

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Arguing with experts: Subjective disagreements on matters of taste

Permalink

<https://escholarship.org/uc/item/8921n58s>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Kaiser, Elsi

Rudin, Deniz

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Arguing with experts: Subjective disagreements on matters of taste

Elsi Kaiser (emkaiser@usc.edu)

Department of Linguistics, University of Southern California
Los Angeles, CA 90089-1693 USA

Deniz Rudin (drudin@usc.edu)

Department of Linguistics, University of Southern California
Los Angeles, CA 90089-1693 USA

Abstract

When two people disagree about matters of taste, neither is in the wrong: There is nothing contradictory in a dialog where one interlocutor says ‘The rollercoaster was scary!’ and the other responds ‘No, it was not scary.’ This contrasts with disagreements about objective facts. This phenomenon is known as faultless disagreement, and is central for theorizing about subjective expressions. Faultless disagreement is typically assumed to stem from subjective expressions having a special semantics. We present evidence that people’s judgments of faultless disagreement are sensitive not only to the lexical content of a sentence, but also to the broader discourse context (properties of the interlocutors in the dialog) and to extra-contextual factors (participants’ own attitudes about that particular domain). These results problematize arguments that faultless disagreement stems directly from the semantics of subjective lexical items.

Keywords: subjectivity; predicates of personal taste; faultless disagreement; expertise; discourse context; semantics-pragmatics interface

Introduction

The information that we encounter on a daily basis involves both objective facts about the world and people’s subjective opinions. This distinction is also reflected in language: Words that express an individual’s subjective opinion about something (e.g. adjectives like *tasty*, *fun*, *amazing*, *irritating*) differ fundamentally from words conveying more objective facts (e.g. adjectives like *wooden*, *gluten-free*, *Californian*, or descriptions like *grown in Oregon*). By definition, subjective expressions are perspective-sensitive and reflect someone’s opinion or attitude. In other words, they are anchored to an opinion-holder or *judge*, to use a term popularized by Lasnik (2005).¹

Because subjective adjectives express the opinion of a particular individual, in a dialog such as (1a), both speakers can be right in the sense of having made no factual mistakes. In this situation, Sam is expressing the opinion that the wine is tasty-for-Sam, while Alex is expressing the opinion that the wine is not tasty-for-Alex. No one is in the wrong; each interlocutor is entitled to their own opinion and point-of-view. More colloquially: There’s no accounting for taste. Ex.(1a) contrasts starkly with a dialog with an

objective expression, as in (1b). Here, one of the two interlocutors must be wrong about where the wine is grown.

(1a) *Faultless disagreement* => *both people can be right*

Sam: This wine is tasty.

Alex: No, this wine is not tasty. [subjective]

(1b) *‘Regular’ disagreement* => *one of the people is wrong*

Sam: This wine is grown in Oregon.

Alex: No, this wine is not grown in Oregon. [objective]

Faultless disagreement

In the semantics and philosophy literature, the dialog in (1a) is described as licensing an inference of *faultless disagreement*: neither interlocutor is felt to be wrong or at fault (e.g., Koelbel 2004 and many others). This contrasts with ‘faulty’ disagreements (1b), where one person must be wrong (at fault). In essence, presence of faultless disagreement – a situation where two people seem to disagree but actually, neither one is felt to be in the wrong – is widely used by linguists and philosophers as a diagnostic to detect subjective, opinion-based content.

Although the intuition of faultlessness is simple, it has far-reaching theoretical consequences. Despite being a truism since antiquity, faultless disagreement is surprising and problematic for standard semantic assumptions, on which for any proposition *p*, either *p* or $\neg p$ must be false. Indeed, this holds for examples like (1b) – either the wine is grown in Oregon or it is not grown in Oregon. But how can it be that in (1a), neither interlocutor is in the wrong, even though one says the wine is tasty and the other says the wine is not tasty? This poses a challenge for standard semantic assumptions.

As a result, faultless disagreement has led to proposals in both linguistics and philosophy that assign a special semantics to subjective expressions, including the well-known class of subjective adjectives known as *predicates of personal taste* (PPTs, e.g. *tasty*, *fun*, *amazing*) – which are the focus of this paper – as well as many other kinds of subjective expressions including epistemic modals (e.g. *must*, *might*), and aesthetic and moral terms (e.g. *beautiful*, *wrong*, see Lasnik 2005, Stephenson 2007, MacFarlane 2014, McNally & Stojanovic 2017, and many others).

Faultless disagreement has also been put forward as an empirical diagnostic for subjectivity, which is argued to

¹ In this paper, we use the term ‘judge’ in a theory-neutral way.

carry explanatory weight with respect to other linguistic phenomena, including syntactic patterns such as adjective-ordering preferences (*the big red car* vs. *?*the red big car*; see e.g. Scontras et al. 2017).

Despite the central importance of faultless disagreement to motivating exceptional analyses of subjective expressions and its status as a widely-used test for subjectivity, surprisingly little attention has been given to its empirical profile. Given the importance placed on the phenomenon of faultless disagreement as an empirical desideratum for various proposals, it is important to gain a better understanding of what it is actually diagnosing and how it works.

Prior experimental work

In recent work (Kaiser & Rudin 2020), we argued that contrary to what is often (implicitly) assumed by prior approaches, faultless disagreement is not simply a property of the subjective predicate² on its own. We provided experimental data that judgments of faultless disagreement are modulated by the prevalence of opinions in a population, which crucially varies with the object of predication.

In a study with U.S.-based participants, we compared disagreements about widely-liked foods (e.g. chocolate, donuts, pizza) and divisive foods (e.g. blue cheese, anchovies). We hypothesized that if judgments of faultless disagreement are determined only by the (subjective) nature of the predicate on its own, the level of consensus about a particular food's tastiness (or lack thereof) should have no effect. However, our results show that disagreements about divisive foods (ex.2b) were rated more faultless than disagreements about widely-liked foods (ex.2a).

(2a) *widely-liked food (high level of consensus)*

One person says: Chocolate is delicious.

Another person says: Chocolate is not delicious.

(2b) *divisive food (lower level of consensus)*

One person says: Blue cheese is delicious.

Another person says: Blue cheese is not delicious.

Intuitively speaking, participants were more willing to say that both people are in the right when the food is known to elicit disagreements (e.g. blue cheese), but if the food is widely-liked (e.g. chocolate), the disagreement was judged relatively more 'faulty' (i.e. more towards the 'one person is wrong' end of the scale).

These results indicate that judgments of faultless disagreement are sensitive to the lexical content of the entire sentence. We concluded that instead of what is often assumed, judgments of faultless disagreement do not directly reflect the subjectivity of a predicate – rather, participants' judgments of whether a disagreement is faultless are modulated by object of predication. People's

judgments of faultlessness depend on the prevalence of a judgment within the relevant population.

Going beyond the lexical content of the sentence

The results of our earlier work (Kaiser & Rudin 2020) provide initial evidence that judgments of faultless disagreement cannot be derived purely from the subjective adjective itself: we also need to consider the specific object being judged (the object of predication; e.g. *cake* vs. *blue cheese*). Thus, that work indicates we need to consider sentence-level information beyond the adjective.

In the present work, we extend the domain of inquiry beyond the lexical content of particular sentences. As illustrated in Figure 1, we test whether participants' judgments of faultlessness are influenced by (i) properties of the two *interlocutors* having the dialog, and/or by (ii) properties of the *participants themselves* (e.g. participants' own attitudes and opinions).³

We use the phenomenon of *expertise* to investigate these issues. In what follows, we present an experiment showing that people's judgment of faultless disagreement about claims in a particular domain (e.g. wine tasting) is sensitive to the expertise of the interlocutors making the claims, as well as to the experimental participants' own beliefs about the validity of expertise in that particular subjective domain. We test four domains: wine tasting, beer tasting, art and movies. Our results – reported below – show that faultless disagreement is sensitive both to the *context* (properties of the interlocutors) and to *extra-contextual* factors (the judge's own attitudes about the particular domain).

Thus, we conclude that, far from being a fixed property of individual predicates, judgments of faultless disagreement are modulated by factors entirely independent of the lexical content of the sentence. This problematizes the status of faultless disagreement as a desideratum for the semantics of individual predicates; any account of subjective predicates must account for the ways in which faultless disagreement is sensitive to extra-sentential context.

Expertise

It is often implicitly assumed that when faultless disagreement obtains in dialogs like (1a), both interlocutors

² Following semantic tradition, we often use the term 'subjective predicates' (or 'predicates of personal taste.') For the experiments reported here, these are synonyms with 'subjective adjectives.'

³ In an attempt to assess effects of participants' own opinions, in an unpublished study we used the same stimuli as the original Kaiser & Rudin (2020) study to test how participants' own opinions about particular foods modulate their faultless disagreement ratings. We replicated the same pattern we found in our original study: disagreements about divisive foods are rated more faultless than disagreements about high-consensus foods. This occurs both when a participant's own opinion diverges from the norm and when it matches the norm. We also found modulating effects of participants' own opinions. These results, first, corroborate the earlier findings and, second, provide evidence that the properties of the individuals judging the disagreement also play a role, but do not address the question of whether properties of the actual interlocutors themselves matter. We address that question in the experiment reported here.

are on an equal footing in the sense that both are ‘normal tasters’ and have “adequate grounds for their claims” (Koelbel 2009, see also Ninan 2014 on the Acquaintance Inference). Even if the interlocutors differ in terms of their expertise, it is often assumed that faultless disagreement still obtains (see e.g. Lasersohn 2005).




Figure 1. (Some of the) Layers that we propose modulate judgments of faultless disagreement

However, Furey (2017), Wolf (2016) and others suggest that the situation may not be so straightforward. Consider a context like (1c), where only one of the speakers is a wine connoisseur. Is the disagreement in (1c) as faultless as what we saw in (1a)? Or do we defer to the authority of the wine expert, and consequently view the non-expert’s opinion as ‘wrong’? Thus, by looking at contexts where the two speakers differ in terms of their expertise, we can tap into the question of whether judgments of faultless disagreement are sensitive to properties of the interlocutors.

(1c)

Sam (who has no wine expertise): This wine is tasty.
Wine connoisseur: No, this wine is not tasty.

Furthermore, it could be that someone’s judgment of the faultlessness of the disagreement in (1c) is modulated not only by the context in which the dialog occurs, but also by that person’s own attitudes towards wine experts – e.g., whether someone views wine experts’ views as being privileged or believes that both experts’ and non-experts’ opinions are equally valid.

Thus, by asking about participants’ own attitudes, we can tap into the broader question of whether judgment of faultlessness is sensitive to extra-contextual properties of the specific individual who is judging the disagreement.

The present work tests whether factors entirely independent of the lexical content of the sentence, namely the discourse context (properties of the interlocutors) and extra-contextual factors (the judge’s own attitudes about the domain) influence judgments of faultless disagreement.

Predictions regarding Expertise Differential Effects

We use the term ‘Expertise Differential Effects’ when investigating if judgments of faultless disagreement are modulated by whether one of the speakers is an expert. If faultless disagreement is simply a reflex of a semantic property of sentences, it should show no sensitivity to properties of the speakers of those sentences. However, if the interlocutors’ expertise levels modulate participants’ judgements of faultless disagreement – i.e., if we find Expertise Differential Effects – this would mean that faultless disagreement cannot be a ‘fixed’ property of lexical items and instead depends on contextual information about the interlocutors.

Predictions regarding Participant Attitude Effects

We use the term ‘Participant Attitude Effects’ to refer to potential effects of participants’ own attitudes. If judgments of faultless disagreement are modulated by the *judge’s* own attitudes about whether expert opinions in a domain are more valid than laypeople’s opinions, this would provide even more evidence that faultless disagreement does not follow directly from the lexical content of a sentence. More concretely, if Participant A holds the opinions of wine experts in high regard, but Participant B does not, will they attribute different levels of faultlessness to (1c)? If yes, this would indicate that faultless disagreement judgments depend not only on the lexical content of a sentence, and indeed not only on a sentence/context pair, but on the judge’s own attitudes about the subjective domain.

Potential differences between domains

If judgments of faultless disagreement are sensitive to the judge’s attitudes about the specific domain, the judgments may vary from domain to domain. There may be differences between ‘high-culture’ domains where expert status is established and valued (e.g. *wine*, *art*) vs. ‘popular culture’ domains (e.g. *beer*, *movies*) which, while having publicly-acknowledged experts, are viewed as more accessible to untrained consumers. We use the terms *more* vs. *less expertise-oriented* to refer to this distinction.

In our study, we first check whether people’s attitudes about the value of expert opinions in the four different domains actually reflect this proposed distinction. Then, we assess whether this distinction modulates perceived faultlessness: We may find stronger Expertise Differential Effects in more expertise-oriented domains (*wine*, *art*) than less expertise-oriented domains (*beer*, *movies*).

Experiment

We tested dialogs like ex.(3a,b), and manipulated (i) expertise differential (whether the dialog was between two experts or a non-expert and an expert) and (ii) which domain (see Table 1) the dialog concerns. (Only one adjective was displayed to participants for a given item; two example adjectives are given in ex.(3a,b) to highlight that

fact that we tested a mix of positive and negative subjective predicates.)

(3a) [two-expert version]

This is a conversation between two wine experts.

One wine expert says: This wine is {tasty/revolting}.

The other wine expert says: No, this wine is not {tasty/revolting}.

(3b)

[one-expert version]

This is a conversation between Andy (who has no expertise in wine) and a wine expert:

Andy says: This wine is {tasty/revolting}.

The wine expert says: No, this wine is not {tasty/revolting}.

Table 1. The four domains investigated in this experiment
(The art domain items were about paintings)

	<i>More expertise-oriented</i>	<i>Less expertise-oriented</i>
<i>Gustatory</i>	Wine	Beer
<i>Visual</i>	Art (paintings)	Movies

Method

Participants We report data for 84 adult native speakers of U.S. English who participated over the internet.

Materials and design Participants read dialogs like (3a,b) In one-expert conditions, the non-expert spoke first, to avoid creating an odd context where a non-expert ‘corrects’ an expert. We used 8 positive (e.g. *delicious*, *amazing*, *inspiring*) and 8 negative subjective adjectives (e.g. *revolting*, *boring*, *confusing*). The subjective adjectives we used all belong to the semantic class known as predicates of personal taste (PPTs). The study consisted of 16 targets and 24 fillers, presented using a Latin-Squared design.

Across participants, each gustatory-condition adjective was presented with both wine and beer, and each visual-condition adjective with both art and movies. This ensures that any differences between more and less expertise-oriented domains cannot stem from individual subjective adjectives. Each person saw each subjective adjective only once.

Participants saw both positive and negative adjectives in both one-expert and two-expert conditions. Whether the initial sentence used a positive or negative adjective had no clear effect, so we collapse polarity in the following discussion.

Procedure The study consists of two main task types, both implemented using the Qualtrics interface.

Faultlessness ratings During the main experiment, we elicited ratings of faultless disagreement. People used a six-

point scale to indicate whether they thought both speakers could be right or whether one is wrong (see Figure 2), where 1 means ‘One of the two people is wrong’ and 6 means ‘Both people can be right.’ Thus, higher ratings indicate higher faultlessness.

This is a conversation between two beer experts.
One beer expert says: This beer is delicious.
The other beer expert says: No, this beer is not delicious.

One of the two people is wrong. 1	2	3	4	5	Both people can be right 6
--------------------------------------	---	---	---	---	-------------------------------

This is a conversation between Steve (who has no expertise in art) and an art expert.
Steve says: This painting is amazing.
The art expert says: No, this painting is not amazing.

One of the two people is wrong. 1	2	3	4	5	Both people can be right 6
--------------------------------------	---	---	---	---	-------------------------------

Figure 2. Sample items with rating scale (top image: from the beer domain, bottom image: from the art domain)

Participant attitudes: After the main experiment, we collected information about participant attitudes regarding expert opinions. Participants were asked, for each domain, whether they view expert opinions as more valid than non-expert opinions or whether everyone’s opinion is equally valid (see Figure 3). This is step that allows us to check the validity of our assumption that the domains of wine and art are more expertise-oriented than beer and movies.

After the main experiment, we also asked participants if they self-identify as experts in any of the four domains. Very few did (average 4%), so we do not discuss these data further. (We do not have enough statistical power to compare self-identified experts to non-experts, given the low number of self-identified experts).

	Everyone's opinion is equally valid. The opinions of non-experts and experts are equally informative and important.	Experts are better equipped to make judgments than non-experts. I would put more trust in an expert's opinion.
Wine	☐	☐
Beer	☐	☐
Art	☐	☐
Movies	☐	☐

Figure 3. Post-experiment question about attitudes regarding expert opinions in the four domains

Results

We first report the outcomes of the initial sanity-check analysis checking whether the domains of wine and art are more expertise-oriented than beer and movies, and then turn to analyses pertaining to our main questions regarding what kinds of contextual and extra-contextual information impact ratings of faultless disagreement: (i) Expertise Differential Effects: Are judgments of faultless disagreement modulated by whether one or both of the speakers is an expert? (ii) Participant Attitude Effects: Are judgments of faultless disagreement modulated by the *judger's* own attitudes about whether expert opinions in a domain are more valid than laypeople's opinions?

Attitudes about the status of expert opinions in different domains

The results of the post-experiment questions about attitudes regarding expert opinions in the four domains we tested (see Figure 3) are shown in Figure 4. We find that in all domains, most participants feel that everyone's opinions are equally valid (darker grey), but crucially, participants are *more* likely to report that they view expert opinions as privileged (lighter grey bars) in the domains of *wine* and *art* than in the domains of *beer* and *movies*. A chi-squared analysis confirms that participants' preference to give more weight to expert opinions is stronger with *wine* and *art* than *beer* and *movies* ($\chi^2(1)=8.91$, $p<.003$). This confirms our assumptions about expertise-orientation.⁴

Participant Attitude Effects

Figure 5 shows faultless disagreement ratings as a function of domain and participant attitudes about (non)expert opinions in the domain. In expertise-oriented domains (*wine*, *art*), disagreements are rated more faultless (higher bars) by those who view everyone's opinions as equally valid, compared to those who view expert opinions as privileged (*wine*: $\beta = 1.027$, $SE = 0.328$, $t > |3.1|$, *art*: $\beta = 1.59$, $SE = 0.305$, $t > |5.2|$, mixed-effects regression (lmer) with R). Less expertise-oriented domains (*beer*, *movies*) show the same trend, but it does not reach significance (t 's $< |1.6|$). (Following convention, $t > |2|$ is treated as significant.)




Figure 4. Responses to post-experiment questions on participants' own attitudes about expert opinions




Figure 5. Faultless disagreement ratings as a function of domain and participants' attitudes about expertise

⁴ Here, we do not distinguish between different kinds of expertise (e.g. vintner/sommelier); we leave this for future work.

Expertise Differential Effects

Let us now take a closer look at how the one- vs. two-expert context manipulation interacts with participants' attitudes. Figure 5 shows that less expertise-oriented domains (*beer*, *movies*) show no clear effects of expertise differentials (one-

vs. two-experts) on faultless disagreement, confirming our expectation that expertise is not central in these domains.

However, in the *wine* domain, participants who hold experts' opinions in high regard show stronger expertise differential effects than those who view all opinions as equal (1-expert vs. 2-expert x attitude interaction, $\beta = 0.644$, $SE = 0.174$, $t > |3.7|$). Expertise-valuers rate one-expert disagreements as more faulty than two-expert disagreements ($\beta = -0.535$, $SE = 0.267$, $t > |2|$). This can be seen in the fourth pair of bars (from the left) in Figure 5.

There are no effects of expertise differentials in the *art* domain. Even art-expertise-valuers feel that all art-related disagreements are comparably faulty. Notably, this lack of an effect is due to all disagreements being judged relatively faulty, unlike the *movies* and *beer* domains where lack of expertise differential effects is due to all disagreements being judged relatively faultless.

We leave a fuller investigation of the asymmetry between the *wine* and *art* domains to future work. However, it is worth acknowledging the numerical and ratings-based nature of wine appreciation, especially compared to art. The asymmetry between wine and art may stem at least partially from the fact that, compared to the other domains, wine has become relatively more 'objectivized' due its focus on ratings/rankings (arguably more so than movies, for example, see e.g. Zahradka (2020) for related discussion).

Discussion

Our experimental data indicate that judgments of faultless disagreement can be sensitive to (i) individual differences in judges' (participants') views about expertise in a given domain (*participant attitude effects*), (ii) differences between domains, and (iii) properties of the interlocutors (*expertise differential effects*). These results show that faultless disagreement cannot be regarded as a reflex simply of the lexical content of the sentence.

Theories of subjective predicates have generally focused on how to semantically encode information about in whose judgment the predication obtains, and on the role played by perceptual experience in the semantics of these predicates. The factors we have observed here are orthogonal to those issues, and so to the extent that prior theories are attempting to explain faultless disagreement, they have done so inadequately. Our data indicate that faultless disagreement is sensitive both to the *context* (properties of the interlocutors) and *extra-contextual factors* (the judge's own attitudes about the domain). This adds at least two layers mediating between judgments of faultless disagreement and the semantics of subjective predicates (Figure 1), problematizing the argument that faultless disagreement is an empirical desideratum that must be explained directly by the semantics of these predicates.

Our results also have consequences for research that does not focus specifically on faultless disagreement *per se*, but that uses it as a 'tool' or diagnostic for detecting subjectivity. The notion of subjectivity has been argued to be relevant for a wide range of other linguistic phenomena,

and faultless disagreement judgements have been used as a diagnostic for the purposes of identifying subjective expressions and hypothesizing about the consequences of subjectivity for other phenomena, including syntactic patterns (e.g. adjective ordering, Scontras et al. 2017). Our finding that faultless disagreement judgments do not follow directly from the lexical content of a sentence complicates its use as a diagnostic for lexical-level subjectivity.

We suggest that our results make theories where there is nothing particularly special about the semantics of subjective predicates (e.g. Rudin & Beltrama 2019, Willer & Kennedy 2020) more appealing than they might have previously seemed, given that faultless disagreement does not appear to be a stable, consistent property of these predicates, as a theory that hard-codes the licensing of faultless disagreement into their semantics would predict.

Acknowledgements

This material is partly based upon work supported by the National Science Foundation under Grant No. 1749612 awarded to Elsi Kaiser.

References

- Furey, H. (2017). The paradox of gustatory taste. *Ergo*, 4, 481-527. <https://doi.org/10.3998/ergo.12405314.0004.017>
- Kaiser, E., & Rudin, D. (2020). When faultless disagreement is not so faultless. *Proceedings of the Ling Soc of America* <https://doi.org/10.3765/plsa.v5i1.4757>
- Koelbel, M. (2004). Faultless disagreement. *Proceedings of the Aristotelian Society*, 104, 153-173.
- Koelbel, M. (2009). The evidence for relativism. *Synthese*, 166, 375-395
- Laserson, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28, 643-686. <https://doi.org/10.1007/s10988-005-0596-x>
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- McNally, L., & Stojanovic, I. (2017). Aesthetic adjectives. In J. O. Young (Ed.), *Semantics of aesthetic judgment*. Oxford: Oxford University Press.
- Ninan, D. (2014). Taste predicates and the acquaintance inference. *Proceedings of SALT* <https://doi.org/10.3765/salt.v24i0.2413>
- Rudin, D., & Beltrama, A. (2019). Default agreement with subjective assertions. *Proceedings of SALT*. <https://doi.org/10.3765/salt.v29i0.4597>
- Scontras, G., Degen, J. & Goodman, N. (2017). Subjectivity predicts adjective ordering preferences. *Open Mind* 1, 53-65. http://doi.org/10.1162/opmi_a.00005
- Stephenson, T. (2007). *Towards a theory of subjective meaning*. Ph.D., Linguistics and philosophy, MIT.
- Willer, M. & Kennedy, C. (2020). Assertion, expression, experience. *Inquiry* <https://doi.org/10.1080/0020174X.2020.1850338>

- Wolf, L. (2016). Predicates of personal taste and the evidential step. In C. Meier & J. van Wijnbergen-Huitink (Eds.), *Subjective meaning*. Berlin: De Gruyter.
- Zahradka, P. (2020). Evaluating Food and Beverage Experience: Paradoxes of the Normativity. *Eidos* 4(3).
<https://doi.org/10.14394/eidos.jpc.2020.0034>