selectivity*. The political economy of legislation targeting lies about which there's no epistemic disagreement suggests pretty strongly that if legislatures enact any such regulations, they will do so with respect to some lies but not with respect to other seemingly quite similar lies. Though memory laws aren't a good example of permissible regulation of lies, they do illustrate the problem that is associated with permissible ones. Nations with laws against the Holocaust lie have faced pressure from ethnic groups, most notably Armenian, to adopt similar laws about genocides committed against them.⁵³ What, though, about *other* genocides? As far as I know, no nation yet has a law against denying that the Myanmar (formally Burma) government has conducted a campaign of genocide against the Rohingya. This selectivity might well be a general normative problem, but I doubt that it is a problem of free expression. Unless there's some other reason to be suspicious, the fact that a legislature hasn't regulated everything it could doesn't cast doubt on the validity of the "underbroad" regulation.⁵⁴

⁵³ For a discussion of the adoption of a French law recognizing the Armenian genocide (later invalidated by the French Constitutional Council), see Scott Sayare & Sebnem Arsu, *Genocide Bill Angers Turks as It Passes in France*, N.Y. TIMES, Jan. 23, 2012, at A4.

⁵⁴ There are scattered indications in U.S. Supreme Court doctrine of a principle of "First Amendment underbreadth," but they arise in settings where one can fairly suspect that some form of discrimination barred by non-free-expression principles is afoot. See Michael Coenen, *More Restrictive Alternatives*, 96 N.C. L. REV. 1, 20–25 (2017); John Fee, *Greater-or-Nothing Constitutional Rules*, 64 CASE W. RES. L. REV. 101 (2013). The Armenian-Rohingya example might actually present such a setting (with race-based discrimination as the concern). Cf. *Perinçek v. Switzerland*, App. No. 27510/08 (Oct. 15, 2015), <https://hudoc.echr.coe.int/eng?i=001-158235> (opinion of J. Nussbarger, partly concurring and partly dissenting) (rejecting the majority's distinction between Holocaust denial laws and a Swiss Armenian genocide denial law). SHIFFRIN, *supra* note 2, at 124, suggests that the Stolen Valor Act was troublingly underbroad because it penalized lies about military honors but not lies about having received government awards for civilian service and so was viewpoint discriminatory. Perhaps so, but her suggestion points to one general difficulty with a doctrine of underbreadth: We have to determine what the comparison class is—here, why "lies about government-award civilian honors" rather than "lies about socially recognized civilian honors" or even "lies about socially honorable actions whether or not previously recognized officially"? Shiffirin flags the difficulty, *id.* at 131, but doesn't address it directly, writing only that she is "less certain that the path to content-discrimination is inexorable" and offering a few doctrinal tweaks that might limit the scope of the content-discrimination concern. (In the informal discourse of scholars of constitutional law, this problem is referred to as that of "baseline hell," and I think there's general agreement that the field hasn't achieved anything close to an answer about how to escape it.)

(2) *Delegating authority to administrative agencies.* Interest groups can demand regulations of lies they've already identified or of lies of a general type, of which they can present examples. The problem of drafting specificity, which arises when interest groups pursue the latter course, can be addressed by enacting statutes delegating the authority to identify and regulate specific lies within a legislatively defined category to administrative agencies.⁵⁵ Delegation has other advantages as well, but there's no need to rehearse here the general accounts of the political economy of delegations to administrative agencies: They allow legislators to "do something" without forcing them to do anything in particular, they create institutions to which specific interest groups might have greater access than they do to the legislature as a body, and the like. Here, the relevant conclusion is that political-economy considerations suggest that delegations to special-purpose agencies *can* occur.

C. *Administrative Agencies*

Administrative agencies have experience and expertise in their areas of substantive concern. Election commissions know a lot about conducting elections and, depending upon their authority, may know a lot about campaign practices.⁵⁶ Trade regulation agencies know a lot about trade practices, including advertising. Food and drug regulatory agencies know a lot about the characteristics that make medications safe and effective for their intended uses. If charged with identifying false factual statements about basic facts, the agencies' knowledge bases allow them to identify specific false statements with some precision: "Vote at the Shepherd Library" is false; so is "COVID-19 vaccines have microchips in them"; "Only vote in person to make sure that your ballot is counted" isn't false; neither is "COVID-19 vaccines may cause long-term permanent damage to male reproductive organs."

Expertise is different from experience. An agency's expertise lies in the epistemic authority it asserts when making or evaluating *some* factual statements. Expertise doesn't come into play with respect to all such statements because sometimes agencies rely upon "ordinary" knowledge available to everyone: Compare the epistemic basis for the assertion about the Shepherd Library with that for the assertion about microchips—experience for the former, expertise for the latter.

⁵⁵ This is in the present context the form taken by the general defense of administrative agencies as able to respond to unforeseen developments within their general areas of responsibility.

⁵⁶ For a comparative-law focused discussion of electoral management bodies, see MARK TUSHNET, *THE NEW FOURTH BRANCH: INSTITUTIONS FOR THE PROTECTION OF CONSTITUTIONAL DEMOCRACY* 123 (2021).

We're back in trouble when agencies rely upon expertise in contexts of epistemic disagreement.⁵⁷ That's because expertise ordinarily has some ideological and political content.⁵⁸ Electoral management bodies work with an image of sober and deliberative politics. Consensus standards in science result from socialization processes that involve exercises of power. The dissident who asks, "Why are your epistemic authorities better than mine?" is raising a point about how epistemic power is distributed in ways correlated with other forms of power.

D. Forms of Liability

The second-order considerations I've described are driven by concerns about deterring expression that has first-order value: truthful statements in the case of mere falsehoods, false statements that nonetheless provoke valuable reflections on listeners' epistemic judgments in the case of lies. The amount of deterrence may vary with the precise content of the doctrine we develop to deal with it.

For example, we might think that criminal liability will deter more expression than civil liability. If so, we might think that the smaller effect of civil liability might justify imposing it even when criminal liability wouldn't be justified. That might be especially true in the case of lies, where what's being deterred are (by assumption) false statements that test listeners' epistemic judgments: We might think that a decrease in such statements, while a cost, isn't as costly as the deterrence of true statements.⁵⁹ On this view, imposing liability on the dissemination of lies in the form of

⁵⁷ Recall that I'm unsure that the microchip assertion does involve epistemic disagreement. If it doesn't, an agency regulation imposing liability for making it while knowing it to be false would be permissible because the statement is false and, in the absence of epistemic disagreement, there's little risk that the ultimate decision-maker will mistakenly infer from its falsity that it must have been made with knowledge thereof (or believes therein) when in fact the speaker believed it to be true.

⁵⁸ See TUSHNET, *supra* note 56, at 144.

⁵⁹ The U.S. Supreme Court treats civil and criminal liability for mere falsehoods as equally troubling, in part because the liability rules apply to corporate speakers, as to which application of criminal sanctions is difficult, and in part because in the Court's view the consequences for media speakers of a simple finding of liability (accompanied by civil sanctions or even in a declaratory judgment) is substantial. Whether that equivalence holds for other speakers, and in particular for speakers using new information technologies, is a question worth examining. One complexity is that traditional media operations now use new information technologies, making drafting an acceptable regulatory system rather difficult.

a declaratory judgment or an injunction against repeating the lies might be permissible.⁶⁰

Yet we must also take into account the possibly different burdens of proof in civil and criminal cases. We'll get more deterrence of first-order valuable expression if the burden of proof for civil liability is lower than for criminal liability. Clearly, there's no arithmetic method to calculate these amounts, and I wonder whether courts are able to do so intuitively well enough. And there's a further difficulty. Even if the burdens of proof are *formally* equivalent, they might not be equivalent in practice. No matter what the stated law is, ultimate decision-makers might think to themselves, "It's not as if we're going to send this defendant to jail if we find her civilly liable."

We might well disagree among ourselves about how all these factors shake out, with some concluding that the standards for imposing all forms of liability should be the same, others concluding that it should be permissible sometimes to impose civil liability where criminal liability is unavailable.⁶¹

CONCLUSION

Today, epistemic disagreement may be larger than it has been in the past, which suggests that calls for regulating lies will bump up against free expression values more often. It's easy enough to say that disseminating lies can't advance free expression values. The distinction we have to begin with isn't between truth and falsity or even between "mere" falsity and intentional lies; it's between disagreements within an accepted epistemic framework and disagreements about what should be accepted as a source of epistemic authority. With that distinction in hand, we can

⁶⁰ Consider an additional regulatory technique: requiring the disseminator of a statement found (by some decision-maker) to be false to attach a disclaimer (a notional "sticker") to any further dissemination. We can call this compelling speech, but many of the reasons we have for being nervous about compelled speech are inapplicable or only weakly implicated where government-drafted stickers are attached to statements. There's a low risk of attributing the sticker to the speaker, for example, and being forced to attach the sticker is unlikely to induce psychological distress or reconsideration of the original statement in the speaker. For a more complete analysis, see Kenny Chng, *Falsehoods, Foreign Interference, and Compelled Speech in Singapore*, 18 *ASIAN J. COMPAR. L.* 235 (2023).

⁶¹ And because that disagreement is reasonable, the questions about the role of courts and legislatures in devising liability rules that I've put to one side here would return in full force.

ask which if any of our regulatory institutions can reliably make it and impose liability only with respect to the first kind of intentional falsehood.

The justifications for regulating lies and negligent and reckless falsehoods are different, both with respect to first-order and second-order concerns. Perhaps “mere” falsehoods are valuable because they force listeners to clarify for themselves the grounds they have for their beliefs. Further, those who utter “mere” falsehoods may genuinely be relying upon a different set of epistemic authorities than regulators rely upon; regulating such falsehoods would in reality be to punish speakers for epistemic disagreement, which is fundamentally political disagreement.

These concerns, though, disappear or at least diminish substantially when it comes to intentional falsehoods, because the liar isn’t claiming to disagree with the listener’s epistemic authorities, and no clarity results when the listener confronts the lie. Free expression values aren’t violated if and when regulators can reliably distinguish between false statements that reflect disagreement about epistemic authority, which can’t be regulated, and intentional falsehoods. That second-order question requires discrete analyses of particulars: What institution is regulating? What kinds of facts are singled out for regulation? Memory laws might be problematic because of second-order institutional concerns; the Stolen Valor Act shouldn’t have been because it criminalizes intentional statements of fact about which there is no epistemic disagreement.

With the foregoing analysis in hand, we should be in a position to evaluate proposals for regulating the dissemination of falsehoods about elections, about medical treatments for illnesses, and more. That lies are bad isn’t enough reason to regulate them. We have to focus on more discrete questions: What exactly are the lies that are to be targeted and what institutions will be used to identify when such lies have been disseminated? I have no informed views on those questions but am confident that they are the ones we need to ask.⁶²

⁶² The argument also has implications for prohibitions against fraudulent commercial statements. Specifically, it suggests (1) that the core analysis would apply to specifically identified false commercial statements made with knowledge of their falsity, (2) that perhaps a general purpose regulatory agency such as the Food and Drug Administration might permissibly be allowed to identify such specific false statements, (3) that the proper second-order concern is whether decision-makers should be allowed to infer from mere falsity that the statement was made with knowledge or reckless disregard of its falsity, and (4) that if we conclude that first-order concerns support regulation of false commercial statements made with reckless disregard of their falsity, we should continue

to treat commercial speech consisting of factual statements as subject to a lower standard of free expression review than noncommercial speech.