

**Metacontexts and Cross-Contextual Communication: Stabilizing the Content of Documents
Across Contexts**

Alex Davies

Forthcoming in *The Philosophical Quarterly*

ABSTRACT: Context-sensitive expressions appear ill-suited to the purpose of sharing content across contexts. Yet we regularly use them to that end (in regulations, textbooks, memos, guidelines, laws, minutes etc.). This paper describes the utility of the concept of a metacontext for understanding cross-contextual content-sharing with context-sensitive expressions. A metacontext is the context of a group of contexts: an infrastructure that can channel non-linguistic incentives on content ascription so as to homogenize the content ascribed to context-sensitive expressions in each context in the group. Documents composed of context-sensitive expressions can share content across contexts when supported by an appropriate metacontext. The bible has its church, the textbook its education system, the form its bureaucracy, and the manifesto its social movement. Some metacontexts support cross-contextual content-sharing. Some don't. A promising research programme (one with practical importance) would take metacontexts as its unit of analysis.

Keywords: context-sensitivity, content-sharing, incentives, sexual harassment, organizations

1. Cross-Contextual Context-Sensitive Communication

We expect great things of sentences. We produce them in one context, move them (often) to (many) others, and take for granted that their content has remained stable despite the move. For example, at our university, each department has its own review committee, which assesses the progress of the department's PhD students. Sentence (1) is taken from a document entitled "Study Regulations", which regulates, amongst other things, the actions of these committees:

- (1) The review committee assesses the learner's academic progress in research in credit points, considering the following criteria and the corresponding requirements established by the faculty council.

Sentence (1) is supposed to work its regulatory magic by conveying a common content to each person in each department who belongs to some such committee.

The practice of using sentences as vehicles to share content across contexts can seem quite unremarkable because the explanation of how it works can seem so boringly obvious. Users of the same language have learnt to associate the same (standing) meanings with the same words. Those meanings are building blocks of the contents (truth-conditions) of sentences composed using those words. Hence, provided they all speak a common language, different people, located in different contexts, will understand the same sentences to bear the same contents.

However, for many expressions, their standing meaning (something everyone knows about them in virtue of speaking the same language) underdetermines the content those expressions bear. Most obviously this is true of indexicals ("here", "today", "these") – with them, as Gareth Evans (1981, p. 293) once quipped, as one moves through contexts, one has to 'run to keep still' (i.e. *change* one's words in order to say the same thing twice: "Today is my birthday", "Yesterday was my birthday").

But there are also gradable adjectives (Kennedy & McNally, 2005)), modal verbs, modal adverbs, counterfactuals (Kratzer, 2012), quantifiers (Szabó & Stanley, 2000) and many other expressions whose standing meaning underdetermines content. I leave open precisely how wide the range of context-sensitive expressions is. But the following assumption certainly seems plausible: we routinely use context-sensitive expressions to (try to) share content across contexts ((Cappelen & Dever, 2016), (Eisenberg, 1984), (Lanius, 2019)). Look again at (1). As far as its standing meaning is concerned, (1) can be read distributively or collectively (Ritchie, 2017). Read distributively, the review committee assesses a learner’s academic progress (etc.) only if each member of the committee does so. Read collectively, (1) could be true even if labour is divided: e.g. each learner is assigned two members of the committee who review the learner’s progress, whilst the rest of the committee defers to those two members. The difference matters. If we read (1) distributively, then we have to get every member of the committee to read every document that every PhD student in the department submits for review. The likely result? Poorer quality of assessment: partly because each committee member will be overburdened, and partly because a by-stander effect is likely to arise between the committee members—both of which could be avoided with the division of labour permitted by a collective reading.

But if we do routinely use context-sensitive expressions to share content across contexts, then we have to acknowledge that the boringly obvious explanation of how we share content across contexts is significantly incomplete: if nothing takes up the slack left by standing meaning, it would be nothing more than a welcome accident when two users, each in a different context, ascribe the same content to the same context-sensitive expression, used in each context (Cappelen & Lepore, 2005). The incompleteness of the boringly obvious explanation raises what I call “the cross-context question”: how do we share content across contexts using context-sensitive language?

This paper makes a case for the utility of the concept of a *metaccontext* for addressing the cross-context question: viz. the context of a group of contexts within which language users ascribe content to the same context-sensitive expression. The core idea is going to be this. The standing meaning of context-sensitive expressions gives their users freedom to choose between a range of contents to ascribe to each such expression (section 2). But users’ content preferences are susceptible to influence from non-linguistic incentives (section 3). The context of a group of contexts for ascribing content (a.k.a. their *metaccontext*) can serve as an infrastructure that channels these incentives in a way that homogenizes which content is likely to be ascribed to the same context-sensitive expression at each context in the group (section 4).

Consider again (1). As a matter of fact, for years in our department it was stably ascribed a distributive content because of the incentives placed upon its different users. Unwilling to take on additional administrative burdens, academics in the department deferred to an overworked university administrator on how to interpret administrative documents. But this administrator didn’t know whether a division of labour would be considered rule-breaking within the university yet would have been blamed if they had permitted such rule-breaking to take place. Since they were overworked, they didn’t want to investigate the issue. Thus, they were incentivized to simply insist on the distributive reading. Since the academics deferred to this administrator, they too assigned it the distributive reading. Thus, for a time, the organizational structure of the department stably channelled incentives on occupants of these contexts in a way which ensured that, despite the limitations of its standing meaning, (1) was ascribed a distributive content across the contexts within it.

By adopting *metacontext* as a unit of analysis, we can take a cases-and-cases approach to the cross-context question. Instead of choosing between the view that we generally share content across contexts (e.g. (Cappelen & Lepore, 2005), (Szabo, 2006)) and the view that we generally don't share content across contexts (e.g. (Abreu Zavaleta, 2021), (Bezuidenhout, 2002), (Bowker, 2022), (Peet, 2016), (Sperber & Wilson, 1986)), we can say: it depends on the metacontext—some support content-sharing across the contexts within them, and some hinder it. The content ascribed to an expression across multiple contexts can be adjusted by adjusting the metacontext of those contexts—as we'll see in section 5 when we work through a case study of sexual harassment policy documents within a university. We'll explore how different structures of the university differently channel incentives on the ascription of content to "sexual harassment" resulting, in turn, in different contents being ascribed to the expression and different levels of cross-contextual stability in the content ascribed to it: cross-contextual content sharing requires good metacontext design. Section 6 explains how metacontexts can be bigger than individual organizations (lest section 5 leave the impression that only organizations can be metacontexts). Section 7 explains how this unit of analysis permits advances beyond the literature on contextualism and cross-contextual communication of the 2000s and 2010s.

I conclude this section by differentiating between the cross-context question and another related question not addressed in this paper. Picture a scenario where you and I share a context but wherein we ascribe contents (to context-sensitive expressions) that are more informative than our shared context provides us guidance in choosing between. For instance, we are sorting headphones on a production line into two piles: the red and the blue. All the blue headphones are exactly the same shade of blue. All the red headphones are exactly the same shade of red. There are various shades of colour that could be classified as "blue", or not, but which don't appear on the headphones we're sorting. Our context gives us no guidance on how objects with these shades of colour should be sorted ("blue" or not). Nevertheless, in this context, suppose we each ascribe a content (for simplicity's sake: a function mapping objects to truth-values) to the term "blue" which maps not only the headphones we're sorting but also these other objects. Since our context does not guide us in how to do this, it is highly likely that we will ascribe different contents to "blue" in this shared context: we won't share content using this word (see (Abreu Zavaleta, 2021), (Bowker, 2022), (Peet, 2016)). Given that the central features of this example generalize, the question arises: how do we share content within one and the same context? Call this "the one-context question".

I find this second line of thought unconvincing. It assumes that we aim to ascribe to context-sensitive expressions (in a context) contents that are more informative than the context can guide us in ascribing. But this assumption seems optional because typically, when contents are too informationally rich for the context to guide us in selecting between them, we don't need to ascribe contents that are so informationally rich. For example, if, in the production-line context, we ascribed a content to "blue" which sorts only the objects that are obviously relevant in our context (the headphones), then we'd still be able to sort the headphones, but the fact that the context doesn't give us guidance in how to categorize other objects wouldn't in any way imply that we are likely to ascribe different contents to this adjective in this context (for discussion see: (Davies, 2019a), (Davies, 2021)). But regardless of whether you find the one-context question compelling, this paper is about a problem that arises even on the assumption that individual contexts guide our selection of which content to ascribe to context-sensitive expressions used within them. Even assuming this (as we will), the cross-context question remains unanswered: how is content shared *across* contexts with context-sensitive expressions?

2. Linguistic Mistakes and Linguistic Freedom

In principle, you can use a word however you like: “bonjour” where in English we’d use “goodbye”, “Nobody” as the name of a man on the road, or “Who” as the name of the man on first base. But by and large, people don’t do that. It’s not that they can’t. They just don’t want to. People want to speak the language from which the words they utter are taken. But to do that, it’s not enough to utter words from the language: you need to do so without making linguistic mistakes. When Del Boy uses “bonjour” in place of “goodbye”, intending to speak French, although there’s some sense in which he is speaking French (he’s uttering a French word), there’s also an important sense in which he isn’t (he’s not using the word the way you need to in order to be speaking French with it).

I’ll follow Reiland’s (2023) analysis of linguistic mistakes. On this analysis, the standing meanings of expressions provide conditions of correct use for those expressions: a range of conditions in which the standing meaning of the expression permits one to use it. Reiland’s illustrations are as good as any:

...‘Bertrand’ is perhaps for using while you’re thinking or referring to Bertrand and thus its use-conditions are that one has to be performing these mental acts... Similarly, ‘I’ is for using while you’re thinking of yourself in a first-personal or de se way, and ‘is British’ while you’re expressing the property of being British. Finally, a declarative sentence like ‘Bertrand is British’ is for using when one is doing the above acts and further predicating the property of the person, resulting in your entertaining the proposition that Bertrand is British (Soames), or judging it to be the case (Hanks). To take a different sort of example, the expressive interjection ‘Ouch!’ is for using while you’re in pain whereas ‘Oops!’ is for using when you’ve just observed a minor mishap ... Finally, situational terms like ‘Hello!’ or ‘Goodbye!’ might be for using while you’re meeting someone or parting from them etc. (Reiland, 2023, p. 2197)

You misuse an expression—commit a linguistic mistake—if you use it outside of the use conditions provided by its standing meaning.

In some respect, the use of context-sensitive expressions is no different. If, as Reiland suggests, “‘I’ is for using while you’re thinking of yourself in a first-personal or de se way’, then you commit a linguistic mistake in using it where the ordinary English speaker would use “you”. However, at least some context-sensitive expressions give speakers space in which to change how they use the expression without committing a linguistic mistake. We can distinguish between two categories of context-sensitive expression. There are those whose standing meaning completely fixes how context shapes their content—sometimes called *pure indexicals*. The least controversial candidate for such an expression is the word “I”. Then there are those whose standing meaning doesn’t completely fix how context shapes their content—sometimes called *true demonstratives* or *supplementives*. The latter expressions have standing meanings that, very often, give speakers room to select between various contents, where whichever such content they pick, they won’t have made a linguistic mistake ((Buchanan & Schiller, 2022, pp. 13-15), (MacFarlane, 2016, pp. 260-261)).

For example, “smart” should be used to mean *smart* if you’re wanting to speak English. If you use it to mean *small*, then you commit a linguistic mistake. But there are both different scales of smartness and different degrees on such scales to which someone would need to be smart in order for them to count as smart. While we do often criticize people for the contents they assign to words like “smart” (when they’re using “smart” to mean *smart*), we don’t criticize them for having made a specifically *linguistic* mistake. If we’re wondering who’s smart enough to outsmart the professor,

and you say that Harriet's smart, meaning, that she's smarter than 80% of her 6-year-old peers, we can criticize you for being obscure, irrelevant, etc. But we cannot (rightly) accuse you of having used the word against its standing meaning: of sinning against English. The same is true of a great many context-sensitive expressions. We can be criticized for what we count as "tall", "green", the modal base of a modal or conditional, the domain of a quantifier... But there's a wide range of contents we can ascribe to these expressions in context without having made a specifically linguistic mistake. In this sense, many context-sensitive expressions give speakers freedom to choose from a range of different contents.¹

3. Non-linguistic Incentives on language users' content preferences

If a language user merely aims to speak the language from which context-sensitive expressions are taken without committing a linguistic mistake, then *ceteris paribus* that user is free to ascribe to such expressions any content from a variety of different contents without failing to do what they aim to do. Even so, language users' content choices for context-sensitive expressions may nonetheless be influenced by factors beyond what is necessary for avoiding linguistic mistakes. In this section, I will describe 5 non-linguistic incentives that can operate on language users' content preferences. I'm interested in these incentives because they can be channelled to stabilize how language users in diverse contexts ascribe contents to the same context-sensitive expressions (e.g. the same document), as we'll explore in sections 4-7.

3.1. Activity Outcome Incentives

Once a context-sensitive expression is integrated into a broader process, alteration of the content associated with it can often change the outcome of that process (Plunkett & Sundell, 2013). Take a hackneyed but clear example. Imagine we are assessing whether the chilli is "spicy," and the result of this assessment determines whether more chilli powder will be added. The expression "spicy" is plausibly a supplementive: English gives its users freedom to ascribe different contents to "spicy" without risk of linguistic error. However, if you have a specific preference for the spiciness of the finished chilli, your preferred content for "spicy" in this assessment will be influenced by more than just the desire to avoid linguistic mistakes. If you prefer a mild chilli, you will favour a low threshold for the content of "spicy" in this context. If you prefer an intensely hot chilli, you will prefer a high threshold for the content of "spicy" in this context. Your content preferences are influenced by your desired outcome for the process in which "spicy" plays a functional role. Examples of this influence on speaker content preferences abound (see (Gibbons, 2022), (Plunkett & Sundell, 2013), and (Shields, 2021)). I will call incentives to prefer one content over another for a given context-sensitive expression, which arise from understanding the impact of the content choice on the outcome of an associated activity, *activity outcome incentives*.

3.2. Activity Constitutive Incentives

Think again about how the intention to speak a particular language without committing linguistic mistakes constrains how one uses words from that language. One wants it to be true (as one utters

¹ When such expressions are used amongst appropriate other sentences, freedom from linguistic error can dwindle (Stojnić, 2021). But even then, not entirely.

and interprets words from the language) that one is speaking the language without making linguistic mistakes. In order to make that true, one uses the words within their use-conditions. One adapts the way one uses the words in order to meet a constitutive condition on performing an activity one wants to be performing in uttering those words.

We can take that idea and generalize it a little. There are other activities one can be aiming to perform in speaking, which one is performing in speaking only if one assigns the right contents to the context-sensitive expressions one utters in speaking. So, if one wants to be performing those activities, then one will have reason to narrow one's selection of contents for the context-sensitive expressions one uses.² Here's an extended example (modelled on a discussion in (Travis, 1975)).

Take the word "duck". You can use the word "duck" with a content that classifies plastic ducks within the extension of the word. And you can use the word "duck" with a content that doesn't classify plastic ducks within the extension of the word. The standing meaning of the word doesn't tell you which way to go. But the things you mean to be doing in uttering a sentence containing the word "duck" might do. Suppose that in uttering the sentence "That's a duck" whilst pointing at a fairly realistic plastic duck, you mean to be showing your nephew what the huntsman should have been shooting at just after the huntsman accidentally shot such a plastic duck. If that's what you mean to be doing in uttering the sentence, then you really should be using "duck" with a content that excludes plastic ducks. (Consequently, you should consider yourself to have said something false). But now suppose instead that in uttering "That's a duck", pointing at a plastic duck, you mean to be showing your nephew what a duck (as opposed to a goose) looks like, to a certain, rough level of precision. Well, now you don't need to adopt a content that excludes realistic looking plastic ducks in order to be doing what you mean to be doing in uttering this sentence.

This relationship between what activity a speaker means to be doing in uttering a sentence (besides avoiding linguistic mistakes) and what contents context-sensitive expressions in the sentence should be used with, recurs across many examples of linguistic context-sensitivity. Well-formulated examples don't *stipulate* that in two contexts the expression has different contents. They instead provoke us to ascribe different contents by stipulating different activities for language users to pursue in each context. The language users could be irrational and ascribe contents to the relevant expressions which don't serve their purposes (Buchanan & Schiller, 2022). But we assume that they ascribe the contents they should in order to be doing what they want to be doing. You say "There's milk in the fridge." Do you mean to be chastising someone for failing to completely clean the fridge? Or do you mean to be offering them a standard kind of supplement for their coffee? What you mean to be doing changes what content you should be using "milk in the fridge" with. You could do other than what you should be doing. You could be meaning to offer a standard kind of supplement for someone to put in their coffee, and use "milk in the fridge" so that a pool of stale spilt milk at the base of the fridge qualifies as milk in the fridge. But if you did that you wouldn't be offering a standard kind of coffee supplement.

The things you mean to be doing in uttering a sentence that contains a context-sensitive expression provide you with incentives to use context-sensitive expressions with some contents rather than others. I'll call these *activity constitutive incentives*.

² Dobler (2019) and Schoubye and Stokke (2016) propose that the broader purposes of a speaker don't just *incentivize* content ascription in the way I'm describing; they actually determine the content of the sentence used in pursuit of said purposes. This is not a commitment I'm making. For criticism of the Dobler/Schoubye and Stokke view see (Buchanan & Schiller, 2022).

3.3. Incentives to defer: authority and power

Let's say that the user of a context-sensitive expression exhibits semantic deference with respect to that expression if they let others' applications of the expression shape what content they themselves ascribe to the expression. I distinguish two incentives to engage in semantic deference: perceived epistemic authority and social power.

Sometimes one places more confidence in someone else's assessment of the most appropriate content to ascribe to a context-sensitive expression in a particular context than one places in one's own assessment of the same. In this case, one may use the other person's applications of the expression as a guide for determining what content to ascribe to the expression ((Barker, 2002), (Mankowitz, 2021), (Mena, 2023)). For example, let's say Jansen is part of a hiring committee for a pluralistic department that covers both continental and analytic philosophy. However, Jansen was not attentive during the preparatory meetings for the hiring process, so he is unsure whether they are seeking a continental or an analytic philosopher. Consequently, Jansen is uncertain about whether to apply continental or analytic standards when evaluating candidates using expressions like "clear," "compelling," "well argued," etc. Yet, Jansen is confident that other committee members know the desired type of candidate. So he infers the contents they ascribe to these expressions by observing how they apply these expressions to the candidates and follows their lead. I will call this sort of thing an *authority-based incentive*.

However, one might defer for a different reason: the person you defer to can negatively impact your standing if you don't ((Davies, 2016), (Podosky, 2022)). For instance, let's suppose Jansen is well aware of the type of philosopher the department is seeking. However, he also knows that Sue, a powerful member of the department, prefers a different type of candidate and does not respond well to those who challenge her. In this situation, Jansen may defer to Sue on the content of the aforementioned expressions, not because he thinks she knows better what contents would best help the committee identify the candidate the department is supposed to be looking for, but purely out of concern for his career. This is what I'll call a *social-power-based incentive* to defer.

3.4. Filtering

The final incentive that I want to get on the table is what we might call a "Cambridge" incentive (by analogy with "Cambridge changes"). Those who get to use given expressions, in a given context, are sometimes filtered on the basis of which content they are inclined to assign the relevant expressions. For instance, a department might only hire those who use "good philosophers" with a content that excludes philosophers X, Y and Z from the category. Result: the department contains only persons who use "good philosophers" with that exclusive content. This isn't because people in the department are incentivized to adopt that content. You just won't get a seat at the table if you don't have the right content preferences.

4. Organizations as metacontexts

Our problem is how the content of context-sensitive expressions can be stabilized across multiple contexts for ascribing content to said expressions. We've just canvassed incentives (non-linguistic in origin) to whose influence content preferences are susceptible. Let's now return to the notion of

a metacontext as the context of a group of contexts. We're interested in the capacity of a metacontext to channel these incentives to the contexts within it so as to create stable profiles of incentives on content ascription for each such context.

What kind of thing could play this role? What kind of thing assigns users of a document activities whose performance and outcome can be used to guide those users' ascriptions of content to the document? What kind of thing orders users of the document into relations of deference, so that by stabilizing which contents *some* users ascribe to the document we thereby stabilize the contents *other* users ascribe to it?

One thing that can play this incentive-channelling role is an organization. Organizations assign people within them to roles with sets of resources and particular sets of tasks to be carried out with those resources ((Hardimon, 1994), (Rocha, 2011)). Fulfilment of organizational role obligations is materially and socially incentivized, shaping activity and conduct within the organization (Mejia & Skorburg, 2022). An organization could be so structured as to create stable profiles of incentives at a group of contexts within it for the ascription of content to a context-sensitive expression. And in principle, it could do this in a way that promotes ascription of the same content to the same expression at each context.

This proposal has many moving parts. To make the proposal easier to grasp, I present a case study. The case study will include: a particular expression (which, I'll argue, is context-sensitive); two organizational roles within which it may be used; and three alternative organizational structures. I'll explain how the difference between the three structures changes which content occupants of said roles are likely to ascribe to said expression.

5. "Sexual harassment" in the university: the difference metacontext makes

Sexual harassment laws contain expressions which give speakers leeway in interpretation without fear of linguistic mistake (Cahill, 2018[2001], p. 27). Because of this, the organizations within which sexual harassment law is implemented have a significant effect on how those who apply and interpret the law in the organization interpret it ((Cahill, 2018[2001]), (Edelman, 1992), (Edelman & Cabrera, 2020), (Edelman, Erlanger, & Lande, 1993)). I'm going to focus on one expression that appears in sexual harassment documents in universities that implement sexual harassment law: only conduct that is *unwanted* can qualify as sexual harassment.

5.1. The context-sensitivity of desire ascriptions

Here are two ways in which the truth-conditions of desire ascriptions are underdetermined by their standing meaning.

Firstly, there's a difference between what you want *all things considered* and what you want *some things considered*. You have a strong urge to take a nap (you are tired) but you know that bad things will happen to your career if you do. In that case:

there are indeed contexts where it's not true to say that you want to take a nap. (Your partner might ask if he should pull down the shades so it'll be easier to nap, and you could reply by saying, "No, I don't want to take a nap. I need to study.") Nonetheless, there are also contexts where it is true to say that you want to take a nap. (You might say, "I want to

take a nap, but I'm going to do everything I can to stay awake because I need to study.”). (Phillips-Brown, manuscript, pp. 2-3)

The first use is an all-things-considered use: “X wants that P” is true if and only if, given everything that's relevant, X wants that P. The second use is a some-things-considered use: “X wants that P” is true if and only if, given a proper subset of things that are relevant, X wants that P. The two uses have different contents: they're true of different circumstances.

Secondly, suppose you don't know that you need to study for tomorrow. So you stop studying and get ready for bed. The sentence “X wants to study”, made about you, in this predicament, could be used to make two different claims (Jerzak, 2019). Firstly, it could be used to explain your current behaviour. You're going to bed now because you want to study (a false claim). Secondly, it could be used to specify what you should be doing, given the facts: you want to study (said in the same way someone giving directions might say “you want to take a left just after the crossroad”—a true claim). Although Jerzak and others call this latter use “advisory”, it doesn't have to be advice. You could say to yourself, for instance, “I don't know what I want because I haven't ascertained all the facts yet” even when you currently have a strong preference for one particular course of action. It's just that you know that your preferences are likely to change upon receiving further information, and you're using “want” to describe what you would prefer given fuller information. This is still what Jerzak calls an “advisory use”.

The first distinction (all things versus some things considered) doesn't concern ignorance: it's just a question of whether you're bracketing certain pieces of information you have available to you from a description of your own preferences. The second distinction (predictive versus advisory) concerns ignorance: it's a question of whether you're describing your preferences given the information you currently have or instead given information some of which you don't currently have. Desire descriptions can vary in their contents along both dimensions.

These two different ways in which a desire ascription's content can vary can make a difference to how conduct (which some might call “sexual harassment”) is classified. Consider two examples of conduct, which I'll call *Grooming* and *Bounded Choices*. Both kinds of conduct are based on descriptions of conduct routinely experienced by students in UK universities (Bull & Page, 2021).

Grooming

B's professor has, intentionally or not, given her the impression that it's normal for professors and students to engage in various forms of conduct which are not normal for those roles. Thinking that this is normal conduct for students and professors, she has urges to do these activities: e.g. talk over coffee in a secluded café on personal topics including previous sexual and romantic experiences. However, were she fully informed of what really is normal for professor and student relationships (and her professor's history), she would recoil at this activity.

Take the desire ascription “she wanted this” applied to B and her professor's conduct. A predictive content for this ascription is true and thus disqualifies the conduct from being sexual harassment. But an advisory content is false and thus does not disqualify this conduct from being sexual harassment.

Bounded Choices

B's professor has made it clear (not necessarily by intentionally communicating it, but perhaps by it being an obvious consequence of his conduct and dispositions), that B will get access to standard features of graduate education only if she submits to his attempts to make their professional relationship more personal (*vis a vis* time together, topics discussed, physical intimacy). B has no urge to undergo this conduct or participate in it (as she might do with a romantic partner who has no power over her career and who she is attracted to). But out of concern for what would happen to her career if she didn't go along with this conduct, then, all things considered, she thought she should go along with it.

Again, take the desire ascription "she wanted this" applied to B and the conduct of her professor. With a some-things-considered content, this sentence is false and so would not disqualify the conduct from being sexual harassment. But with an all-things-considered content, the sentence is true and so would disqualify the conduct from being sexual harassment.

I'll call contents for "sexual harassment" which disqualify (via the content of "unwanted") Grooming and Bounded Choices as cases of sexual harassment, *exclusive contents*. I'll call contents that don't disqualify this conduct as cases of sexual harassment, *inclusive contents*.

5.2. Two contexts in Three Metacontexts

Two contexts in a university in which one will have opportunity and reason to decide whether "sexual harassment" applies to Grooming and/or Bounded Choices are: the investigator of whether a professor has sexually harassed a student and the accuser. I'll describe three metacontexts: Universities A, B and C. University A incentivizes occupants of each role to ascribe exclusive contents to "sexual harassment." University B does not do this for either role. University C incentivizes exclusive contents for the accuser role but not for the investigator role. The differences between University A and University B show how we can change the *de facto* content of a context-sensitive expression by changing the metacontext in which the document is interpreted. The difference between Universities A and B on the one hand, and University C on the other, illustrate the difference between metacontexts that support an expression's capacity to share content and metacontexts which don't.

5.2.1. University A: incentivizing exclusive contents

In University A, the investigator is the accused professor's head of department. There are no sexual harassment accusation training preconditions on who gets to occupy the position of department head nor on anyone else in the department—including students (a not uncommon predicament (Bull & Rye, 2018, pp. 19-20)). This metacontext will generate the following incentives for occupants of each role.

Consider first the investigator role.

Firstly, the investigator will have a conflict of interest built into their role (Stark, 2001): the investigator will experience negative outcomes if the investigator finds that sexual harassment has

taken place.³ It will be embarrassing for the department. It'll threaten recruitment of students—which in turn harms the financial security of the department. It could involve loss of funds if those funds are dependent on the accused professor—either because he is a name that pulls in students or because he has a grant attached to him. The investigator will be a colleague of the accused and there may be collegial pressures not to find a colleague to have sexually harassed a student. These negative outcomes are activity outcome incentives to prefer exclusive contents for “sexual harassment”.

Secondly, insofar as it is a job expectation of the investigator to protect the resources of the university, the investigator will also experience an activity constitutive incentive to adopt an exclusive content for “sexual harassment”. To use “sexual harassment” in a way that is not exclusive would be a way to fail to be defending the resources of the university (jeopardising financial and human resources). So, much as a lawyer, to be doing her job, should look for ways to interpret rules, principles and guidelines which favour her client, an investigator who is employed to protect the resources of the university will look for an interpretation of “sexual harassment” which allows him to count as defending those resources.

Thirdly, the absence of any training prerequisites upon entrance into the investigator role will mean there is no filtering of who enters the role on the basis of content preferences for “sexual harassment”: nothing informed by legal precedent, nothing to ensure that they take into consideration the perspective of the accuser.

Consider now the accuser role.

Firstly, because she hasn't received any relevant education or training, and she is in an environment in which she normally defers to the expertise of university staff, she's likely to have more confidence in the ability of university staff to identify the appropriate content for “sexual harassment” than her own, resulting in authority-based semantic deference to such staff ((Bull & Rye, 2018, p. 13), (Whitley & Page, 2015, p. 39)). Since the staff are incentivized to favour exclusive contents for “sexual harassment”, the student will be likely to end up doing the same.

Secondly, accusers are dependent upon staff for various goods that can be withheld from them if they lose the favour of staff—grades, access to educational opportunities and resources, letters of reference, networking opportunities ((Bull & Rye, 2018, p. 14), (Quinn, 2000)). Moreover, if the person they are to report sexual harassment to has an in-role conflict of interest, that will probably be obvious to the accuser. So any accuser who seeks to challenge uses of “sexual harassment” with an exclusive content by the investigator will know that in doing so they risk losing the favour of someone who has the ability to negatively influence their future. This is a social power-based incentive on the accuser to defer to the investigator's exclusive use of “sexual harassment” (at least, in official settings).

In University A, content is shared across investigator and accuser role. Both are incentivized to prefer exclusive contents for “unwanted” and therefore “sexual harassment.”

³ Not all universities have in-role conflicts of interest of this extremity. But this isn't essential to the example. Anyone who is paid by the university to protect the resources of that university and who is at the same time asked to investigate whether a member of staff has sexually harassed a student faces the same conflict as described above (see (Bull & Page, 2022, p. 21), (Edelman & Cabrera, 2020, p. 374), (Kihnley, 2000, p. 72), (Marshall, 2005, pp. 99-100, p.118)).

5.2.2. University B: disincentivizing exclusive contents

At University B, the investigator is drawn from an Equity and Human Rights Office (EHRO), whose employees are paid expressly, not to protect the (human and other) resources of the university, but instead to safeguard social justice within the university. These employees are hired only if they have training which favours inclusive contents for “sexual harassment”. All students are trained by the Equity and Human Rights Office to accept inclusive contents for “sexual harassment” and told to report sexual harassment only to employees of the EHRO. Students are made fully aware of the incentives to which employees of this office are subject.

Consider the investigator role.

Firstly, the activity constitutive and activity outcome incentives experienced by the investigator in University A are absent. The investigators do not have in-role conflicts of interest and their job descriptions make clear that their job is not to protect its resources, but instead, to ensure justice.

Secondly, in University B those entering the investigator role are filtered based, in part, on their content preferences *vis-à-vis* “sexual harassment”, so that you get a seat at this table only if you prefer inclusive contents.

Consider the accuser role.

The authority-based incentive (that exists in University A) to defer to staff who prefer exclusive contents is replaced by an incentive to defer to staff who favour inclusive contents. The social-power-based incentive to defer to department staff who favour exclusive contents is neutralized insofar as those professors are incentivized by the powers of the EHRO (in a way students themselves understand) not to harm students who ascribe inclusive contents to “sexual harassment”.

In University B, content is again shared across investigator and accuser roles. But whereas in University A this content was exclusive, in University B it's inclusive.

5.2.3. University C: incentivizing different contents in different contexts

In University C, the investigator role is as described in University B, except that the EHRO defers the job of training students to individual departments. Seizing this opportunity to hide incidents from the EHRO that it would most likely classify as sexual harassment, the individual departments avoid giving their students an inclusive understanding of “sexual harassment”, avoid informing the students of the powers the EHRO has over them, and direct students to report in the first instance to the head of department: who has the same in-role conflict of interest (and absent training) as in University A.

Occupants of the investigator role will be incentivized to favour an inclusive content for “sexual harassment” – as in University B. But students will be incentivized to favour an exclusive content. They'll defer to people who have in-role conflicts of interest that drive preference for an exclusive content—as in University A. In University C, investigators are incentivized to adopt inclusive contents whilst accusers are incentivized to adopt exclusive contents.

5.3. Summary and Clarification

We began with a proposal with many moving parts. We've now worked through a case study of this proposal. The case study illustrates how ascribed linguistic content and organizational design can interact: both how organizational design can change the content that persons occupying different roles in the same organization are likely to ascribe to a given expression, and how it can change whether those persons will be likely to ascribe the *same* content to this expression.⁴

I now want to make two points of clarification.

Firstly, the differences between universities A-C are to be understood as generating differences in trends. I'm not suggesting, for instance, that no one in the accuser or investigator role in University A would ever adopt an inclusive content for "want" (and hence "sexual harassment"). I'm saying that the differences between the universities make it likely that persons occupying the relevant roles will go for the corresponding contents. My interest is in how metacontexts change likelihoods of content-sharing. A concept of metacontext understood in any other way would be impractical.

Secondly, we've been concerned with what contents language users will be incentivized to ascribe to an expression. Such contents aren't necessarily the same as the content that is best qualified as the content of the expression i.e. the content that it is correct to ascribe to the expression. To see how the two can come apart, think of a bank employee whose job is to apply the bank's criteria for appraising loan clients. Suppose a person requesting a loan sleeps with this employee and this incentivizes the employee to interpret the terms of appraisal more liberally than the bank and its legal context permit. Despite the incentive, the employee could be disciplined or prosecuted for misapplying the terms of appraisal. The existence of the sexual incentive didn't change the content of the criteria even if *ex hypothesi* it did change what content the employee ascribed to the criteria.

Should I be focusing on the correct content of a context-sensitive expression rather than on which content language users will be inclined to ascribe to the expression? My interest has been in cross-contextual content-sharing: what might help and hinder it. This purpose requires us to consider what contents language users actually ascribe to the expressions they use (i.e. produce and consume). For what matters is what contents 2+n language users ascribe to the same expressions. So that's what I focus on.

6. Beyond organizations as metacontexts

The case study shows how organizations can function as metacontexts. But if organizations were our sole model of metacontext, we would be likely to think that the walls of the organization are the limits of the cross-contextual stability that metacontexts can provide. This would be to underestimate the range of ways in which a metacontext can be realized. In this section, I explain how a metacontext can extend beyond the confines of a single organization.

The internal structure of an organization is shaped by its environment (Daft, 2010). Hence, appropriate design of the environment of a set of organizations can serve to spread a characteristic

⁴ For further case studies, I recommend the literature on policy implementation. See for example Degn (2020) on the content ascribed to "research integrity", Talesh (2012) on the content ascribed to consumer-rights legislation, and Molnár-Gábor, et al. (2022) on the content ascribed to the General Data Protection Regulation.

internal structure across those organizations, and thereby a characteristic profile of incentives accompanying parallel roles within those organizations.

Return again to the sexual harassment case study. Universities have sexual harassment policies largely because they are vicariously liable for sexual harassment executed by their employees ((Hawkesworth, 1997), (Thomas, 2004), (Whitfield & Dustin, 2015)). Commonly laws that make universities liable allow them to protect themselves from this liability without protecting students from harassment by constructing what Lauren Edelman calls “symbolic structures”: procedures that look OK legally but are dysfunctional *vis a vis* justice ((Edelman, 1992), (Edelman & Cabrera, 2020), (Edelman, Erlanger, & Lande, 1993)). This legal context helps proliferate universities resembling our University A: the cross-contextual stability of the exclusive content for “sexual harassment” witnessed *within* University A is to be expected *across* universities.

If we wanted to increase the prevalence of universities that resemble University B, we would need a regulatory system that doesn’t tolerate symbolic structures. A good start would be a regulatory system that increases the risk of legal censure for the creation of in-role conflicts of interest. One feature of University B that made it conducive to inclusive contents for “sexual harassment” was the fact that investigators are chosen by an Equity and Human Rights Services unit (rather than by those who have, to varying degrees, in-role conflicts of interest). University B is based on Western University (Western University, Equity and Human Rights Services, 2017). Western University is in Ontario. The Ontario Human Rights Policy (which regulates sexual harassment in Ontario) includes guidelines on the design of sexual harassment procedures that explicitly recommend that the procedure include ‘an assurance that the person handling complaints should be independent, expert, etc.’ (section 8 of (Ontario Human Rights Commission, 2013)). Requiring university units which specialize in justice to appoint investigators is an effective means to ensure this. But they are expensive: we shouldn’t expect smaller universities in Ontario to have such units. Of those universities which are comparable in size to Western, six out of seven have their investigators chosen by such a unit.⁵ Ontario promotes the prevalence of universities like University B and thereby contexts for interpretation of “sexual harassment” which incentivize inclusive contents for that expression. If a metacontext is a context of a group of contexts which channels incentives for contents in each context in the group, then the jurisdiction of the Ontario Human Rights Policy (something which is much larger than a single organization) has a claim to being a metacontext.

Notice that the Ontario Human Rights Policy doesn't help standardize "sexual harassment" across organizations by adding more text to its definition. Instead, it shapes how individuals in those organizations ascribe content to the existing definition of "sexual harassment" by shaping organizational design. That the same laws which define “sexual harassment” also promote cross-organizational content stability is incidental to the example. If we don’t appreciate this, then (once again) we’ll underestimate the range of ways in which a metacontext can be realized.

To drive this point home, here’s a brief example of a metacontext that extends beyond a single organization but not by means of supra-organizational legal regulation. Think of temporal language

⁵ Western has approximately 37,000 students. I define universities in Ontario of comparable size as those that have at least 30,000 students (McMaster University, Toronto Metropolitan University, University of Ottawa, University of Toronto, University of Waterloo, Western University, York University – <https://www.univcan.ca/universities/facts-and-stats/enrolment-by-university/>). The University of Waterloo is a borderline case: technically the associate dean selects the investigator, but this is done under the guidance of the Conflict Management and Human Rights Office (see 7.1 of <https://uwaterloo.ca/secretariat/sexual-violence-response-protocol-and-procedures-re>).

("6pm", "noon", "midnight"). It wasn't always the case that people living in different settlements (but in what we would nowadays call the same "time zone") referred to the same times with these expressions: there was no standardization across different settlements in the contents assigned to these expressions (Zerubavel, 1982). Before any law was created that required inter-settlement standardization, the railway system came into existence. The railway operator needed a stable inter-settlement way of referring to time in order to coordinate the movement of its trains across settlements. Settlement occupants needed to synchronize clocks with the railway operators in order to facilitate the catching of trains. Laws regulating temporal language came only later. What standardized temporal language across settlements was not in the first instance a law, but rather the development of a new technology and accompanying activity that rechannelled incentives on content ascription to temporal language.

Rather than thinking of metacontexts solely as organizations, it is better to think of them more generally as a kind of infrastructure that accompanies a text across the different contexts in which it is interpreted. To paraphrase Bruno Latour (1983): just as a train can travel through a cornfield, *provided that you build the track for it first*, so a context-sensitive text can bear a stable content across many a context, *provided that you build an appropriate metacontext for it first*. The bible has its church, the textbook its education system, the form its bureaucracy, and the manifesto its social movement. These social institutions are not just containers. They help stabilize the content that diversely located users of these texts ascribe to them, and thereby play a crucial role in making them cross-contextually stable vehicles of content.

7. Beyond the old contextualism debate

The literature of the 2000s and 2010s that concerned the cross-context question can be divided into two sides: the contextualists who attempted to show that linguistic context-sensitivity is widespread; and those who attempted to protect appearances (effortless, uncomplicated cross-contextual content-sharing) against the tension between these appearances and widespread context-sensitivity. Contextualists' primary concerns were two: to show that context-sensitivity is widespread and to replace the idea of a community of language co-speakers as a group which shares a fixed set of contents, with the view that language co-speakers constantly construct contents ("ad hoc concepts", a "dynamic lexicon") on the fly, conversation by conversation. Contextualists spent comparatively little time seriously addressing the question their conclusions raised: the cross-context question—see the space given to this issue in (Carston, 2002), (Recanati, 2010), and (Sperber & Wilson, 1986) relative to the space spent defending widespread context-sensitivity. This was an eminently reasonable way to allot argumentative resources given how controversial context-sensitivity was at the time. But it didn't help bring about consensus on the cross-context question.

In two ways, the concept of metacontext allows us to advance beyond the debate as it was.

Firstly, the concept allows us to see that each side of the debate got something right. Even if standing meaning gives freedom of content ascription to language users, metacontexts can channel incentives across those contexts in a way that directs users' ascriptions toward the same content to the same expression across multiple contexts. Moreover, because it's the metacontext that brings this about, individual language users who find themselves in such metacontexts won't be put out by the effort: content-sharing therein will feel easy. In these metacontexts, context-sensitivity is compatible with effortless cross-contextual content-sharing. But other metacontexts will cause

systematic divergences in the contents that occupants of different contexts are likely to assign to the same expressions: as the likes of Cappelen and Lepore foresaw.

Secondly, the problem of cross-contextual content-sharing with context-sensitive language was posed originally at the level of human psychology and general linguistic competence. You couldn't solve the problem in one case without solving it in every case because there was no division of the problem into "cases and cases". Adoption of metacontext as a unit of analysis makes this division possible. It allows us to see content-sharing troubles not as direct consequences of human psychology and general linguistic competence, but as consequences of contingent design choices and accidents in circumscribable social settings. Take for instance the problem of designing the organizations and legal contexts that implement the GDPR so that the spirit of the GDPR is preserved in practice (see footnote 4). This is clearly a problem that can be solved without having to solve every content-sharing problem at once. The concept of metacontext thus has the potential to convert a quandary about human communication into a progressive research programme (the regular solving of solvable puzzles). Such a programme would combine the empirical and conceptual tools of linguistics and experimental philosophy of language with theory and findings from studies of organizations, public administration, and bureaucracy. It would study particular pairings of expression and metacontext in order to ascertain how a metacontext's structure incentivizes fragmented interpretation of an expression that should be understood homogeneously. It could offer recommendations on reform to improve the content-sharing capacity of a targeted metacontext. However, as content relativists (like (Cappelen, 2008)) and scholars of the language of law and policy (like ((Lanius, 2019), (Lipsky, 2010 [1980]), (Soames, 2011), (Väyrynen, Lanamäki, Laari-Salmela, Iivari, & Kinnula, 2022)) have noted, documents which are ascribed different contents in different contexts can be useful. For instance, if the appliers of a document are supposed to pursue a given purpose in applying it, but the contents that best serve that purpose are difficult to discern without detailed knowledge of the context of implementation, then it's best to give appliers of the document discretion in ascribing it content. The envisaged research programme would therefore also explore the utility of such metacontexts.⁶

References

- Abreu Zavaleta, M. (2021). Communication and Variance. *Topoi*, 40, 147-169.
- Barker, C. (2002). The Dynamics of Vagueness. *Linguistics and Philosophy*, 25, 1-36.
- Bezuidenhout, A. (2002). Truth-Conditional Pragmatics. *Philosophical Perspectives*, 16, 105-134.
- Bowker, M. (2022). Ineliminable underdetermination and context-shifting arguments. *Inquiry*, 65, 215-236.

⁶ I presented versions of this paper at the University of Tartu Troubling Gender Winter School (2020), the Words Workshop reading group (2021), the Vienna Forum for Analytic Philosophy Graduate Conference (2021), the ALEF Research Seminar (2022), the Oslo Testimony and Communication workshop (2022), and the University of Reading Department of Philosophy Seminar (2023). It has also received constructive feedback from Nat Hansen, Paul Podosky, William Tuckwell, Tomasz Zyglewicz and reviewers for the *Philosophical Quarterly*. I'm grateful to these audiences and persons for helping me to think and say better what I first tried to articulate some three years ago.

- Buchanan, R., & Schiller, H. (2022). Pragmatic Particularism. *Philosophy and Phenomenological Research*, 105(1), 62-78.
- Bull, A., & Page, T. (2021). Students' accounts of grooming and boundary-blurring behaviours by academic staff in UK higher education. *Gender and Education*, 33, 1057-1072.
- Bull, A., & Page, T. (2022). The Governance of Complaints in UK Higher Education: Critically Examining 'Remedies' for Staff Sexual Misconduct. *Social and Legal Studies*, 31, 27–49.
- Bull, A., & Rye, R. (2018). *Silencing Students: Institutional Responses to Staff Sexual Misconduct in UK Higher Education*. Portsmouth, UK: The 1752 Group/University of Portsmouth.
- Cahill, M. (2018[2001]). *The Social Construction of Sexual Harassment Law: The Role of National, Organizational and Individual Context*. Abingdon: Routledge.
- Cappelen, H. (2008). The Creative Interpreter: Content Relativism and Assertion. *Philosophical Perspectives*, 22(1), 23-46.
- Cappelen, H., & Dever, J. (2016). *Context and Communication*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Oxford: Oxford University Press.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Oxford University Press.
- Daft, R. (2010). *Organization Theory and Design*. Cincinnati: South-Western College Pub.
- Davies, A. (2016). How to Silence Content with Porn, Context and Loaded Questions. *European Journal of Philosophy*, 24(2), 498-522.
- Davies, A. (2019a). Testimonial Knowledge and Context-Sensitivity: a New Diagnosis of the Threat. *Acta Analytica*, 34(1), 53-69.
- Davies, A. (2021). Communicating in contextual ignorance. *Synthese*, 199, 12385-12405.
- Degn, L. (2020). Integrating Integrity: The Organizational Translation of Policies on Research Integrity. *Science and Engineering Ethics*, 26(6), 3167-3182.
- Dobler, T. (2019). Occasion-sensitive semantics for objective predicates. *Linguistics and Philosophy*, 42(5), 451–474. doi:<https://doi.org/10.1007/s10988-018-9255-x>
- Edelman, L. (1992). Legal Ambiguity and Symbolic Structures: Organizational Mediation of Civil Rights. *American Journal of Sociology*, 97(6), 1531-1576.
- Edelman, L., & Cabrera, J. (2020). Sex-Based Harassment and Symbolic Compliance. *Annual Review of Law and Social Science*, 16, 361–383.
- Edelman, L., Erlanger, H., & Lande, J. (1993). Internal Dispute Resolution: The Transformation of Civil Rights in the Workplace. *Law and Society Review*, 27(3), 497-534.
- Eisenberg, E. (1984). Ambiguity as Strategy. *Communication Monographs*, 51, 227-242.
- Gibbons, A. (2022). Rational conceptual conflict and the implementation problem. *Inquiry*.
- Hardimon, M. (1994). Role Obligations. *The Journal of Philosophy*, 91(7), 333-363.

- Hawkesworth, M. (1997). Challenging the Received Wisdom and the Status Quo: Creating and Implementing Sexual Harassment Policy. *NWSA Journal*, 9(2), 94-117.
- Jerzak, E. (2019). Two ways to want? *Journal of Philosophy*, 116(2), 65-98.
- Kennedy, C., & McNally, L. (2005). Scale Structure, Degree Modification, and the Semantics of Gradable Predicates. *Language*, 81(2), 345-381.
- Kihnley, J. (2000). Unraveling the Ivory Fabric: Institutional Obstacles to the Handling of Sexual Harassment Complaints. *Law and Social Inquiry*, 25(1), 69-90.
- Lanius, D. (2019). *Strategic Indeterminacy in the Law*. Oxford: Oxford University Press.
- Latour, B. (1983). Give Me a Laboratory and I Will Raise the World. In K. Knorr-Cetina, & M. Mulkay (Eds.), *Science Observed: Perspectives on the Social Study of Science* (pp. 141-170). London & Beverly Hills: C.A. Sage.
- Lipsky, M. (2010 [1980]). *Street-Level Bureaucracy: Dilemmas of the Individuals in Public Services*. New York: Russell Sage Foundation.
- MacFarlane, J. (2016). Vagueness as Indecision. *Aristotelian Society: Supplementary Volume*, 90(1), 255-283.
- Mankowitz, P. (2021). How to have a metalinguistic dispute. *Synthese*, 199, 5603-6522.
- Marshall, A.-M. (2005). Idle Rights: Employee's Rights Consciousness and the Construction of Sexual Harassment Policies. *Law and Society Review*, 39(1), 83-124.
- Mejia, S., & Skorburg, J. A. (2022). Malleable character: organizational behavior meets virtue ethics and situationism. *Philosophical Studies*, 179, 3535-3563.
- Mena, R. (2023). Metalinguistic effects. *Mind and Language*, 38(2), 545-565.
- Molnár-Gábor, F., Sellner, J., Pagil, S., Slokenberga, S., Tzortzatou-Nanopoulou, O., & Nyströmb, K. (2022). Harmonization after the GDPR? Divergences in the rules for genetic and health data sharing in four member states and ways to overcome them by EU measures: Insights from Germany, Greece, Latvia and Sweden. *Seminars in Cancer Biology*, 84, 271-283.
- Ontario Human Rights Commission. (2013, May). Policy on preventing sexual and gender-based harassment. Retrieved May 13, 2021, from <http://www.ohrc.on.ca/en/policy-preventing-sexual-and-gender-based-harassment-0>
- Peet, A. (2016). Testimony and the epistemic uncertainty of interpretation. *Philosophical Studies*, 173(2), 395-416.
- Phillips-Brown, M. (manuscript). Some-things-considered desire.
- Plunkett, D., & Sundell, T. (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, 23, 1-37.
- Podosky, P. (2022). Agency, Power, and Injustice in Metalinguistic Disagreement. *The Philosophical Quarterly*, 72(2), 441-464.
- Quinn, B. (2000). The Paradox of Complaining: Law, Humor, and Harassment in the Everyday Work World. *Law and Social Inquiry*, 25(4), 1151-1185.

- Recanati, F. (2010). *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Reiland, I. (2023). Linguistic Mistakes. *Erkenntnis*, *88*, 2191–2206.
- Ritchie, K. (2017). Plural and Collective Noun Phrases. In M. Jankovic, & K. Ludwig (Eds.), *The Routledge Handbook of Collective Intentionality* (pp. 464-475). London: Routledge.
- Rocha, J. (2011). Autonomy Within Subservient Careers. *Ethical Theory and Moral Practice*, *14*, 313-328.
- Schoubye, A., & Stokke, A. (2016). What is said? *Nous*, *50*(4), 759-793.
- Shields, M. (2021). Conceptual Domination. *Synthese*, *199*, 15043–15067.
- Soames, S. (2011). What Vagueness and Inconsistency Tell US about Interpretation. In A. Marmor, & S. Soames (Eds.), *Philosophical Foundations of Language in the Law* (pp. 31-57). Oxford: Oxford University Press.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Stark, A. (2001). Comparing Conflict of Interest Across the Professions. In M. Davis, & A. Stark (Eds.), *Conflict of Interest in the Professions* (pp. 335-351). Oxford: Oxford University Press.
- Stojnić, U. (2021). *Context and Coherence*. Oxford: Oxford University Press.
- Szabo, Z. (2006). Sensitivity Training. *Mind and Language*, *21*(1), 31-38.
- Szabó, Z. G., & Stanley, J. (2000). Quantifier Domain Restriction. *Mind and Language*, *15*(2-3), 219-261.
- Talesh, S. (2012). How Dispute Resolution System Design Matters: An Organizational Analysis of Dispute Resolution Structures and Consumer Lemon Laws. *Law & Society Review*, *46*(3), 463-496.
- Thomas, A. (2004). Politics, policies and practice: assessing the impact of sexual harassment policies in UK universities. *British Journal of Sociology of Education*, *25*(2), 143-160.
- Travis, C. (1975). *Saying and Understanding: a generative theory of illocutions*. Oxford: Basil Blackwell.
- Väyrynen, K., Lanamäki, A., Laari-Salmela, S., Iivari, N., & Kinnula, M. (2022). Policy Ambiguity: a Problem, a Tool, or an Inherent Part of Policymaking? *ICIS Proceedings*, *7*, 1-9.
- Western University, Equity and Human Rights Services. (2017, July 13). POLICY 1.35 – Non-Discrimination/Harassment Policy. Ontario, Canada. Retrieved 05 13, 2021, from https://www.uwo.ca/univsec/pdf/policies_procedures/section1/mapp135.pdf
- Whitfield, L., & Dustin, H. (2015). *Spotted: Obligations to Protect Women Students' Safety & Equality*. End Violence Against Women Coalition.
- Whitley, L., & Page, T. (2015). Sexism at the Centre: Locating the Problem of Sexual Harassment. *new formations: a journal of culture/theory/politics*, *86*, 34-53.
- Zerubavel, E. (1982). The Standardization of Time: A Sociohistorical Perspective. *American Journal of Sociology*, *88*(1), 1-23.

