The Structure of Open Secrets

Sam Berstler
Massachusetts Institute of Technology

Abstract. In conversation, we often do not acknowledge what we jointly know to be true. My aim in this paper is to identify a distinctive kind of non-acknowledgment norm, *open secrecy*, and analyze how such norms constrain our speech. I argue that open secrecy norms are *structurally* different from other everyday non-acknowledgment norms. Open secrecy norms iterate: when *p* is an open secret, then there's a norm not to acknowledge that *p*, and this norm is itself an open secret. The non-acknowledgment at issue in open secrecy norms, I argue, motivates a more complex understanding of discourse. When interlocutors are conforming to open secrecy norms, they rely on at least two disjoint common grounds, one of which has a privileged status. To understand why and how it is privileged, I develop Erving Goffman's notion of *defining a social interaction*. Finally, I show how strategic speakers can exploit the structure of open secrecy in order to both communicate about the open secret and shield themselves from retaliation for what they communicate.

Keywords: open secret; non-acknowledgment; implausible deniability; strategic speech; common ground; conversational norms

Word count: 15,410 words

I. Introduction

Jane's meeting with her human resources representative, Wilcock, is turning into a nightmare (*The Assistant* 46:44-57:40). Jane is a naïve recent college graduate and assistant to a powerful movie producer. She's also just realized that her boss has hired a new "girl" in order to have sex with her.

Jane has requested a meeting with Wilcock in order to report her boss's misconduct. But Wilcock doesn't seem to understand any of Jane's insinuations. When Jane becomes blunter, he changes tactics. He implies that Jane is endangering her current job, insists that her complaint has no evidential merit, accuses her of jealously, and finally exhorts her to remember that the "girl" is a "grown woman" who can make her own decisions. It becomes clearer and clearer that Wilcock understood Jane the whole time. As a defeated Jane abandons her complaint and walks out the door, we learn the truth:

HR 1
Wilcock: I don't think you have anything to worry about.
Jane: Hmmm?
Wilcock: You're not his type.

(*The Assistant* 57:36)

The Assistant depicts a thinly fictionalized Miramax, in the years before the Harvey Weinstein scandal broke. As Scott Rosenberg, a Hollywood producer, would later acknowledge "Everybody-f**king-knew" that

Weinstein sexually harassed female actors (Kiefer 2017). It was "Hollywood's biggest open secret" (*Axios* staff 2017). According to media reports, people didn't talk openly about it because of a "tacit code," "wall," or "conspiracy" of silence, around Weinstein specifically and sexual harassment generally (Zacharek 2017; Twohey et al 2017; Nyong'o 2017).[1]

But *The Assistant* doesn't depict a workplace literally silent on the issue. The boss' misdeeds are the constant subject of off-color jokes and gossipy office murmurings. "*Never* sit on the couch," one production executive counsels another (*The Assistant* 44:25). In the real world, inside jokes about Weinstein worked their way into the scripts of major comedy shows like *30 Rock* and *Entourage* (Yahr 2017). When Jane and Wilcock face-off, Wilcock doesn't, as you might expect, seek to convince Jane that her boss isn't a sexual harasser. He continuously hints at and then confirms the truth to Jane. This isn't a literal silence but a kind of group-wide refusal to state too much too openly.

To explain both the nature and pattern of this pseudo-silence, it's tempting to focus on two factors: first, individuals' fear that Weinstein would retaliate against them; and second, in the case of someone like Wilcock, individuals' desire to minimize their liability for violating workplace sexual harassment laws. On this analysis, the industry-wide pseudo-secrecy around Weinstein was an emergent phenomenon, arising out of speakers' individual desires to preserve their own plausible deniability for what they were communicating.[2] But while I don't deny the importance of these factors, I will argue that they aren't the full explanation either. Cultures of silence like the one in Hollywood, I will argue, are *literally* a matter of culture, arising because of social norms enforced through patterns of shame, disesteem, reputation management, and retaliation.

In this paper, I aim to identify what these norms are and analyze how they constrain our speech. I'll call these norms *open secrecy norms*. I'll say that the information that they target is the *open secret*. In particular, I'll aim to understand the sense in which a group conforming to open secrecy norms both does and does not acknowledge the open secret. The non-acknowledgment at issue in open secrecy norms, I argue, motivates a more complex understanding of discourse. In particular, I complicate the presupposition and common ground framework that Stalnaker has developed throughout his oeuvre (Stalnaker 1984, 1999, 2002, 2014). To do this, I develop ideas from the work of mid-century sociologist Erving Goffman (1959, 1967, 1974). Thus, while this paper is about open secrecy narrowly, I aim to use my analysis in order to advance our understanding of the structure of conversation broadly.

In the first part of the paper (sections II and III), I'll hypothesize that open secrecy norms are a species of what I call non-acknowledgment norms. Non-acknowledgment norms are a type of social norm that guides our conversational behavior. Non-acknowledgment norms sometimes require us to use polite euphemisms, avoid conversational subjects likely to trigger conflict, engage in paper-thin pretenses of mutual goodwill, or otherwise navigate around a "background of…'unmentionables' " such as our differing financial status, our sexual lives, or taboo bodily functions like defecation (Camp 2018: 57). But open secrecy norms, unlike run-of-the-mill non-acknowledgment practices, iterate. When p is an open secret, there's a social norm not to acknowledge that p, and this norm is itself an open secret. As one sociologist of denial explained, the original silence around Weinstein's misconduct, a father's alcoholism, or a government's torture program infinitely spirals, creating "meta-silences" upon "meta-silences" (Zerubavel 2006: 53).

---

[1] Press reports about Jeffrey Epstein used similar language. Thus, it was an "open secret" in Palm Beach, Florida that Epstein hired underage girls for sex work (Ramadan 2019). The secret became even more open when, as part of a 2007 plea deal, Epstein was placed on the sex offender registry but still retained his place in high society (Reilly 2019). In a recent complaint filed against the United States government, attorneys for Epstein's victims allege that even the FBI "disregarded" that "Epstein's penchant for teenage girls was an open secret" (quoted in Ruiz 2024).

[2] For recent philosophical work on plausible deniability, see Walton (1996); Pinker (2007;(Pinker et al (2008); Lee and Pinker (2010); Fricker (2012); Peet (2015); Camp (2018); Berstler (2019); Davies (2019); and Dinges and Zakkou (2023).

In the second part of the paper (sections IV and V), I ask: what is the relevant notion of "acknowledge" in the injunction "do not acknowledge the open secret"? It turns out that this is a significantly challenging question. Simple norms like, "Do not assert the open secret" or "Pretend that the open secret isn't true" don't work. To understand the nature of open secrecy, I argue, we must better understand the nature of conversation itself.

I propose that in an open secrecy conversation, interlocutors exploit multiple common grounds, which bear a hierarchical relationship to each other. To do this, I borrow the notion of *activity layering* from Goffman (Goffman 1974, Clark 1996). Then, I argue that we need a way to privilege one of these common grounds. To do this, I argue, we need Goffman's independently motivated notion of *defining the social interaction* (Goffman 1959, 1967). We thus arrive at this analysis: if *p* is an open secret, do not define the social interaction as one in which *p*. Finally, I close by showing how strategic speakers sometimes exploit the structure of open secrecy in order to both communicate about the open secrecy and protect themselves from retaliation for what they communicate.

II. Bringing open secrecy into focus

Open secrecy norms are a kind of non-acknowledgment norm. Suppose that you know that your departmental chair is embezzling money. You might refuse to acknowledge this fact to others in the sense that you refuse to *tell* or *reveal* this fact to them. That's not the kind of acknowledgment that's at issue in this paper. As I'll understand them, non-acknowledgment norms tell you how to handle knowledge that you and your interlocutor already share.

If you are an outsider to a group, open secrecy can be difficult to diagnose. This is not just because open secrecy norms prohibit us from openly discussing their existence. It is also because individuals conforming to open secrecy norms can seem behaviorally indistinguishable from individuals suffering from denial, which in the popular imagination is a "twilight" zone between knowing and not knowing (Cohen 2001: 9).

Consider this story:

> **Production staff**
> Ethan is a senior member of a production staff. He repeatedly cautions his staff not to send female actors to Russell's dressing room. At parties, Ethan watches Russell like a hawk and guides inebriated female actors away from Harvey. But if asked, Ethan will deny that Russell is a sexual predator, and his denial will seem sincere.[3]

One way to explain this pattern of behavior is to attribute some kind of epistemic dysfunction, fragmentation, or mental opacity to Ethan. For example, maybe Ethan literally believes, in one sense, that Russell is a predator and literally believes, in another sense, that Russell isn't (Elga and Rayo forth).[4] Maybe Ethan believes that Russell is a predator but also has some belief-like state that Russell isn't (Gendler 2008). Maybe he in-between believes both that Russell is and isn't a predator (Schwitzgebel 2010). Or maybe Ethan knows that Russell is a sexual predator but doesn't know that he knows this.[5] If the majority of people in a group have one of these epistemic problems, then we're not dealing with an open secret.

---

[3] This example is adapted from details in the podcast *Pod Save the UK,* in which the hosts described the comedy world's "open secret" about Russell Brand's alleged sexual misconduct (Kumar and Khan 2012). Thanks to [removed for blind remove] for drawing my attention to it.

[4] For example, on Arendt's interpretation of the prolific use of Nazi euphemisms, the euphemisms functioned to fragment Nazi officials (Arendt 1994 [1965]: 85). When the euphemisms worked, then relative to one body of knowledge, a Nazi higher-up knew he was murdering but didn't know it was wrong, and, relative to another body of knowledge, knew murder was wrong but didn't know he was murdering.

[5] For a nuanced comparison of these strategies, see Greco (2014).

Open secrecy, in my sense, also doesn't arise under conditions of pluralistic ignorance. Here's a kind of silence that isn't an open secrecy silence:

> **Holiday repression**
> Coralie's extended family used to have explosive and fruitless politic arguments at holiday gatherings. So at some point, they tacitly adopted a rule: at holiday gatherings, don't bring up a political subject if it will cause acrimonious disagreement. Coralie believes that the president's latest immigration policy is appalling. Coralie *also* believes that each member of her extended family disagrees with her. So Coralie doesn't discuss it at family gatherings. Unbeknownst to Coralie, each member of Coralie's family also believes that the president's immigration policy is appalling. But they also each believe that they are the only person to believe this. So everyone in the family disapproves of the policy but no one discusses it.

Pluralistic ignorance is a form of ignorance, and this ignorance is what explains the family's silence.[6] But when *p* is an open secret, the majority of the group shares their knowledge that *p*, and they can rationally rely on *p* in order to coordinate. Open secrecy norms guide them in *how* to make use of this shared information. For simplicity, I'm going to assume that this means that the majority of the group *commonly* or *jointly knows* that *p*:

> **Common / Joint Knowledge**
> A and B jointly know that *p* iff
> A and B know that *p*,
> A and B know that A and B know that *p*,
> A and B know that A and B know that A and B know that *p*,
> *et cetera*
>
> (Lewis 1969; *cf.* Stalnaker 2001)

Often, open secrecy norms on *p* arise when group members are continuously in contact with overwhelming evidence that *p*:

> **Alcoholism**
> A mother, Eloise, arrives home with her son, Mark. Eloise's husband, Jerry, is asleep on the couch. The house is a mess. Alcohol bottles are strewn everywhere. What's happened is obvious: Jerry drank too much and passed out. Eloise and Mark silently clean up the house. The next day, they don't talk about what happened. This scene repeats over and over again throughout Mark's childhood.
>
> (adapted from Zerubavel 2006: 7)

Episodes of joint perception—in which I'm aware of *p* and you're aware of *p* and we're both aware of each other's awareness of *p*, and so on—paradigmatically generate common knowledge of *p* (Peacocke 2005). Since Mark and Eloise commonly know that there's overwhelming evidence that Jerry is an alcoholic, they plausibly also commonly know that Jerry *is* an alcoholic.

---

[6] The norms that emerge under conditions of pluralistic ignorance are a major topic of study within economics and the social sciences. Representative discussions include Katz and Schank (1938), O'Gorman (1986), Noelle-Neumann (1993), Miller and Prentice (1994), and Bicchieri (2016).

The information that the open secrecy norm targets isn't always obvious. Rosenberg claims that what everyone jointly knew wasn't that Weinstein was a rapist but something weaker: that Weinstein had "a certain pattern of overly-aggressive [sexual] behavior" (Kiefer 2017). The hit HBO show *House of the Dragon* offers a fascinating case study. In the story, the open secret is not that the crown princess' children are illegitimate but rather that they overwhelmingly *look* like they are illegitimate:

> **Illegitimate heir**
> In the fictional country of Westeros, the children of the crown princess, Rhaenyra, are obviously illegitimate. In this fantasy world, two parents with silvery blond hair (associated with a particular magical dynasty) always have children with silvery blond hair. Rhaenyra and her husband both have this hair. Their children have brown hair. Moreover, while Rhaenyra is racially white, her husband is racially Black. None of Rhaenyra's children have racially Black features. Finally, Rhaenyra's children bear striking resemble to Rhaenyra's favored knight, Harwin Strong. Despite some quiet grousing, there is a long period of time during which most nobles pretend that there is no reason to question the legitimacy of Rhaenyra's children.
>
> (*House of the Dragon* 2022)[7]

appearance:	the appearance of Rhaenyra's children is evidence that they are the children of Harwin Strong.
illegitimacy:	Rhaenyra's children are illegitimate.

How can we tell that the open secret is *appearance* and not *illegitimacy*? First, as I'll elaborate in the next section, the hallmark of open secrecy is the fact that we can't talk about the open secrecy norm. But the king openly institutes a rule against discussion of *illegitimacy*. In Westeros, it's acceptable to say, "You must not question the legitimacy of the crown princess' children!" But it's not acceptable to say, "You must not draw attention to the visual resemblance between the crown princess' children and Harwin Strong!"

Second, courtiers cannot and do not discuss *appearance* even in order to deny *illegitimacy*. For example, an assertion like (1) is socially unacceptable in Westeros:

> 1. Yes, it's true that Luke and Jace look an awful lot like Harwin Strong. But Rhaenyra would never bear a child out of wedlock. That's a ridiculous conspiracy theory.[8]

Courtiers primarily know *illegitimacy* on the basis of knowing *appearance*. So since they cannot assert *appearance,* they characteristically cannot assert *illegitimacy* either. If someone were to ask them how they knew *illegitimacy,* they could not justify their assertion by asserting *appearance*.

In principle, open secrecy norms can target even subtler information. Here's an example of a meta-linguistic open secret:

> **Torture**
> At war with sectarian rebels, a government initiates a program to "intensively interrogate" detainees. If asked whether they are torturing detainees, the government and interrogators will deny that they are. When talking to each other, they also all describe what they are doing as intensive interrogation. But everyone involved in the program understands and understands that everyone understands that what they are really doing is torture.

---

[7] On my analysis, media coverage that asserts "their parentage is…an open secret (McAtee 2022) and "the truth of Rhaenyra's children's parentage seems to be King's Landing's open secret" (Gugliersi 2022) is strictly speaking incorrect.
[8] In one scene in which the king is all but forced to confront the issue, he communicates (1) only through a roundabout parable in which a black mare and white stallion produce a brown horse ("The Princess and the Queen" 20:31).

(adapted from Cohen 2001: 107)

The open secret, in this case, is something like:

> *torture:* "torture" is an accurate way to describe what we are doing.

Sometimes the open secrecy is neither factual nor meta-linguistic but normative. This category corresponds to what the sociologist Stanley Cohen calls "interpretative" and "implicatory" denial (Cohen 2001: 7-8). For example, the group might acknowledge that rape happened during the Bosnian genocide but, since "war is war," refuse to acknowledge that it was wrongful. Or they might acknowledge that a military intentionally killed civilians but refuse to apply thick terms like "massacre" or "slaughter" to the killings. Again, we are dealing with open secrecy norms only if the majority of the group simultaneously jointly knows that the rapes were wrongful and that the killings were massacres and slaughters.

Finally, open secrecy norms are not necessarily morally pernicious. Sometimes, given non-ideal background conditions, an open secrecy norm may have more good consequences than bad. They may even be morally justified. Here's an example of a not-obviously-pernicious open secrecy norm:

> **Unfair will**
> Amelia's grandmother Rose recently passed away. Rose left a will that unfairly advantages her favorite son, Chet. No one, least of all Chet, thinks that Chet deserves the extra money. He wasn't a good son to Rose, he has much more money than the rest of the family, he didn't invest in Rose's end of life care, and he's a jerk. The family jointly knows that the will is unfair but also jointly knows there's no point in confronting Chet about it. Chet would just cause a scene and then withdraw from the extended family.
>
> After the funeral, Chet says, "I'm glad to see that Rose left such a fair will. It would be a shame to squabble about money." Without any sarcasm in her voice, Amelia replies, "Yes. Thank goodness." For years, the entire family keeps up the pretense that the will is fair, even when Chet isn't in the room.

(adapted from Michaelson 2018: 184)[9]

Of course, it would be better for everyone if Chet had the courage to own up to his own bad behavior. But given that Chet won't, how should the rest of the family respond? As Thomas Nagel emphasizes, sometimes refusing to acknowledge obvious conflict points is the only way to preserve important relationships (Nagel 1998). The family's silence on the unfairness of the will is a way to stave off "a direct collision, filled with reproaches and counterreproaches, guilt and defiance, anger, pity, humiliation, and shame, which their intimacy would not survive" (Nagel 1998: 15-16).

[9] In Michaelson's version, the family members openly complain when Chet isn't present. I've changed it so that the case clearly describes open secrecy. The anecdote is drawn from Michaelson's personal life (Michaelson 2018: 205fn4).

III. The iterative structure of open secrecy

In order for *p* to qualify as an open secret, certain social norms must exist. Therefore, just because a group doesn't talk about something they jointly know and doesn't talk about the fact that they don't talk about it, it doesn't follow that *p* is an open secret in the group. Maybe they simply have no need to discuss it.

Social norms are not just conventions, regularities, or defaults. How to distinguish the former from the latter is contentious; I'll try to remain somewhat neutral on the matter. Here's how, in general, I'm thinking about social norms.[10] First, we generally *enforce* social norms against each other. We do this through directly sanctioning people who fail to comply—sometimes subtly, as when I give people the stink eye for talking too loudly on the subway and sometimes more aggressively, as when we ostracize someone. But we also enforce social norms indirectly, through esteeming conformers and disesteeming non-conformers, on the one hand, and revealing our opinion to others through gossip, on the other. We often conform to social norms, in other words, to build a good reputation—as a team player, as other-oriented, or even just as a predictable actor. Finally, if we're properly socialized into our community, then when we break a social norm, we characteristically feel shame or self-revulsion (Elster 1989, 1999). That's not to say that we *should*, in a moral sense, feel such shame. Many social norms are morally pernicious.[11]

As a first pass, here's my analysis of open secrecy. I'll leave a detailed analysis of the (c) clause until the next section. For now, I'll rely on an intuitive understanding of "acknowledge:"

Open secrecy
P is an open secret in *g* iff
(a) *p*,
(b) the majority of *g* jointly believes *p*,
(c) there is a social norm in *g* not to acknowledge *p*, and
(d) (b) and (c) are open secrets.

Here are some details about the definition.

First, clause (a) guarantees that open secrecy is factive: if *p* is an open secret in *g*, then *p*. This brings my definition into line with ordinary language usage:

2. # In that family, it's an open secrecy that Joe is an alcoholic, but he's not.

Of course, it's possible that a group jointly believes that *p*, conditions (c) and (d) otherwise hold, and *not p*. In such a case, we can say that the group merely *treats p* as an open secret.

Second, I am going to assume this principle:

---

[10] Two developed philosophical accounts of social norms, which are both broadly compatible with the sketch in this paragraph and which are associated with different corners of philosophy, are the ones in Brennan and Pettit (2004) and Pettit (2019), on the one hand, and in Haslanger (2021, 2023), on the other. For the centrality of gossip in incentivizing norm conformance, see Dunbar (2004, 1996).

[11] Why might someone conform to a social norm? My account is compatible with many possible answers, and different people might conform for different reasons. Someone might conform for prudential reasons: she might want to avoid the penalties that norm-violators incur and reap the rewards that norm-conformers attain. She might conform for broadly moral reasons: she might believe (rightly or wrongly) that she is morally obligated to follow the social norm. She might conform out of unreflective habit. Or she might conform because she values the existence of or conforming to the social norm for some further reason: because she values what the social norm expresses about the group, for example, or what the norm enables the group to do. Thanks to an anonymous referee for encouraging me to discuss this.

Social norm transparency
If *p* is a social norm in *g*, then a majority of *g* jointly knows that *p* is a social norm in *g*.

This principle captures the assumption that when *p* is a social norm in *g*, then characteristically interlocutors can rationally rely on this fact in order to coordinate their behavior. If social norms didn't have this public status, they couldn't play one of their most essential functions. However, this principle doesn't require most members of *g* to have the capacity to *articulate* that *p* is a social norm in *g*. Their knowledge may be tacit or implicit, in the same way much linguistic knowledge is tacit or implicit.

Social norm transparency enables us to ignore condition (b) when we check whether the iteration condition in (d) holds. Assume (as seems reasonable) that knowledge entails belief. If we establish that there is a social norm not to acknowledge that there is a social norm not to acknowledge *p,* then, by the social norm transparency principle, we have established that the majority of *g* jointly believes there is a such a social norm. In other words, if condition (c) holds, so does condition (b). Similarly, if we establish that there is a social norm not to acknowledge that the majority of the group jointly believes *p*, we have established that the majority of *g* jointly believes that there is such a norm.

From this definition, we can also derive the following principle:

Open secrecy iteration
If *p* is an open secret, then it's an open secret that *p* is an open secret.

To see why, suppose that *p* is an open secret. Is it an open secret that *p* is an open secret? Given our supposition, clause (a) of the open secrecy definition follows immediately.

Second, does clause (b) hold? By the definition of joint belief, if the majority of *g* jointly believe *p*, then they jointly believe that they jointly believe that *p*. By the social norm transparency principle and the assumption that knowledge entails belief, they jointly believe that there's a social norm not to acknowledge that *p*. We can repeat the same reasoning for each norm that the (c) clause generates. So clause (b) holds.

Now let's establish clause (c): if *p* is an open secret in *g*, then there is a social norm in *g* not to acknowledge that there's an open secret in *g*. All we need here is one observation and one assumption. The observation: that if *p* is an open secret in *g*, then *p*. This follows from clause (a). The assumption: if you must not acknowledge *p* and *q* entails *p,* then you must not acknowledge *q*. This assumption is independently plausible. So if there is a social norm not to acknowledge *p*, and *p* is an open secret, then there is a social norm not to acknowledge that *p* is an open secret. Finally, we can establish clause (d) using the same procedure we used above.[12] So if *p* is an open secret, then it's an open secret that *p* is an open secret.[13]

Given this structure, here are the sorts of speech acts that open secrecy norms prohibit. Members of the group sometimes *do* make assertions like (3, 4, 5). But, if they do, they are generally acting in a socially unacceptable manner:

3. Russell is a sexual harasser.
4. Around here, we don't talk about the fact that Russell is a sexual harasser.

---

[12] To see why (b) is an open secret, observe that if the majority of the group jointly believes that it's an open secret that *p*, the majority of the group jointly believes that *p*. Using the same principle above, it follows that if there's a norm not to acknowledge that the majority of the group believes *p*, there's a norm not to acknowledge that the majority of the group believes that *p* is an open secret.

[13] We could derive this principle even if we dropped condition (d) in the original definition. However, this analysis would incorrectly predict that assertions like (6) and (7) do not violate the open secrecy norm.

5. In our industry, it's an open secret that Russell is a sexual harasser.

Open secrecy iterates over the (b) and (c) clauses. What we want is for open secrecy norms to prohibit discussion of the open secrecy norm, *even when* the speaker doesn't presuppose the truth of the open secret. For example, someone who says (6) and (7) should violate the open secrecy norm, even though the speaker of (6) doesn't presuppose the truth of the accusation:

6. We don't tolerate the accusation that Russell is a sexual harasser. Those kinds of accusations are bad for morale.
7. Around here, everyone is completely convinced[14] that Russell is a sexual harasser. But I assure you that that's the result of a vicious smear campaign. It's nonsense.

Right now, we can't achieve this result, because in ordinary English, I acknowledge that *p* only if I presuppose the truth of *p*. In the next section, when we analyze the target notion of "acknowledgment," we will need to find an analysis without the factivity presupposition.

Open secrecy norms do permit members to say this:

8. You shouldn't say that Russell is a sexual harasser. That accusation is totally false.

On the most plausible reading, the speaker of (8) is admonishing her interlocutor to conform to the knowledge norm on assertion (Williamson 2000). She is pretending that she herself knows that Russell isn't a sexual harasser. It's acceptable to exhort group members to conform to norms not to discuss *p*, so long as those norms aren't the norms that constitute the open secrecy. When we refine the target notion of "acknowledgement" in the next section, we'll have a better idea of how this works.

Nonetheless, under conditions of open secrecy, many speakers will often choose not to say (8). If the speaker says it out-of-the-blue or if she fails to sound appropriately serious in saying it, others may interpret her as conversationally implicating that Russell *is* a sexual harasser or that there's a widely shared rumor to this effect. ("The lady doth protest too much, methinks.") If she makes either implicature, then she violates the open secrecy norm. Thus, a speaker who wants to minimize her social risk, especially one who isn't confident that she can keep a straight face while telling egregious lies, may try to avoid the subject as much as possible. Alternatively, some speakers who otherwise conform to open secrecy norms will not say (8) for moral reasons. They're willing to *refrain* from telling the truth but not to lie. (If the open secrecy norm ever falls apart, they can also insist, in a somewhat lawyerly fashion, "Well at least I never *denied* that it was true!")

A group member who either doesn't know that *p* is an open secret or doesn't want to conform to the open secrecy norms can quickly find herself frustrated. Imagine a conversation between Amelia, who is conforming to the open secrecy norms, and her cousin Gerald:

Unfair will 2

Gerald:	Look, can we just talk openly about the will?
Amelia:	I don't know what you're talking about. That's what we're doing right now.
Gerald:	We both know that the will was unfair.
Amelia:	That's not what I think.

[14] I'm assuming that in context, the speaker of (7) intends to convey that nearly everyone *jointly* believes that Russell is a sexual harasser. We don't generally use this term in idiomatic speech, so the assumption is reasonable. Note that if the speaker had said "Everybody secretly believes that Russell is a sexual harasser, but they think they're the only one who believes it. But it's all a big misunderstanding," she wouldn't violate the open secrecy norm. But see the discussion in text for why many speakers would still choose not to say this.

Gerald: I know that you're just saying that to avoid conflict.
Amelia: This conversation isn't going anywhere.

Here's how one real-life whistleblower described the open secrecy in her company:

> I said I'd do anything [my boss] wanted—keep silent, resign, ask for a transfer. All he had to do was discuss the issue with me. But he wouldn't do it, and he wouldn't talk about not doing it. My insubordination was the only issue.
>
> (Alford 2001: 21)

She later explained "You think I'm some kind of hero 'cause I blew the whistle. The only reason I spoke up is because I didn't want to go crazy" (Alford 2002: 3).

For comparison, consider how a few other types of non-acknowledgment norms work. Here are three. I've labeled them closing-the-book norms, exclusion norms, and polite euphemism norms.

Closing-the-book norm
A group *g* has a closing-the-book norm or agreement on *p* iff
(a) the majority of *g* jointly knows that there's significant disagreement whether *p*, and
(b) there's a social norm or agreement in *g* not to discuss whether *p*.

Closing-the-book norms or agreements arise when continued discussion about *p* is either unlikely to produce agreement or isn't worth the bad social effects it will generate. Closing-the-book norms and agreements only affect conversations in which the group members disagree about *p*. If you and I jointly believe that *p,* it's acceptable for us to presuppose it, since we're not *discussing* whether *p*.

Exclusion norms
A group *g* has an exclusion norm or agreement on *p,* for some set of actions *A*, iff:
(a) the majority of *g* jointly knows that *p*, and
(b) there's a social norm or agreement in *g* not to treat *p* as a reason to *A*.

Exclusion agreements are common under conditions of formal, collective deliberation. For example, a hiring committee might jointly know that Annie has a PhD from Princeton and Barry has a PhD from Indiana University. In order to combat prestige bias, however, they might agree that they will not treat these facts as reasons in favor of hiring either Annie or Barry.

Polite euphemistic norms
A group *g* has a polite euphemistic norm on some referent *r* in a context *c* iff:
(a) the majority of *g* jointly knows that some phrase 's' refers to *r*, and
(b) there's a social norm or agreement in *g* not to use 's' to refer to *r* in *c*.

For example, in professional contexts, you should say "had sex with" and not "fuck." You should also say "go to the bathroom" and not "take a dump."

None of these norms are self-camouflaging in the way that open secrecy norms are. Speakers who assert the following don't violate the relevant norm:

> 9. I know that we agreed that we wouldn't discuss our disagreements about our hiring policies anymore. But I really think we should reconsider that agreement. Otherwise, we're just going to have the same issue when we come to graduate admissions this year.

> 10. As a reminder, please don't bring up the fact that Annie has a PhD from Harvard. We agreed not to let this fact bear on our decision.
>
> 11. You really need to change the sentence, " 'Rhaenyra fucked Harwin' to 'Rhaenyra had sex with Harwin.'" It's not appropriate to use the word "fuck" in this context.

Because of this, open secrecy norms have a kind of inertia that other non-acknowledgment norms lack. If you want to lift or alter a social norm, one way to do so is to persuade fellow group members that the norm is morally unjustified, socially pernicious, or just suboptimal given the groups' other ends. But if *p* is an open secret, then group members often have prudential reason to either broach the subject only with trusted intimates or not broach it at all. This is because in broaching the subject, they risk others' social censure for violating the open secrecy norm.

The prohibition on talking *about* the open secrecy norm seems like an obviously bad-making feature. In a modern liberal society, we tend to value open debate about the nature of our shared social world. But the self-camouflaging nature of open secrecy isn't necessarily pernicious. Sometimes, a group has good moral reason to "close off" a matter permanently. Here's a story adapted from something I witnessed:

> **Fursona**
> Through Internet sleuthing, Hazel has identified the anonymous twitter account of her colleague, Paris. Hazel then discovers that Paris has a "fursona," a harmless but mildly taboo practice in which members dress as anthropomorphized animals, often for sexual gratification. Hazel gossips about this to everyone else in the office. Pretty soon, everyone jointly knows that Paris is a furry. Paris is furious and embarrassed.

In this case, given that the workplace has no other legitimate reason to discuss Paris' sex life and given the psychological relief the open secrecy norms may bring Paris and others, it is not implausible that open secrecy norms are justified here.

Open secrecy norms are an effective stopgap solution in such cases because they systematically impede our ability to transmit and pool information related to the open secret. Most obviously, open secrecy norms prohibit group members from revealing or confirming that *p* to people outside the group, even when those individuals already know that *p*. In that sense, open secrecy norms are similar to other, simpler secret-keeping practices:

> **Group secrecy norm**
> A group *g* has a group secrecy norm or agreement on *p* iff there's a social norm in *g* not to reveal *p* to non-members of *g*.
>
> **Omerta norm**
> A group *g* has an omerta norm or agreement on *p* iff there's a social norm in *g* not to acknowledge *p* to non-members of *g*.

If a group is keeping a group secret or omerta on *p*, its members may still choose to speak amongst themselves about *p* in an oblique or coded manner or on an infrequent basis. But this is generally because they are worried that those outside the group will overhear them, read their documents, or otherwise intercept their group-internal communications. Hannah Arendt, for example, hypothesized that the Nazi elite so rarely acknowledged the true nature of the Final Solution only because they were so often in the company of stenographers and junior staff (Arendt 1994 [1965]: 85). Amongst themselves, they were free to baldly state the truth.

Why would a group develop an open secrecy norm on *p*, rather than a simpler group secrecy or omerta norm on *p*? One answer is that open secrecy norms, unlike group secrecy norms and omerta norms, play a role in *habituating* group members into maintaining the secret from outsiders. This is because if I am habituated into never acknowledging *p*, I am *ipso facto* habituated into never acknowledging *p* to those who don't know that *p*. Similarly, if I am habituating into speaking about *p* only in oblique ways, I am habituated into speaking about *p* obliquely in the presence of those who might overhear me.

But the more obvious answer is that open secrecy norms don't just function to keep information internal to the group. They also obstruct our capacity to collectively address, deliberate, or share further information about the open secret.[15] If Ethan cannot acknowledge that Russell is a sexual harasser, he cannot openly justify firing Russell on this basis. Nor can he openly deliberate with his fellow production staff about whether they should fire Russell on the basis of these allegations. And if Ethan knows additional information that confirms Russell's predatory ways—for example, the existence of a second or third or fourth victim, whom his productive staff know nothing about—it's difficult for him to share it without violating the open secrecy norms.

Given this, open secrecy norms also dramatically and holistically change our epistemic environment. In doing so, they render our knowledge of *p* more vulnerable than it otherwise would have been. Surrounded by people who keep insisting that *not p* and deprived of confirmatory evidence that we might have otherwise had, we may come to doubt our knowledge that our fellow group members know that *p* or even come to doubt *p* itself. This process can occur both rationally, when our evidence that *not p* starts to mount too high, or irrationally, when we become gaslit and so stop believing that *p*. Arguably, in cases like **Torture** and **Alcoholism**, the open secrecy norms arise precisely *because* they function to erode knowledge in this way. Since to undermine joint knowledge that *p* is to undermine an open secrecy norm on *p*, the open secrecy norms are here functioning to undermine themselves.

Finally, open secrecy norms can come to play a role in signaling and constructing group identity. For a newcomer, obtaining the knowledge that your group is maintaining an open secret on *p* is a significant epistemic achievement. It requires not just realizing that *p* and not just realizing that the group jointly knows that *p* but also realizing that the group is managing this information in a peculiar way. All of this must generally happen without anyone in the group explaining, "Around here, we don't talk about the boss' sexual misdoings…" in the way they might explain other unwritten rules like, "Around here, it's technically required to submit the forms within three days but nobody looks at them until a week later…" Because of that, when someone clearly and intentionally conforms to the open secrecy norm, she provides good evidence that she has mastered the group's implicit rules.

This is one of the more surprising ways in which the self-camouflaging nature of open secrecy norms is morally and socially dangerous. It provides individuals with additional reason not just to conform to the norms but to perpetuate the conditions that give rise to them. Thus, I might conform to an open secrecy norm in my workplace not because I fear retaliation or because I want to participate in a cover-up as such but because I want to show that I belong. In some cases, a group may even individuate itself *according* to the open secrecy norm. Thus, a department or industry might have a group of "insiders" who both are and think of themselves as insiders only in virtue of the fact that they jointly know and keep an open secrecy norm on *p*. I might then keep this open secret in part because I value belonging to this clique. If the open secrecy becomes widely known and discussed, I lose an identity that I perceive as status-conferring.

---

[15] Thus, Nagel writes that the "essential function of the boundary between what is acknowledged and what is not is to admit or decline to admit potentially significant material into the category of what must be taken into consideration and responded to collectively…" (Nagel 1998: 12).

IV. The non-acknowledgment in open secrecy

If *p* is an open secret, then by definition, there is a social norm prohibiting group members from acknowledging that *p*. But what does it mean to acknowledge that *p*? We've already seen one constraint: whatever it is, it isn't factive. Assertions like (6) violate the constraint:

> 6. We don't tolerate the accusation that Russell is a sexual harasser. Those kinds of accusations are bad for morale.

I am going to ultimately argue for this analysis:

> **Open secrecy social norm (final)**
> If *p* is an open secret, then do not define the social interaction as one in which *p*.

As we'll see, complying with this norm generally requires interlocutors to maintain a privileged kind of conversational pretense. But in order to articulate the nature of this pretense, I will need to introduce and develop some concepts from Goffman (1959, 1967, 1974).

I am going to start by showing why three simpler, seemingly more obvious strategies won't work. First, you might think that acknowledging that *p* is a matter of avoiding certain words. On this view, non-acknowledgment is a kind of euphemistic practice:

> **Open secrecy social norm: euphemism**
> If *p* is an open secret and if *S* is some sentence that conventionally means *p*, don't utter *S*.

We can quickly rule out this hypothesis. It incorrectly predicts that Amelia's reply to Gerald below violates the open secrecy norm. But Amelia is impeccably upholding the norms:

> **Unfair will 3**
>
> Gerald: The will is unfair.
> Amelia: Saying that the will is unfair is a vicious smear!

Instead, we might try restricting the open secrecy norm so that it governs acts of saying or asserting, not acts of uttering:

> **Open secrecy social norm: saying**
> If *p* is an open secret, then don't say that *p*.

> **Open secrecy social norm: assertion**
> If *p* is an open secret, then don't assert that *p*.

But now the open secrecy norm is too weak. It incorrectly permits speakers to make assertions like (12, 13):

> 12. Because Russell is a sexual predator, we can't hire him.
> 13. Dad is no longer an alcoholic.

In response to these difficulties, you might think that we need a principle that governs more than group members' linguistic behavior:

Open secrecy social norm: presupposition
If *p* is an open secret, then don't presuppose that *p*.

Open secrecy social norm: acting as if
If *p* is an open secret, then don't act as if *p*.

These norms seem like an improvement. They prohibit assertions of (12) and (13). But they are, once again, too strong. They predict that when Ethan watches women at parties and subtly guides them away from Russell, he is violating the open secrecy norms. But while Ethan is not promoting Russell's interests or straightforwardly upholding patriarchy, he isn't exactly breaking the silence either. At a more basic level, this analysis makes paradoxical demands on group members. If Amelia is conforming to open secrecy norms, then Amelia is acting as if or presupposing that the will is unfair. But she is doing so *because* she is treating the will as an open secret. Since open secrecy is factive, it follows that when Amelia presupposes or acts as if *p* is an open secret, she presupposes or act as if *p*. These norms thus require her to both act as if and not act as if *p*.

Restricting these norms so that they only govern what we presuppose for the sake of *joint* interaction doesn't fix the problem. Suppose that at a party, Ethan and Miranda can jointly perceive that Russell is aggressively bothering a woman. Without acknowledging what they are doing, they gradually move themselves towards Russell, hoping that their presence will interrupt him. Our modified norms predict that Ethan and Miranda are violating the open secrecy norms. And once again, that seems incorrect. The norms similarly predict that if two interrogators in **Torture** work together to obstruct an anti-torture NGO's investigation or if Mark and Eloise in **Alcoholism** work together to avoid discussing any topic related to Jerry's addiction, they are violating the open secrecy norms.

The next move is to combine the best of the latter two strategies, so that the norm doesn't narrowly rule out certain speech acts but nonetheless targets only linguistic behavior:

Open secrecy social norm: presupposition
If *p* is an open secret, then don't presuppose that *p* for the sake of the conversation.

Open secrecy social norm: acting as if
If *p* is an open secret, then don't act as if *p* for the sake of the conversation.

This is on the right track. But in order to make this idea work, we'll need a heavy-duty notion of conversation. This is because we need to ensure that just because two people are relying on their shared knowledge of *p* in order to *guide* their conversational behavior—as when Mark and Eloise work together to avoid certain topics—it doesn't follow that they are acting as if *p for* the conversation. To accomplish this, I'll need three sets of tools: first, the apparatus of conversational *presupposition* and the *common ground,* which is a kind of *lingua franca* in philosophy of language (Stalnaker 1984, 1999, 2002, 2014); second, the concept of activity *layering*, broadly rooted in Goffman's ontology of activity (Goffman 1974, Clark 1994); and finally, the concept of a social situation definition, again broadly rooted in Goffman (Goffman 1959, 1967).

The basics of the former are well-known, so I'll run through them only briefly. First, Stalnaker hypothesizes that when we converse, we do so against a shared and evolving background of information called the *common ground* (Stalnaker 2002, 2014). The common ground roughly corresponds to what we take for granted in the conversation. This body of information plays a central role in both explaining what we're doing in a conversation and justifying why this behavior is rational.

In the simplest and most straightforward kinds of conversations, what is common ground is just what we jointly know (Stalnaker 2014: 24-25). But often we aren't in a simple conversation. Our conversations are chockful of pretense, assumption, supposition, and other non-truth-aiming attitudes. In its most recent

instantiation, therefore, Stalnaker identifies the common ground not with what interlocutors jointly know or believe but what they jointly accept.

> **Common ground**
> *P* is common ground in a conversation *c* iff all interlocutors in *c*
> accept (for the sake of *c*) that *p*,
> accept (for the sake of *c*) that they accept (for *c*) that *p*,
> accept (for the sake of *c*) that they accept (for *c*) that they accept (for *c*) that *p*,
> *et cetera.*
>
> (Stalnaker 2014: 4, 25)[16]

To accept that *p* is roughly to treat *p* as true for some purpose or other (Stalnaker 1984: 84, Stalnaker 2002: 716). One way to accept that *p* for the sake of the conversation is to believe that *p*. But just because you *don't* believe that *p*, it doesn't follow that you can't or don't accept that *p* for the sake of the conversation. For example, you might assume that *p* in order to explore what follows (Stalnaker 2002: 717, 2014: 25). And just because you believe that *p*, it doesn't follow that you accept that *p* for the sake of conversation (Stalnaker 2014: 46). You might pretend that *not p* (Stalnaker 2002: 718) or that you don't know whether *p* or *not p*.

Stalnaker defines presupposition as a propositional attitude that interlocutors take towards the common ground. In a non-defective conversation, interlocutors' bodies of presupposition are identical with each other's and the common ground (Stalnaker 2014: 46-47):

> **Presupposition**
> An interlocutor in *c* presupposes that *p* for the purposes of *c* iff she accepts that *p* is common ground in *c*.
>
> (Stalnaker 2014: 4, 25)

The concept of activity layering meshes well with this framework. Goffman argues that whenever we act, our action has one or more *frames,* which we can use in order to taxonomize action (Goffman 1974). For our purposes, we only need to worry about *layered joint actions*. In these activities, interlocutors are acting in ways that presuppose multiple, hierarchically ordered layers, which are a special type of frame (Clark 1996: 354-358). These interlocutors use the "higher" layer in order to coordinate activity at the "lower" layer. Information at the lower layer is, in some sense, replicated at the higher layer but not vice versa.

We'll say that each layer presupposes its own *common ground*.[17] To do this, we'll need to adjust our terminology. Stalnaker defines the common ground in terms of what we accept *for the conversation*. But plausibly, if I accept that *p* for the sake of some layer that partially constitutes the conversation, I accept that *p* for the conversation. This would entail that interlocutors routinely presuppose both *p* and *not p*. So we'll need to index each attitude of acceptance to the relevant layer:

---

[16] Both this definition and the subsequent definition for presupposition differ from the earlier ones in Stalnaker (2002: 716-717). For our purposes, the difference won't matter.

[17] Why am I saying that these conversations involve multiple common grounds and not just multiple sub-contexts? This is because any sub-context that characterizes an activity layer must have the iterative attitudinal structure constitutive of a common ground. For example, the sub-context that corresponds to the William and Kate pretense (in the main text below) must entail that William (Sean) and Kate (Polina) accept and accept that they accept and accept that they accept…*p & q*. It is this iterative structure that guarantees that William (Sean) and Kate (Polina) are having a normal, rational conversation within the pretense. But a sub-context needn't have this iterative attitudinal structure (Stalnaker 2014: 147). For example, you and I can construct and access a sub-context in which neither of us exist, as when I assert, "If humanity hadn't existed, then the Great Pyramids wouldn't have existed."

Common ground
P is common ground in a layer *l* of a conversation *c* iff all interlocutors in the conversation *c*
accept (for the sake of *l*) that *p*,
accept (for the sake of *l*) that they accept (for *l*) that *p*,
accept (for the sake of *l*) that they accept (for *l*) that they accept (for *l*) that *p*,
et cetera.

Presupposition
An interlocutor in a conversation *c* presupposes that *p* for the purposes of a conversational layer *l* of *c* iff she accepts that *p* is common ground in *l*.

In order to coordinate, interlocutors simultaneously access information within both common grounds. To see how this works, consider two actors, Sean and Polina, improvising a scene in which they are Prince William and Princess Kate. Relative to the *lowest* common ground, the conversation occurs between William and Kate in Kensington Palace. That is, *p* and *q* are common ground:

p: We are Prince William and Princess Kate.
q: We are in a living room in Kensington Palace.

As evidence, suppose that Sean asserts (14):

14. Our eldest son attends Lambrook.

In context, he asserts (or purports to assert) something true. This is because "my eldest son" picks out Prince George, who attends Lambrook. Common ground information standardly plays a role in resolving context-sensitive information like this.

But in order to guide the entirety of their action, Sean and Polina also need to access information like:

r: This conversation is occurring between Sean, who is playing Prince William, and Polina, who is playing Princess Kate.
s: This conversation is occurring on a stage.

For example, Sean and Polina need to coordinate so that they are always facing towards the audience, so that they are speaking loudly enough for the audience to hear, and so that they don't introduce any elements into the scene that they would struggle to convincingly portray. $R \& s$ are logically incompatible with $p \& q$. So Sean and Polina must be presupposing a second common ground that entails $r \& s$ but not $p \& q$.

The higher common ground contains the lower common ground in the following sense:

Common ground layering
Interlocutors in a conversation *c* layer common ground cg_2 over cg_1 just in case:
(1) interlocutors in a conversation *c* simultaneously rely on cg_1 and cg_2, and
(2) if cg_1 entails *p*, then the interlocutors aim to ensure that cg_2 entails [that cg_1 entails *p*], and
(3) cg_1 does not entail the existence of cg_2.

The layering principle captures the fact that Sean and Polina aren't just pretending *p*. They also simultaneously jointly *believe* that they are pretending that *p*, and they act in order to ensure that they jointly believe this. These joint beliefs enable them to coordinate their pretense. Moreover, layering is a hierarchical relation. The lower common ground does not have access to the existence of the higher common ground.

This accounts for the fact that *in the pretense* William and Kate are not pretending that they are in Kensington Palace. They really are in Kensington Palace.

The advantage of the common ground layering approach is that it is infinitely recursive.[18] It easily handles pretenses embedded within pretenses. Consider a case in which Sean and Polina pretend that they are William and Kate, who then pretend that they are Harry and Meghan in order to mock them. Then Sean and Polina are simultaneously accepting three jointly incompatible propositions:

p:	We are Sean and Polina.	*accepted for l_3 in virtue of believing* p
q:	We are William and Kate.	*accepted for l_2 in virtue of pretending* q *in pretense$_1$*
r:	We are Harry and Meghan.	*accepted for l_1 in virtue of pretending* r *in pretense$_2$*

Since Sean and Polina are accepting both *q* and *r* in virtue of pretending that *q* and *r*, we cannot define each common ground by appealing to the particular "flavor" or "tone" of acceptance that constitutes it.[19] We need to index the acceptances to a layer.

It's tempting to now analyze the open secrecy norm like this:

> **Open secrecy social norm: acceptance 1**
> If *p* is an open secret, then don't accept that *p* for the purposes of the lowest layer *l* of an interaction.

And, indeed, there is much to recommend this analysis. Suppose that Eloise and Mark carefully steer a conversation about Jerry away from any discussion of his recent black outs. Eloise and Mark are relying on and so presupposing that Jerry's alcoholism is an open secret. (This explains both what they're doing and how they're doing it.) But Eloise and Mark aren't thereby breaking the open secret. Structurally, this conversation is no different from the one between Sean and Polina on stage. There Sean and Polina relied on their knowledge that they were merely pretending to be William and Kate in order to steer their conversation away from any topics that they, Sean and Polina, don't know much about but that William and Kate would. Just like the information that Sean and Polina are on stage is located at a higher layer in the interaction, so too is the information that Jerry's alcoholism is an open secret. At the lowest level, it isn't common ground that Jerry is an alcoholic.

But there are still serious problems with this analysis. Because I've analyzed the norm in terms of the *lowest* layer of the interaction, the norm makes false predictions about more obvious forms of conversational pretense. For example, suppose that two interrogators in **Torture** pretend that they are United Nations officials in order to mock them:

---

[18] It also explains puzzling data that philosophers have debated elsewhere. Suppose that I tell a journalist, "I am saying this only to you. And I am going to say it only once. If you repeat it or anything that presupposes it, I will deny it. The Attorney General was behind the cover-up" (Fallis 2013: 350). Name the proposition that I assert *p*. Am I proposing to make *p* common ground or not? Some philosophers say obviously not (Fallis 2013; Keiser 2020). I'm not prepared to treat it as background for the conversation. Other philosophers say yes, because my interlocutor can felicitously presuppose that *p* (Stokke 2018). For example, she can ask, "How did the Attorney General arrange the cover-up?" On my analysis, both camps are right. I am proposing that my interlocutor and I adopt a layered common ground, the primary layer of which does not entail either *p* or that I communicated *p* but the secondary layer of which does. The journalist is acting uncooperatively in one sense, since she is not adopting my pretense. However, I have no problem *interpreting* her since we jointly have access to the relevant presupposition at a higher layer of the common ground.

[19] Thus we cannot individuate these common grounds according to what Yalcin calls "conversational tone" (Yalcin 2007: 1008).

Torture 2

George: *(in a ridiculous voice)* You're torturing the precious detainees! Oh no! Someone call the Security Council!
Teddy: *(in a ridiculous voice)* Yes, you will be receiving a letter from us telling you how very angry we are at you!

Here the lowest common ground is the common ground that entails George and Teddy are UN officials and that the government is torturing detainees. So according to the above norm, George and Teddy are violating open secrecy. But precisely by mocking the very possibility that they are torturing detainees, George and Teddy are conforming to the open secrecy norms. For similar reasons, the analysis incorrectly predicts that in **Torture 3**, George and Teddy *are* conforming. Given the exaggerated and phony nature of their pretense, it should not:

Torture 3

George: *(in a ridiculous voice)* Ohh yeah, we're *tootallllyyy* not torturing anybody. Wink wink.
Teddy: *(in a ridiculous voice)* Yup, no torture here. Nosiree. The UN would be proud of us.

Finally, even when we exclude the kinds of obvious pretenses in **Torture 2** and **3**, the norm still fails to prohibit even basic assertions about the open secret. This is because there is no in principle reason why our speech acts must target the lowest layer of a conversation. For example, if Polina can find a discrete way to do so, she might say this to Sean:

15. a. We need to prepare for the photographers!
b. *(a whisper)* Speak up. The audience isn't hearing you.

Or while pretending that they're UN officials, Teddy might say this:

16. a. Yes, you will be receiving a letter from us telling you how very angry we are at you!
b. Wait, hang on, George, is that what they do? Do they send letters? Or do they just pass resolutions?

In both cases, the speaker clearly intends her speech act in (b) to target the higher common ground, which entails, respectively, that the interlocutors are on stage and that they are not UN officials. Since both speakers can then seamlessly move back into the pretenses, it's *ad hoc* to claim that the common ground that corresponds to the pretense suddenly stops existing. But if the pretense common ground persists, this incorrectly entails that George in (17) is conforming to the open secrecy norm.

17. a. We're only engaged in intensive interrogation. It's illegal to torture in this country.
b. Wait, hang on, Teddy, we're just saying that, right? We're really torturing and just pretending that we're not, right?

This is because he continues to accept, for the lowest conversational layer in the conversation, that they are not torturing. His question in (b) reveals that he is accepting the open secret only at the higher layer. But (17b) is exactly the kind of speech act that open secrecy norms prohibit.

What we need is a way to determine which pretenses do and don't count, for the sake of open secrecy norms, and an explanation for why speakers cannot ever permissibly "break out" of those pretenses. To do this, I will once again borrow from Goffman's ontology of activity (Goffman 1959: 3-7 and *passim*, 1967: 5-7 and *passim*). What open secrecy norms prohibit is *defining the social situation* such that it entails *p*:

Open secrecy social norm (final)
If *p* is an open secret, then do not define the social interaction as one in which *p*.

As I interpret Goffman, a social interaction[20] is not just a situation involving two or more people or a situation in which social norms govern individuals. Rather, like pretenses, demonstrations, performances, and competitions, a social interaction is a *type* of joint activity, with its own constitutive rules and norms. But it is usually the most basic type of interaction that we have. Thus, Ethan demonstrates his knitting skills to Miranda and Sean and Polina improvise their pretense only within and as part of a broader social interaction. Social interactions are special not just because of the constraints they impose on our activity but because of their normative role. Our *face,* which is something like our social self or social standing, is primarily sensitive not to what we know or jointly know but to what we mutually take for granted in the social interaction (Goffman 1967: 5-8 and *passim*).

In the multiple common ground framework that I have been developing, defining the social situation is just a special form of accepting. Again, this form of acceptance is privileged because of the privileged status of the activity to which it corresponds. And so we can speak of a privileged social interaction common ground.[21]

Defining the social interaction
A speaker defines the social interaction as *p* if she accepts that *p* for the sake of the layer *l* of the conversation *c* that corresponds to the social interaction.

When the social interaction common ground entails *p*, then the social interaction *is* defined as one in *p*:

The definition of the social interaction
A social interaction is defined as one in which *p* iff it's common ground in the layer *l* that corresponds to the social interaction that *p*.

(This means that even if you define the social interaction as one in which you're the Prince of Morocco, the social situation is not yet defined this way. We must both opt into the pretense.)

If the social interaction contains phony elements, we should not, at any layer of the conversation, draw excessive attention to that fact. In fact, for Goffman, this is something like a constitutive norm:

---

[20] In Goffman (1961), he replaces the term "interaction" with "encounter." I retain the earlier but better-known term from Goffman (1959).

[21] Two progenitors to my proposal are those independently developed in Camp (2018) and Stokke (2018). Camp distinguishes between a conversational record and common ground, Stokke (2018) between an official and an unofficial common ground. My framework has advantages over both of these approaches.

First, Camp (2018)'s characterization of the conversational record requires us to accept a theory of conversational commitment, which many intentionalists will find dubious. More worryingly, Camp's view entails that when I communicate only within a common ground that's layered over the conversational record (i.e., the social interaction common ground), I am not committed to what I communicate. But intuitively, that doesn't seem right. On this point, see fn 23.

In Stokke's framework, unofficial common grounds are what we rely upon and update when we tell elaborate jokes or act on stage (Stokke 2018: 59-60). The official common ground is what we rely upon and update when we make genuine assertions and other serious speech acts. The problem, as Michaelson (2018) points out, is that we sometimes seem to rely on *two* official common grounds. On my analysis, one of those common grounds is the social interaction common ground, and the other is layered over it. On this point, see fn 18.

In distinguishing between the official and unofficial common ground, Stokke intends to offer an analysis of lying. On his analysis, to lie about *p* is, *inter alia*, to propose to update the official common ground with *p*. Using the framework I've sketched, Stokke could instead hold that to lie is, *inter alia*, to propose to update the social interaction common ground with *p*. This adjustment may deflate some of the objections that Michaelson (2018), Keiser (2018), and others raise. Unfortunately, it would take us too far afield to discuss the details here.

Verisimilitude principle
In any social interaction, if you accept that *p* for the sake of the social interaction layer *l* but do not believe that *p*, do not draw excessive attention to this fact.[22]

Thus, assertions like (18-20) are not just offensive. The speakers are violating a constitutive rule of interaction itself:

18. That's fascinating. What else did you do on your vacation? I mean, look, I'm not actually interested in it, but I've got to say something to you to just keep passing the time.

19. You aren't actually busy after this, but it's okay. I don't mind if you leave now.

20. You're trying to project the image of a really erudite philosopher. To be clear, I'm going along with it just to politely humor you and also because as my advisor, I need to flatter you.

Central to Goffman's thesis is the observation that we often conform to this principle not by deceiving each other about what we really believe but by mutually collaborating on a pretense. It is often common knowledge that we're feigning interest in a topic, concocting a false reason to leave a gathering, or merely going along with another's self-indulgent persona. In such cases, the speaker of (18-20) violates the verisimilitude principle not because of what she reveals but because of what she acknowledges.

Just as Sean and Polina rely on the information that they're on stage in order to coordinate their pretense, interlocutors often represent this phoniness at a higher layer of the conversation (Goffman 1967: 30 and *passim*). If we both know that you're only pretending that you're authentically happy that I received a job offer, we can rely on this fact in order to work together to steer our conversation away from that topic. Or suppose that my interlocutor makes what I believe is an outrageous request of me. If I reply, "Sure," not quite sarcastically but still letting my hesitation shine through my tone of voice, I simultaneously communicate, for the purposes of the social interaction, that I am happy to perform the request and, for the purposes of the higher layer that corresponds to what we jointly know, that I am unhappy to perform the request.[23] If you are cooperative, you will then concoct some reason to withdraw your request. I enabled you to save face, because I didn't enter my displeasure into the social interaction layer.[24]

---

[22] Other Goffman-inspired versions of this principle include:

1. Do not appear to believe that *p*. (Brown and Levinson 1978: 212).
2. Act so that a virtual audience would believe that you believe that *p*. (Lee and Pinker 2010: 7896, Camp 2018: 50)

The problem with (1) is that it's unclear how to interpret "appear." The problem with (2) is that it's too strong. In the text below, for example, a virtual audience would have access to my tone of voice and so would know that I don't want to fulfill the request. But that case is canonically one in which I'm conforming to the verisimilitude principle (Goffman 1967: 30).

[23] On Camp's analysis (2018), I don't commit to the proposition that I don't want to perform the request. By this, she means that I'm not liable to defend the truth of this proposition in other contexts (Camp 2018: 59). But that's the wrong way around. If someone were watching my interaction and later asked me, "Did you really not to accept that request?" I would need to acknowledge and defend that I had committed to that proposition in the previous context. I only seem to escape liability in the current context because you can't directly address the matter without violating the verisimilitude principle.

[24] An anonymous referee points out that we often use a similar strategy to avoid officially acknowledging changes to our social relationships. Consider two friends, Marie and Aaron, who jointly know that their friendship has run its course. When Marie runs into Aaron, she says, "Let's get lunch sometime," and he replies, "Let's." Marie and Aaron only make these requests because they jointly know that the other won't follow up on them. They've defined the social interaction as one in which they're friends. But to coordinate, they are relying on their shared knowledge that they are not. Whether Marie and Aaron are maintaining an open secret depends upon whether they are enforcing this informal practice as a norm against each other and whether that norm has the right structure.

So whatever it is to draw "undue attention" to the gap between what I accept and what I believe, it doesn't prohibit me from accepting *p* for the purposes of the social interaction and simultaneously insinuating *not p* at a higher common ground. Luckily, for the sake of explaining the open secrecy data, we don't need a full analysis of the verisimilitude principle. We only need to accept that whatever its analysis, it entails this principle:

Social Interaction is Mandatory
If I utter some *U* in order to make one or more speech acts *S*…*S*,* then
(a) I necessarily aim to update either
(i) the social interaction layer *l* or
(ii) some common ground *l** such that *l* is layered over *l**
(b) with at least one of *S*…*S.**

This explains what goes wrong with (17). When George asks, "We're really torturing and just pretending that we're not, right?" he is only making one speech act with his utterance. In this interaction, George is not asking this question as part of a pretense within the social interaction. So there is no common ground *lower* than the social interaction layer that George could target. That leaves the social interaction layer and the layer above the social interaction, relative to which they are torturing people. According to **Mandatory**, George's speech act must target the former. Therefore, George is necessarily accepting for the sake of the social interaction that they are torturing detainees. And therefore, he violates the open secrecy norm.

Mandatory also helps us to see why, in both **Torture 2** and **Torture 3**, the lowest common ground does not correspond to the social interaction. To do this, we need to check whether a speaker who asserts (21) and (22), respectively, would violate the rules of the social interaction *qua* social interaction:

21. We're not really in the UN.
22. We're torturing people.

This is trickier than it sounds for two reasons. First, if two interlocutors are engaged in a pretense, it's often uncooperative to suddenly not engage. In George were to assert (21) in **Torture 2** out-of-the-blue, his assertion would be socially inappropriate in the sense that he'd be a killjoy. Second, if Teddy were to say (22) in **Torture 3**, he'd violate the open secrecy norm. So he'd assert something socially unacceptable, but it wouldn't necessarily be socially unacceptable because of the nature of social interaction itself. So, let's imagine that in **Torture 2**, George exclaims, as their pretense winds down, "Thank God we're not really in the UN!" And let's imagine that in **Torture 3**, George and Teddy are close friends who hate the open secrecy in their workplace and in private regularly flout the norms.

In both cases, the assertions of (21) and (22) are not socially inappropriate and do not violate the verisimilitude rule. Moreover, the fact that in both **Torture 2** and **Torture 3**, George and Teddy interact in highly theatrical, exaggerated, even parodic ways is further evidence that they are speaking in a pretense *within* a social interaction. They are not directly updating the social interaction common ground. Such exaggerated action surely violates the "undue attention" constraint of the verisimilitude principle. Moreover, generally we use this kind of exaggeration precisely in order signal to each other that we are operating within a pretense *within* a social interaction and not the social interaction itself.

Finally, we can now explain why some forms of indirect speech, but not others, break the open secrecy norms. For example, normally if Ethan asserts (23) in order to implicate (3) to Miranda, then he is violating the open secrecy norms:

23. Russell is, ummmm, a bit of a climate problem.
3. Russell is a sexual harasser.

This is because the mere fact that Ethan implicates, rather than says, (23) doesn't entail that Ethan is *pretending*, at any layer, that he didn't implicate (3). For example, Miranda can non-problematically reply:

24. I appreciate that you're looking out for my safety.

Since Miranda's single speech act in (24) presupposes that Ethan implicated (3), Miranda is both accepting and presupposing that Ethan implicated (3) for the purposes of the social interaction. So Ethan is also accepting that he implicated (3). So Ethan and Miranda are both violating the open secrecy norm.

On the other hand, many cases of what Elisabeth Camp calls "deep insinuation" *do* conform to open secrecy norms (Camp 2018: 57). Here's a snippet of Jane and Wilcock's conversation. At this point, it's become clear that Wilcock knows that Jane has insinuated that her boss has hired the "girl" as a sexual *quid pro quo*. Jane has been emphasizing the girl's inexperience in the industry:

HR 2

Jane:	I was just—I was worried for this girl.
Wilcock:	She's a woman. She's a grown woman.
Jane:	Sorry, yes.
Wilcock:	(a) You think a grown woman can't make her own choices?
Jane:	I never said that.
Wilcock:	Because she's a waitress?

(*The Assistant* 55:42)

Wilcock is conforming to the open secrecy norms because he is continuously engaged in multiple simultaneous speech acts. For example, by his utterance in (a), Wilcock is asking *both* (25) and (26):

25. You think a grown woman can't make her own choice to take a job for which she lacks experience?
26. You think a grown woman can't make her own choice to accept the boss' sexual *quid pro quo*?

Since Wilcock's question in (25) is aimed at the social interaction common ground, he is free to ask (26) at the higher common ground. While he presupposes the open secret in asking (26), he is only accepting the open secret for the sake of this higher common ground. At the common ground that characterizes the social interaction, he is conforming to the open secrecy norm.

It is not always clear which speech act a speaker is issuing or what information she is presupposing, and so it's not always clear whether she is conforming to the open secrecy norm. But this is just as it should be: the fuzziness in our theoretical judgments corresponds to the fuzziness in our intuitive judgments. For example, in **HR1** when Wilcock says, "I don't think you have anything to worry about…You're not his type," he is clearly implicating that the boss, though a sexually harasser, isn't likely to sexually harass Jane. The question is whether his assertions, which necessarily target the social interaction common ground, break the open secret norm. To answer that question, we need to know precisely what he is presupposing for the sake of the social interaction. Is it intelligible that he's presupposing that the boss isn't a sexual harasser but that Jane has nothing to worry about because *even if he were* Jane isn't his type? It's unclear, which is why it's unclear whether he's violating the norm.

Finally, it is important to once again emphasize the difference between the run-of-the-mill non-acknowledgment that the verisimilitude principle requires and the extreme forms of non-acknowledgment that open secrecy demands. While social interaction constitutively imposes certain non-acknowledgment requirements on us, they only require us not to draw attention to the pretense-laden nature of *this very*

interaction. It does not prohibit speakers from acknowledging the pretense-laden nature of other conversations. So long as the assertions don't draw undue attention to any *ongoing* pretense within the conversation in which they are uttered, it's acceptable for a speaker to assert (27, 28):

> 27. I know you just heard me tell Daniel that I loved his ribs, but I was just being polite. They were too dry.
>
> 28. Sometimes people say that they have to go walk the dog, but really they're just tired and want to go home.[25]

But when it comes to open secrecy, there is *no* context relative to which a group member can find reprieve, either to state what was really happening in a prior conversation or to formulate the general open secrecy rule that governed it.

V. Strategically talking about open secrets

Open secrecy norms are often moral disasters. As the Weinstein case exemplifies, they can serve as shields for powerful people guilty of serious, even criminal wrongdoing. The norms can compound the harm that befalls their victims. Victims and their allies find they don't *just* have to contend with the perpetrator's financial resources, political might, and interpersonal capital. They must go up against an entire social arrangement. Group members, even those who dislike the perpetrator or have no fear of him, may deem a silence-breaker to be crass, uncooperative, unprofessional, naïve, or gossipy. They may find themselves distrustful of the victim: if she breaks *this* social norm, which other ones might she break? Or they may simply do nothing, depriving the victim of even the minimal solace, affirmation, or social recognition that she deserves.

Open secrecy norms are also, unsurprisingly, often disasters for those who want to *change* their dysfunctional groups. Someone who thinks, "If *only* I had the right kind of evidence to convince the CEO that Russell is a sexual predator! Then she'd fire Russell immediately!" does not understand the fundamentally social nature of what she faces. But that's a feature, not a bug. A newcomer, appalled to learn about what's happening in her organization or friend group, may waste time trying to "inform" the "right people" about it. By the time she comes to the sickening realize that everybody already knew the whole time, the newcomer has become burned out or ostracized.

---

[25] These examples show that at least some everyday politeness norms are not open secrets. But it's certainly possible that some everyday politeness norms are. An anonymous referee raises this intriguing possibility: could a politeness norm ever become an open secret *because* of the verisimilitude principle? In principle, yes. Suppose that whenever we talk to each other, we are always conforming to some politeness norm and that this norm always requires us to maintain some pretense. Then there will never be a context in which we can both conform to the verisimilitude principle *and* state the content of the norm. Articulating the norm would draw undue attention to a pretense operative in that very conversation.

But it's unclear to me whether there are any such politeness norms. Whatever the norm is, surely at least sometimes people won't conform to it. Suppose that I stare at my long-time enemy in stony silence while he drones on about how everyone loves his great ideas. I refuse to nod or give him any cursory uptake. In such a context, if I were to say, "You think everyone agrees with you because you're an idiot. You don't realize that sometimes, people just nod at other people out of politeness," I would not violate the verisimilitude principle.

It remains the case that the verisimilitude principle may make it quite difficult to talk about many of the deeper and more widespread politeness norms.

For all that, it would be a mistake to think that open secrecy norms always protect the subject of the open secret. I want to close by considering how open secrecy norms sometimes protect those who *talk* about the open secrecy. Consider the following story:

Aemond's toast

Recall that in *House of the Dragon,* it's an open secret that the visual appearance of Rhaenyra's children, Jace, Luke, and Joffrey, is evidence that they are the children of Harwin Strong. Strong is not their legal father. The boys' cunning uncle, Aemond, delivers the following toast:

Aemond: To the health of my nephews. Jace. Luke. And Joffrey. Each of them handsome, wise….*strong*…
Alicent: Aemond!
Aemond: (a) Come! Let us drain our cups to these three *strong* boys.
Jace: I dare you to say that again.
Aemond: Why? 'Twas only a compliment. Do you not think yourself strong?

Jace and Luke attempt to attack Luke. Alicent, Aemond's mother, then takes Aemond aside and privately asks him:

Alicent: Why would you say such a thing before these people?
Aemond: I was merely expressing how proud I am of our family, Mother.

He then turns, speaking loudly to Jace and Luke.

(b) Though it seems my nephews aren't *quite as proud of theirs.*

("The Lord of the Tides" 1:01:00)

Much like Wilcock, Aemond conforms to open secrecy only because he simultaneously issues multiple speech acts. At (a), Aemond simultaneously asserts (29, 30):

29. Let us drain our cups to these three strong boys.
30. Let us drain our cups to these three Strong (as in, the son of Harwin Strong) boys.

At (b), Aemond more or less asserts (31, 32):

31. Though it seems my nephews aren't quite as proud of me, their family member.
32. Though it seems my nephews aren't quite as proud of their father, Harwin.

But Aemond has laid a trap for his nephews. On the one hand, since it's common knowledge that the boys look like Strong, it's common knowledge that Aemond is insinuating (31, 32). At the banquet, it's as public as can be that Aemond has insulted them and so impugned their honor. On the other hand, in order to retaliate against Aemond, Jace and Luke need to break the open secrecy norm.[26] So they are stuck: they can either quietly accept that Aemond has called their parentage into question (and so seem weak in the face of challenge), find a creative way to respond to the insult through insinuation (which, again, may seem weak), or openly respond and so break the open secrecy. As hot heated teenage boys, they choose to the latter.

[26] Berstler (MS) argues against Dinges and Zakkou (2023) that in virtue of this fact, Aemond has non-epistemic or implausible deniability for insinuating (30, 32).

To see how Aemond has created this trap, remember that Aemond's insinuation targets a common ground higher than and layered over the social interaction common ground. Relative to this common ground and only this common ground, the boys look like their father is Harwin Strong. When Jace and Luke then openly attack Aemond, their interaction is only interpretable under the assumption that they are presupposing that Aemond insinuated that they are illegitimate. And they can only intelligibly presuppose that Aemond insulted them if they also presuppose that they look like Harwin Strong. So Jace and Luke, in attacking Aemond, accept the open secret for the sake of the social interaction. So they break the open secrecy norms.

Thus, the extreme, spiraling silence that open secrecy imposes upon groups doesn't uniformly protect the subject of an open secret. It also leaves her systematically vulnerable to hints and insinuations that she cannot openly address. And the more communicatively creative and subtle the speaker, the more she can talk about the open secret. In fact, we can now see that even the metaphor of "silence" is no longer quite apt. People can communicate loudly and constantly about an open secret—so long as their speech acts are located on the periphery of the social interaction.

VI. Bibliography

Alford, C. Fred. *Whistleblowers: Broken Lives and Organizational Power.* Ithaca: Cornell University Press, 2002.

Arendt, Hannah. *Eichmann in Jerusalem: a Report on the Banality of Evil.* New York: Penguin, 1994 [1965].

Axios staff. "Harvey Weinstein was Hollywood's Biggest Open Secret." *Axios.* 6 October 2017. Accessed 27 November 2023. https://www.axios.com/2017/12/15/harvey-weinstein-was-hollywoods-biggest-open-secret-1513306011.

Berstler, Sam. "Non-Epistemic Deniability." MS 2024. Accessed 10 September 2024. On file at https://samberstler.com/research/.

--------. "What's the Good of Language? On the Moral Distinction between Lying and Misleading." *Ethics* 130 (2019): 5-31.

Bicchieri, Cristina. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms.* New York: Oxford University Press, 2016.

Brennan, Geoffrey and Philip Pettit. *The Economy of Esteem: an Essay on Civil and Political Society.* New York: Oxford University Press, 2005.

Brown, Penelope and Stephen C. Levinson. *Politeness: Some Universals in Language Usage.* Cambridge, UK: Cambridge University Press, 1978.

Camp, Elisabeth. "Insinuation, Common Ground, and the Conversational Record." In *New Work on Speech Acts.* Edited by Daniel Fogal, Daniel W. Harris, and Matt Moss. 44-46. New York: Oxford University Press, 2018.

Clark, Herbert. *Using Language.* New York: Cambridge University Press, 1994.

Cohen, Stanley. *States of Denial: Knowing about Atrocities and Suffering.* Malden, MA: Polity Press, 2001.

Davis, Alex. "Testimony, Recovery and Plausible Deniability: a Response to Peet." *Episteme* 16 (2019): 18-38.

Dinges, Alexander and Julia Zakkou. "On Deniability." *Mind* 132 (2023): 372-401.

Dunbar, Robin. "Gossiping in Revolutionary Perspective." *Review of General Psychology* 8 (2004): 100-110.

--------. *Grooming, Gossip, and the Evolution of Language.* London: Faber and Faber, 1996.

Elga, Adam and Agustín Rayo. "Fragmentation and Information Access." In *The Fragmented Mind.* Edited by Cristina Borgoni, Dirk Kindermann, and Andrea Onofri. New York: Oxford University Press, forthcoming.

Elster, Jon. *Alchemies of the Mind: Rationality and the Emotions.* Cambridge, UK: Cambridge University Press, 1999.

Fallis, Don. "Davidson was Almost Right about Lying." *Australasian Journal of Philosophy* 91 (2013): 337-353.

Fogal, Daniel, Daniel W. Harris, and Matt Moss. "Speech Acts: the Contemporary Theoretical Landscape." In *New Work on Speech Acts.* Edited by Daniel Fogal, Daniel W. Harris, and Matt Moss. 1-39. New York: Oxford University Press, 2018.

Fricker, Elizabeth. "Stating and Insinuating." *Aristotelian Society Supplementary Volume* 86 (2012): 61-94.

Gendler, Tamar Szabó. "Alief and Belief." *Journal of Philosophy* 105 (2008): 634-663.

Goffman, Erving. *Frame Analysis: an Essay on the Organization of Experience.* Cambridge, UK: Harvard University Press, 1974.

--------. *Encounters: Two Studies in the Sociology of Interaction.* Indianapolis: Bobbs-Merill, 1961.

--------. *Interaction Ritual: Essays on Face-to-Face Behavior.* New York: Random House, 1967.

--------. *The Presentation of the Self in Everyday Life.* New York: Anchor Books, 1959.

Greco, Daniel. "Iteration and Fragmentation." *Philosophy and Phenomenological Research* 88 (2014): 656-673.

Grice, H.P. "Meaning." *Philosophical Review* 66 (1957): 377-388.

Gugliersi, Antonella. "Why Viserys Refuses to Believe the Truth about Rhaenyra's Kids." *Screenrant.* 29 September 2022. Accessed 27 November 2023. https://screenrant.com/why-viserys-doesnt-believe-truth-rhaenyra-harwin-children/.

Haslanger, Sally. "Systemic and Structural Injustice: Is There a Difference?" *Philosophy* 98 (2023): 1-27.

--------. *Ideology in Practice: What Does Ideology Do?* The Aquinas Lectures. Milwaukee, WI: Marquette University Press, 2021.

House of the Dragon. Created by G.R.R. Martin and Ryan Condal. HBO. 2022.
Katz, Daniel and Richard L. Schank. *Social Psychology*. New York: Wiley, 1938.
Keiser, Jessica. "The Limits of Acceptance." *Inquiry: an International Journal of Philosophy* (2020): 1-18.
Kiefer, Halle. "Screenwriter Scott Rosenberg Recalls His Time with Harvey Weinstein at Miramax: 'Everyone F*cking Knew.'" *Vulture. New York Magazine*. 16 October 2017. Accessed 27 November 2023. https://www.vulture.com/2017/10/writer-scott-rosenberg-everyone-knew-about-weinstein.html.
Kumar, Nish and Coco Khan. "The Russell Brand Scandal and Fixing the Police." *Pod Save the UK*. 21 September 2023. Accessed 27 November 2023. https://crooked.com/podcast/the-russell-brand-scandal-and-fixing-the-police/.
Lee, James J. and Steven Pinker. "Rationales for Indirect Speech: the Theory of the Strategic Speaker." *Psychological Review* 117 (2010) 785-807.
Lewis, David. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (1979): 339-359.
--------. *Convention: a Philosophical Study*. Cambridge, MA: Wiley-Blackwell, 1969.
McAtee, Riley. " 'House of the Dragon' Episode 8 Breakdown: Bastards and Broken Kings." *The Ringer*. 11 October 2022. Accessed 27 November 2023. https://www.theringer.com/house-of-the-dragon/2022/10/11/23397534/house-of-the-dragon-episode-8-breakdown.
Michaelson, Eliot. "The Lies We Tell Each Other." In *Lying: Language, Knowledge, Ethics, and Politics*. Edited by Eliot Michaelson and Andreas Stokke. 183-205. New York: Oxford University Press, 2018.
Miller, Dale T. and Deborah Prentice. "Collective Errors and Errors about the Collective." *Personality and Social Psychology Bulletin* 20 (1994): 541-550.
Nagel, Thomas. "Concealment and Exposure." *Philosophy and Public Affairs* 27 (1998): 3-30.
Noelle-Neumann, Elisabeth. *The Spiral of Silence: Public Opinion—Our Social Skin*. New York: Chicago University Press, 1993.

Nyong'o, Lupita. "Speaking Out about Harvey Weinstein." *New York Times*. 19 October 2017. Accessed 27 November 2023. https://www.nytimes.com/2017/10/19/opinion/lupita-nyongo-harvey-weinstein.html.
O'Gorman, Hubert J. "The Discovery of Pluralistic Ignorance: an Ironic Lesson." *Journal of the History of the Behavioral Sciences* 22 (1986): 333-347.
Peacocke, Christopher. "Joint Attention: Its Nature, Reflexivity, and Relation to Common Knowledge." In *Joint Attention: Communication and Other Minds*. Edited by Naomi Eilan, Cristoph Hoerl, Teresa McCormack, and Johannes Roessler. 298-324. New York: Oxford University Press, 2005.
Peet, Andrew. "Testimony, Pragmatics, and Plausible Deniability." *Episteme* 12 (2015): 29-51.
Pettit, Philip. "Social Norms and the Internal Point of View: an Elaboration of Hart's Genealogy of Law." *Oxford Journal of Legal Studies* 39 (2019): 229-258.
Pinker, Steven. "The Evolutionary Social Psychology of Off-Record Indirect Speech Acts." *Intercultural Pragmatics* 4 (2007): 437-461.
--------, Martin A. Nowak, and James J. Lee. "The Logic of Indirect Speech." *Proceedings of the National Academy of Sciences* 105 (2008): 833-838.
Ramadan, Lulu. " 'Open Secret' at Royal Palm High: At Least 15 Students Were Lured to Jeffrey Epstein's Palm Beach Mansion." *The Palm Beach Post*. 26 July 2019. Accessed 10 September 2024. https://www.palmbeachpost.com/story/news/crime/2019/07/26/jeffrey-epstein-case-open-secret-at-royal-palm-high/986590007/
Reilly, Katie. "How the #MeToo Movement Helped Make New Charges Against Jeffrey Epstein Possible." *Time*. 9 July 2019. Accessed 10 September 2024. https://time.com/5621958/jeffrey-epstein-charges-me-too-movement/
Ruiz, Michael. "Jeffrey Epstein Victims Sue FBI for Alleged Failure to Investigate 'Sex Trafficking Ring for the Elite.'" *Fox News*. 15 February 2024. Accessed 10 September 2024. https://www.foxnews.com/us/jeffrey-epstein-victims-sue-fbi-alleged-failure-investigate-sex-trafficking-ring-elite

Traister, Rebecca. "Why the Harvey Weinstein Sexual-Harassment Allegations Didn't Come Out Until Now." *The Cut. New York Magazine.* 5 October 2017. Accessed 27 November 2023. https://www.thecut.com/2017/10/why-the-weinstein-sexual-harassment-allegations-came-out-now.html.

Twohey, Megan, Jodi Kantor, Susan Dominus, Jim Rutenberg, and Steve Eder. "Weinstein's Complicity Machine." *New York Times.* 5 December 2017. Accessed 27 November 2023. https://www.nytimes.com/interactive/2017/12/05/us/harvey-weinstein-complicity.html?smid=tw-nytimes&smtyp=cur.

Stalnaker, Robert. *Context.* New York: Oxford University Press, 2014.

--------. "Common Ground." *Linguistics and Philosophy* 25 (2002): 701-721.

--------. *Content and Context: Essays on Intentionality in Speech and Thought.* New York: Oxford University Press, 1999.

--------. *Inquiry.* Cambridge, MA: MIT Press, 1984.

Schwitzgebel, Eric. "Acting Contrary to Our Beliefs or the Gulf Between Occurrent Judgment and Dispositional Belief." *Pacific Philosophical Quarterly* 91 (2010): 531-553.

Stokke, Andreas. *Lying and Insincerity.* New York: Oxford University Press, 2018.

The Assistant. Directed by Kitty Green. Symbolic Exchange Production Company. 2019.

Walton, Douglas. "Plausible Deniability and Evasion of Burden of Proof. *Argumentation* 10 (1996): 47-58.

Williamson, Timothy. *Knowledge and Its Limits.* New York: Oxford University Press, 2000.

Yahr, Emily. "Harvey Weinstein's Behavior was a Dark Inside Joke on Shows like "Entourage' and '30 Rock.'" *The Washington Post.* 10 October 2017. Accessed 27 November 2023. https://www.washingtonpost.com/news/arts-and-entertainment/wp/2017/10/10/harvey-weinsteins-behavior-was-a-dark-inside-joke-on-shows-like-entourage-and-30-rock/.

Yalcin, Seth. "Epistemic Modals." *Mind* 116 (2007): 983-1026.

Zacharek, Stephanie. "How Do You Solve a Problem Like Harvey Weinstein?" *Time Magazine.* 12 October 2017. Accessed 27 November 2023. https://time.com/4979250/how-do-you-solve-a-problem-like-harvey-weinstein/.

Zerubavel, Eviatar. *The Elephant in the Room: Silence and Denial in Everyday Life.* New York: Oxford University Press, 2006.