





From Individual Rights to Community Obligations: A Jewish Approach to Speech

🔗 Jessica Hammer and Samantha Reig, Carnegie Mellon University

Hate speech, disinformation, doxing—these are all information harms spread through online speech. Questions of what is acceptable for online speech are typically framed as issues of rights. For example, as of this writing, Spotify is facing heavy criticism for hosting Joe Rogan’s podcast, in which he disseminates misinformation about vaccination. Spotify has responded by arguing that Rogan’s podcast did not violate their terms of service, and therefore he has the right to remain on the platform.

Rights-based framings around online speech are important. They make room for marginalized perspectives and ensure that people can protest injustice. The fact that the same conceptual tools can be used to argue for, say, giving Nazis access to public platforms does not mean they are unimportant. It does suggest, however, that perhaps we need complementary framings to help create a healthy online sphere.

Jewish thinking focuses not on rights, but on obligations. *Mitzvot* are the commandments that guide

Jewish behavior. Observing *Shabbat* (the Jewish Sabbath), treating employees fairly, keeping a kosher home, and helping take care of a friend in need are all mitzvot. Collectively, mitzvot form the basis for *halacha*, a complex religious-legal code that defines people’s responsibilities to themselves, to their families, to the community, and to the divine. Jews do not always agree on the nature of these responsibilities, how to prioritize them, and how flexible they are in response to social change. However, the orientation toward obligation is a common theme across Jewish communities [1].

The rules governing speech are complex, and speech laws in Judaism include rules about blessings, prayers, contracts, vows, promises, and more. In this article, we examine some of the *halachot* (laws) around interpersonal communication. These laws were codified and summarized by the *Chofetz Chaim* [2], and have since been further explicated by popular interpreters such as Joseph Telushkin [3]. Here, we argue that these rules can provide

Insights

- Jewish thinking emphasizes obligations over rights, and centers communities rather than individuals.
- The laws of Jewish speech help translate a complex and ambiguous topic into specific behaviors and practices.
- Jewish perspectives can give us new design tools to tackle wicked problems in HCI.

a novel conceptual approach, and potentially new design frameworks, for social media and other types of online speech.

LASHON HARA: THE EVIL TONGUE

A common defense around harmful speech is to claim that it is true, so it is justified. In a Jewish framework, even true speech is prohibited if it might harm another person. *Lashon hara*—harmful true speech about a person—is a transgression distinct from *hotza'at shem ra* (defamation). True speech about another person can be lashon hara whether or not the person who might be harmed is aware of it, even if the thing being said does not appear negative on the surface. For example, if you are commenting on someone's beautiful house, it might still count as lashon hara if the context means it would arouse resentment in the listener.

A modern analogy might be doxing. Doxing involves sharing true information about someone, such as their home address. In some contexts, that information is neutral or even positive. You might want to share your address with a visiting friend, or give it to a store so that they can deliver your order. However, being doxed implies that the information is being shared to threaten or harm. Consequences of doxing can include death threats, “swatting” (having armed police sent to the home), or stalking.

Halacha asks us to consider impact in light of context instead of trying to divine what someone else intended. Intent is unknowable. The person self-reporting their intent is also the one with the most incentive to lie about it, if they are indeed being malicious. In determining whether speech is lashon hara, instead of putting ourselves in the role of the speaker, we are asked to put ourselves in the role of a listener, and honestly evaluate whether the speech in question might have had a negative impact on us in some way.

Understanding lashon hara as a contextual judgment illustrates one of the challenges with the design of modern speech technologies such as social media platforms: *context collapse* [4]. Context collapse is when digital content intended for one online audience (community, social group, etc.) ends up on the screens of people outside that audience. Evaluating whether a true statement is lashon hara requires an analysis of social context. Who is listening? Whom

might they repeat the statement to? What conclusions might they draw? On current social media platforms, it is impossible to know for sure who is “listening.” Resharing, retweeting, screenshotting, and tagging other users all provide easy ways for a statement to go beyond the context in which it was originally said.

Context is also important in terms of who is being discussed.

For example, some groups, such as Black people, Muslims, and trans people, experience hypervisibility online. The underrepresentation of these groups makes individuals more salient targets for surveillance and other types of harm. A statement intended to be innocuous made about a white man might be very differently received when made about a Black woman.

In addition to taking responsibility for your speech, the laws of lashon hara make you responsible for what you listen to. The one who speaks, the one who listens, and the one who repeats a statement are all violators of the law. Unfortunately, our current online speech technologies do a poor job of giving us discrimination in listening. We can choose not to reshare harmful information, but what power do we have to avoid hearing it in the first place? We might mute a topic keyword or block a person, but existing tools do not map well to the constraints of lashon hara, such as distinguishing statements made about another person.

Once you do hear lashon hara, it comes with specific obligations. You are, at minimum, required to remove yourself from the situation. Online, this might mean using muting or blocking tools. However, removing oneself from the situation also serves to rebuke the person who is speaking lashon hara—and ensuring that you do not hear the lashon hara can be counterproductive with respect to this goal. A recent test of Reddit's new blocking tools, for example, found that when people “walked away” from someone sharing hate, it only made it easier for that person to find an audience who appreciated their posts, and to be boosted to the top of the subreddit where new members are sure to encounter them. This is exactly the opposite of how community response to lashon hara is supposed to work.

If we designed online speech systems with lashon hara in mind, they might instead empower communities to identify and enforce *behavioral* and *contextual* boundaries, rather than emphasizing individual intent. For example, new members can be enculturated into appropriate norms through shared activities that do not require exposing them to derogatory information about others. They could converse about a relevant topic, be exposed to exemplars of community behavior in a structured way, or engage with a community symbol or artifact (see [5]).

WHY SO SERIOUS?

Jewish law treats violations of speech ethics very seriously. In stark contrast to the saying “Sticks and stones may break my bones, but words can never harm me,” Jewish law compares the lashon hara to murder. In part, this comparison is meant to emphasize the severity of the violation; in a similar vein, Islam compares *gheebah* (backbiting) and *nameemah* (slandering) to eating the other's flesh. In Judaism, however, there are specific *halachic* reasons for the comparison, and understanding them can give us further insight into treating lashon hara as a community problem.

A rabbinic parable about lashon hara goes, more or less, like this: A person visits their rabbi to ask why lashon hara is taken so seriously. In response, the rabbi asks them to rip open a feather pillow. When they comply, feathers go everywhere—up the chimney, under the couch, out the window, you name it. The rabbi then asks the visitor to collect all the feathers. “That's impossible!” exclaims the

Halacha asks us to consider impact in light of context instead of trying to divine what someone else intended. Intent is unknowable.

visitor, and the rabbi responds, “Exactly. You will never find all the feathers. Nor can you track down all the impacts of your speech.”

To fully understand the lesson of this story, and the nature of the comparison to murder, it is necessary to understand the Jewish process of atonement. For violations that affect other humans (as opposed to violations of ritual law, e.g., eating pork), the 12th-century scholar Maimonides lays out a series of steps [6]. Before you may even apologize, let alone be forgiven, you must admit your wrongdoing, make restitution, and refrain from doing wrong in a similar situation. You have to both make right what you did and provide evidence that you are committed to not doing it again, before you can atone.

The parable of the feathers illustrates what murder and lashon hara have in common: It is not possible to make restitution for what you did. In the case of murder, the person you have wronged is dead. The dead by definition cannot forgive, nor may others forgive on their behalf. In the case of lashon hara, it is not possible to track down all the ramifications of your speech. You do not know who has heard it, nor what use they made of it, nor to whom they passed it on. The harm cannot be successfully addressed.

In an analog world, this analogy holds. Just as it is impossible to speak with the dead, it is impossible to identify all the ripple effects of speech. A digital context, however, might create an interesting distance between the two cases. The “right to be forgotten,” for example, uses searching and tracing of digital data to ensure that past information about a person is not available, even if it is true. One might imagine that if every feather from the pillow had an individual tracker, the person who needed to collect them would have a much easier job of it. This approach would not necessarily reduce the negative impacts of the lashon hara—others might be affected by the information before the one atoning tracks it down—but it might provide new avenues for repair and help prevent further harm. It is even possible that if people knew their lashon hara could be traced back to them, they might sometimes refrain from speaking it in the first place.

WHEN TO SPEAK

Not all true statements about another person are forbidden, even if they are disparaging. The *Chofetz Chaim* identifies cases where lashon hara is permissible:

- To get help for the subject of the statement
- To protect someone else from the subject
- To help someone already harmed by the subject
- To help end an ongoing dispute *if* it could escalate to a community-wide conflict
- To help others learn from the subject’s mistakes.

The *Chofetz Chaim* recognizes that these exceptions can be misused, for example by seeking “help” from someone’s enemies. It outlines a set of conditions that must be met before lashon hara becomes permissible, such as trying other avenues of correction first. The speaker must also identify why they are saying what they are saying, such as articulating that their goal is to seek help and they have already tried other avenues. If all of the *Chofetz Chaim*’s conditions are met, the speaker may speak true statements no matter how derogatory, and a listener may listen without fear of transgression.

What is notable about these exceptions is that they are about *care for the community*. They recognize that an individual’s behavior has an impact on others. Preventing them from causing harm, and remediating harm they may have done, are acceptable reasons to risk derogatory speech. Additionally, community norms are set both by action and inaction. Failing to deal with bad behavior not only sends a signal that the community approves but can also attract similar bad actors, who think (correctly!) they have identified a safe place to misbehave.

This approach lines up with the “whisper network” often used by marginalized people when they are experiencing abuse by someone from a dominant group. In contemporary culture, these networks are often dismissed as “just gossip” or are seen as an unacceptable alternative to public confrontation. From a halachic point of view, those networks are in fact fulfilling a community obligation.

Although a detailed treatment is beyond the scope of this paper, these ideas connect to the concept of *tochacha*, or one’s obligation to correct another. On the one hand, if someone is speaking lashon hara in front of you, you have an obligation to correct them. On the other hand, permission to speak lashon hara can be part of the correction process.

FROM THEORY TO PRACTICE

What does it mean to live by these laws in practice?

Translating rules to behavior can be very challenging. For example, it is arguably nearly impossible to make a statement about another person on social media without violating lashon hara. Posting online means losing control of where one’s words can go, and at what scale. Some Jews respond to this by avoiding social media—either as an individual decision, or by forbidding it at the communal level. This choice can be isolating, particularly when restrictions on social media are used as a tool of social coercion. Other Jews retreat to private online spaces where they trust that others will behave ethically, but their success relies on the strength of their judgment about who is allowed in. Still other Jews engage in public online life, but focus on talking about issues or causes rather than people, and accept the risks that their words may be twisted. And, of course, there are Jews who ignore the laws of lashon hara entirely. Rules to protect against harmful speech, implemented by communities, individuals, or digital platforms, won’t be interpreted or applied by all Jews in the same fashion, and can’t address every possible case of lashon hara.

At the same time, we believe that rules matter, even in the breach. For HCI researchers, this is often easiest to see when those rules are instantiated in code. For example,

It is arguably nearly impossible to make a statement about another person on social media without violating lashon hara.

Twitter's misinformation labels indicate when a tweet is misleading or deceptive. Individuals can ignore or circumvent those labels, but the existence of the feature shifts the center of gravity of the larger community and conveys expected community norms. But when it comes to interpersonal ethical conduct, people sometimes claim that rules don't matter. What's needed is individual introspection, sometimes framed as an "inward turn" or, worse, as "godliness."

Judaism agrees that an "inward turn" is critical for developing values. However, Judaism also argues that it is not a practical route to improving behavior for most people. The core philosophy of Judaism is about learning through action. In the biblical story of Mount Sinai, the Jews pledged *na'aseh v'nishma*—we do and *then* we hear. By embodying our rituals and regulating our daily actions, we enculturate ourselves to the mindset behind them. Maimonides sums this up by arguing there are two paths to enlightenment. One is to be a philosopher king who can attain enlightenment through reason alone. Most people can't do that, so the alternative is to grow into it through the lived and embodied system of Jewish law.

Godliness is a much more difficult concept for Judaism. The Jewish notion of godliness is not like that of other religions. Our role models argue with God, contradict God, wrestle with God—a famous parable even has God rejoice when the rabbis rebuke him. When people say godliness, they often are actively excluding us. But there's also the fact that Judaism does not require belief in God. Judaism is the religion of the Jewish *people*, and does not fit well into modern religious definitions. Agnostic and atheist Jews are common, both among secular and religiously observant groups. Our relationships with the divine do not change what we observe, our identities as Jews, or the fact that our lives are shaped by Jewish values. "Belief in one or fewer gods" is the hard barrier; if you do believe in a God, it must be the God of the Jews.

Given all this, should we use Jewish approaches to inspire new designs? On the one hand, designing from non-dominant perspectives can help us innovate, see new perspectives, and center the needs of the marginalized. On the other hand, Jewish obligations, the mitzvot, are incumbent only on Jews. Non-Jews are expected to follow the seven Noahide Laws, which include prohibitions on murder and a requirement to set up a just court system. However, speech laws (other than those relating to false testimony) are not included in these requirements.

We would argue that even though the concept of lashon hara is not binding on gentiles, it can nonetheless help us collectively create wiser, healthier, and kinder communities. Judaism is a closed practice, limited to the Jewish people. Outsiders are not welcome to perform their own *seder* (the Passover ritual) or to claim Shabbat for themselves. However, Jews also live in a society. We want to participate in communities that are successful, vibrant, and ethical. That means sharing our ethical perspectives with those who share our social context so that we can co-create communities where Jews are welcome—including those Jews who are multiply marginalized, whose identities as Black or trans or queer benefit even more from Jewish

ethical insights into communal speech obligations. None of us can build a welcoming online community alone.

CONCLUSION

As the examples show, a Jewish framing of speech ethics focuses not on what one has the *right* to do, but rather on one's positive obligations and negative prohibitions. This framing helps turn a complex and ambiguous topic into behavioral guidelines. The examples also show a focus on the community impacts of individual behavior. Allowing someone to speak lashon hara in your presence not only spreads potentially derogatory information but also creates a community norm that it is acceptable to speak ill of others. While Jewish communities often struggle to implement these laws in practice, they can still inspire new designs and change our baseline assumptions about how online speech should work.

Before closing this article, a note about our perspectives. Between the two of us, we have experience with the Orthodox, Conservative, and Reform denominations of Judaism (which gives new meaning to the saying "two Jews, three opinions"). Having perspectives from multiple denominations has helped us see commonalities across different types of Jewish thought and practice. However, Judaism is a religion of pluralism and dissent. This article is not intended to be fully definitive of Jewish approaches to speech, but rather to begin a conversation. We look forward to continuing this work together with all of you.

ACKNOWLEDGMENTS

With profound thanks to Kaitlin B. Heller, a role model for speaking up to safeguard others.

ENDNOTES

1. Hammer, J. Envisioning Jewish HCI. *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2020, 1–10.
2. Pliskin, Z. *Guard Your Tongue: A Practical Guide to the Laws of Lashon Hora*. S. Weissman, Brooklyn, NY, 1977.
3. Telushkin, J. *Words That Hurt, Words That Heal: How to Choose Words Wisely and Well*. Harper Collins, 2010.
4. Marwick, A.E. and boyd, d. Networked privacy: How teenagers negotiate context in social media. *New Media & Society* 16, 7 (2014), 1051–1067.
5. Seering, J., Luria, M., Ye, C., Kaufman, G., and Hammer, J. It takes a village: Integrating an adaptive chatbot into an online gaming community. *Proc. of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2020, 1–13.
6. Abramson, H.M. *Maimonides on Teshuvah: The Ways of Repentance*. Henry Abramson, 2017.

✉ **Jessica Hammer** is the Thomas and Lydia Moran Associate Professor of Learning Science at Carnegie Mellon University, jointly appointed in the Human-Computer Interaction Institute and the Entertainment Technology Center. Her research focuses on transformational games; she is also an award-winning game designer.
→ hammerj@andrew.cmu.edu

✉ **Samantha Reig** is a Ph.D. student in the Human-Computer Interaction Institute at Carnegie Mellon University. She researches human-agent interaction in socially complex contexts, and is advised by Jodi Forlizzi and Aaron Steinfeld. Reig received her B.A from Cornell University in 2017, and an M.S. from CMU in 2020.
→ sreig@andrew.cmu.edu

