2024. *Proc Ling Soc Amer* 9(1). 5739. https://doi.org/10.3765/plsa.v9i1.5739.

The fact that these are opinions: Processing and acceptability patterns of subjective vs. objective information embedded under 'the fact that'

Haley Hsu & Elsi Kaiser*

Abstract. Language can convey both objective, fact-based information and subjective, opinion-based information. Previous research has focused on linguistic contexts that are associated with subjective information, such as information embedded under *find*. We aim to complement this existing work by exploring whether particular linguistic contexts are specifically associated with objective information. We report two psycholinguistic experiments testing the acceptability and processing of subjective and objective predicates embedded under *the fact that.* As a whole, the results suggest that subjective predicates embedded under *the fact that* are as acceptable and as easy-to-process as objective predicates, suggesting that this construction does not create a context that requires objective, fact-based information.

Keywords. objectivity; subjectivity; predicates of personal taste; nominal embedding; faultless disagreement; experimental semantics; discourse processing

1. Introduction. Language allow us to express both objective, fact-based information ('matters of fact') and more subjective, opinion-based information ('matters of opinion'). The notion of subjectivity has been investigated from various perspectives, including semantics and philosophy, by a variety of researchers. The present work seeks to complement existing work by exploring whether some linguistic expressions, in particular the embedding nominal construction *the fact that*, are geared towards objective information. Prior corpus studies of this phrase typically take as their starting point the intuition that *the fact that* relates to objectivity (e.g. Granath 2001; Guo 2022). However, this intuition is challenged by corpus examples showing that both subjective and objective information can seemingly be felicitously embedded under *the fact that* (see e.g. Granath 2001; Gentens 2019 for discussion). This is illustrated by the (constructed) examples in (1-2). Both an objective predicate (*triangular*) and a subjective predicate (*pretty*) sound felicitous in this context.

(1) The fact that the table is triangular is the reason it sold out quickly.
(2) The fact that the table is pretty is the reason it sold out quickly.

In this paper, we complement existing corpus-based findings from a psycholinguistic perspective. We report two experiments investigating objective vs. subjective information embedded under *the fact that*, compared to *the opinion that*, in order to assess whether there are differences in the naturalness and ease of processing subjective and objective information in these two contexts. Experiment 1 uses a naturalness rating task to yield offline judgment data, while Experiment 2 uses self-paced reading to tap into real-time incremental processing load. As a whole, the results suggest that a subjective predicate embedded under *the fact that* is as acceptable as and as easy to process as an objective predicate in the same context, whereas *the*

*Many thanks to the LSA100 audience for helpful feedback. Also, thanks to Deniz Rudin and Zuzanna Fuchs for their additional feedback and help throughout the various stages of this project. Authors: Haley Hsu, University of Southern California (haleyhsu@usc.edu) & Elsi Kaiser, University of Southern California (emkaiser@usc.edu).

opinion that shows a sharper distinction in its willingness to embed subjective vs. objective predicates. This paper contributes to our understanding of the lexical semantics of the expressions *the fact/opinion that* as well as to previous discussions about the asymmetry between objectivity and subjectivity.

2. Background.

2.1. LINGUISTIC CUES TO SUBJECTIVITY. Intuitively, there is a distinction between 'matters of fact' and 'matters of opinion.' However, how to capture this intuition in semantic theories is a complex question and subject to ongoing debate (see e.g. Lasersohn 2005; Glanzberg 2007; MacFarlane 2007; Stephenson 2007; Kennedy & Willer 2022; and many others), and a question that this paper does not seek to answer. Nevertheless, it is clear that some linguistic environments are especially conducive to subjective information rather than objective information. For example, *find* felicitously embeds subjective adjectives (ex.3a) but is infelicitous with objective adjectives (ex.3b, see e.g. Sæbø 2009; Kennedy & Willer 2022; and many others).

(3) a. I find this drink delicious. (from Kennedy & Willer 2022)
b. # I find this drink fermented. (from Kennedy & Willer 2022)

Another context where subjective and objective adjectives pattern differently is in disagreement contexts (e.g. Kölbel 2003). Disagreeing over a factual matter (ex.5) means that one of the speakers is in the wrong, but disagreeing over a subjective matter (ex.4) yields faultless disagreement: neither speaker is incorrect, or 'at fault.'

(4) Speaker 1: The cookies are tasty.
Speaker 2: No, the cookies are not tasty.

(5) Speaker 1: The cookies are gluten-free.
Speaker 2: No, the cookies are not gluten-free.

2.2. LINGUISTIC CUES TO OBJECTIVITY. Given the existence of expressions like *find* that seem to be associated with subjective information, one might wonder if there are other expressions that are associated with objective information. One place to look is in the domain of so-called certainty markers, which have been noted to front-load attitude meanings and convey the certainty of a truth statement by providing the author's stance, even if the statement following contains subjective information (e.g. Rubin & Kando 2006). One way to convey stance is with noun complement structures such as the fact that and the assumption that (e.g. Jiang and Hyland 2015). In the present paper, we focus on the expression the fact that and use an experimental approach to explore what kind of information can be embedded under it. In the rest of this section, we consider the intuition, from prior work, that the expression the fact that is linked to objectivity.

There is a large body of corpus-based work focusing on the question of what can be embedded under noun-complementizer phrases such as 'the fact that' using various American and British English corpora, including written and spoken text across many text genres (e.g. Schmid 2000; Granath 2001; Gentens 2019; Salkie 2017; Guo 2022). These researchers often start with the observation that there is an intuition that in contexts like *the fact that X*, X conveys factual information (Schmid 2000; Gentens 2019; Salkie 2017). Some researchers even start with a more specific and stricter intuition, namely that in these contexts, X should be objective and verifiable (Granath 2001; Guo 2022). While these intuitions turn out to be overly simplistic, they

highlight the intuition that expressions like *the fact that* are often perceived as being used to express factual, objective information. As noted by Kiparsky & Kiparsky (1970), the opposition between matters of fact and matters of opinion is also reflected in "the common question: *Is that a fact or is that just your opinion?*" (pg. 168–169).

2.3 CORPUS STUDIES. Previous corpus results suggest that the intuition linking *the fact that* to objectivity, though widespread, is an oversimplification (Granath 2001; Gentens 2019; Schmid 2000). For example, consider (6) and (7) from Granath (2001). These examples should not be felicitous given the claim that the complement following *the fact th*at should be objectively verifiable. In (6), whether something is 'woeful' is presumably not objectively verifiable, yet the sentence sounds fine. In (7), people can disagree about whether a particular album is powerful musical landmark or not, yet the sentence is felicitous.

(6) …the most obvious encumbrance on this picture is *the fact that* it is woefully late (Freiburg – LOB Corpus of British English, FLOB, G49 140, cited by Granath 2001)

(7) He ignores *the fact that* Geffen reputedly offered her $1million for it, and that her previous album Pretty On The Inside was a powerful musical landmark two years previously. (The Observer, 5 July 1998, p. 7, cited by Granath 2001)

Similar points are raised by Gentens (2019), who shows, using corpus examples, that *the fact that* can embed both subjective and objective information.

Thus, previous corpus studies suggest that *the fact that* is not limited to embedding objective information. However, one might still wonder whether more systematic investigation or finer-grained methods could detect a preference for *the fact that* to embed objective information. To test if objectivity plays a role in the processing of different types of information, we conducted two psycholinguistic studies. Before turning to these studies, let us first consider our hypotheses.

3. Hypotheses. So far we have seen corpus evidence showing that *the fact that* can occur with both subjective and objective information. However, the existence of corpus examples does not mean that a particular construction is easy to process; there exist many constructions that are relatively hard to process even though they are grammatical and felicitous. Thus, one possibility is that *the fact that* can embed either subjective or objective information, but that it nevertheless exhibits congruency effects, such that *the fact that* embedding objective information is easier to process and more natural than *the fact that* embedding subjective information. In this paper, we compare *the fact that* to *the opinion that*. As far as we know, there is no evidence suggesting that *the opinion that* can embed objective information; we regard this expression as an environment associated with subjective information.

If both *the fact that* and *the opinion that* show equally strong congruency effects – i.e. if *the fact that* 'demands' objective information as strongly as *the opinion that* 'demands' subjective information – then we should find congruent configurations involving these two expressions pattern alike. More specifically, if the congruency of the noun and the embedded information matter across the board, a configuration with [opinion+subjective information] should be rated more natural and be easier to process than [fact+subjective information], and [opinion+objective information] should be less natural and harder to process than [fact+objective information]. In other words, under this view we predict lower naturalness and a heavier processing load when the noun type and the information type are *incongruent* (compared to when they are *congruent*), and this should hold equally for both *the fact tha*t and *the opinion that*. We call this the **Overall Congruency Hypothesis**.

However, it may be congruency effects do not hold equally across the board. It could instead be that one information type or noun type exhibits stronger congruency effects than the other. We call this the **Asymmetrical Congruency Hypothesis.** The earlier corpus work suggests that subjective adjectives can more freely occur under both *fact* and *opinion* than can objective adjectives. Under this approach, then, we predict that a decrease in naturalness ratings or processing difficulties stemming from a lack of congruence is less severe for subjective information than for objective information. More specifically, the prediction is that the difference in naturalness and processing load between [opinion+subjective information] and [fact+subjective information] is *smaller* (and perhaps even entirely absent) compared to the difference between [opinion+objective information] and [fact+objective information].

4. Aims of this paper. There have not been, to the best of our knowledge, systematic psycholinguistic experiments testing the kinds of information (subjective vs. objective) that can be embedded under *the fact that*. We aim to complement previous corpus work by testing, in a carefully-controlled context, the naturalness and processing ease of subjective and objective information embedded under *the fact tha*t as compared to *the opinion that*. The structure of the paper is as follows. Section 5 reports Experiment 1, a naturalness rating study, which provides offline acceptability judgments. Section 6 reports Experiment 2, a self-paced reading study, which offers a more sensitive measure of real-time processing. Section 7 is the general discussion. As we will see, both studies converge to suggest that *the fact that* is a flexible expression that felicitously embeds both subjective and objective information, in contrast to *the opinion that* which prefers subjective information. Thus, our results support the Asymmetrical Congruency Hypothesis**.**

5. Experiment 1: Naturalness rating.

5.1. PARTICIPANTS. We report data for 52 native adult U.S.-English speakers recruited via Prolific (Palanab & Schitter 2018) who completed the study on Qualtrics (Provo, UT) and received $5. Participants were excluded if they were not born in the United States, did not speak English as their first language, or did not have normal/corrected-to-normal vision (7 participants were excluded for these reasons; information based on self-report). Participants were also excluded if they made more than two errors on five unambiguous practice items (6 participants were excluded based on this criterion).

5.2. MATERIALS AND DESIGN. We manipulated information type by adjective (objective vs. subjective) and the embedding noun (th*e fact that* vs. *the opinion that*), yielding a 2x2 within-subjects design (Figure 1). A sample item is shown in (10).

	Noun Type	
Adj. Type	**Fact**	**Opinion**
Objective	*The fact that + obj*	*The opinion that + obj*
Subjective	*The fact that + sub*	*The opinion that + sub*

Figure 1. Experiment 1 design

This experiment included 24 targets, presented using a Latin Square design. As shown in (10a-d), each target starts with the phrase *the fact/opinion that*… and contains an embedded noun that

is described with either an objective or subjective adjective and followed by a wrap-up section in the passive voice (e.g. …*is sometimes endorsed by the CEO of the company*).

(10) a. [fact+objective]
The fact that the holiday parties are mandatory is sometimes endorsed by the CEO of the company.
b. [fact+subjective]
The fact that the holiday parties are boring is sometimes endorsed by the CEO of the company.
c. [opinion+objective]
The opinion that the holiday parties are mandatory is sometimes endorsed by the CEO of the company
d. [opinion+subjective]
The opinion that the holiday parties are boring is sometimes endorsed by the CEO of the company.

The subjective adjectives were selected based on prior work, and the vast majority were predicates of personal taste (PPTs), e.g. *fun, tasty, stunning, hideous, gross*. The objective adjectives were mostly quality-, shape-, or process-based (e.g. *ceramic, triangular, handsewn, plastic, floral, braided*), in other words, adjective that refer to verifiable properties. We used 24 different objective and 24 different subjective adjectives, to avoid lexical repetition. All of the subjective adjectives and none of the objective adjectives allow for faultless disagreement, and all subjective adjectives and none of the objective adjectives can be felicitously embedded under *find* (see Section 2.1).

The experiment also contained 32 filler sentences for a total of 56 items. Out of the 32 filler sentences, 16 were designed to be fully natural, 6 were designed to receive low naturalness ratings, and 10 were designed to be rated in-between. Fillers included a variety of structures, including target-like sentences (with *the N that* constructions), garden paths, number attraction errors, and other structures.

5.3. PROCEDURE. Participants saw each target individually, on a separate screen, and rated its naturalness using a 7-point scale, with 1 being very unnatural and 7 being very natural. Participants competed the task at their own pace. Figure 2 shows an example screenshot.

Figure 2. Example target from Experiment 1

5.4. PREDICTIONS. Let us first consider the predictions of the Overall Congruency Hypothesis. If the congruency between the embedding noun and the adjective matters across the board, we expect that the 'congruent' [fact+objective] condition should be rated more natural than the 'incongruent' [opinion+objective] condition, and that the 'congruent' [opinion+subjective] condition should be rated more natural than the (supposedly) 'incongruent' [fact+subjective] condition. In other words, we expect lower naturalness ratings when the noun type and the adjective type are incongruent. Under this view, both of the incongruent conditions should show equal decreases in acceptability relative to the congruent conditions.

However, according to the if Asymmetrical Congruency Hypothesis, congruency effects may not hold equally across the board. In this case, we may find that one adjective type or noun type exhibits stronger congruency effects than the other. In fact, based on corpus work, one might hypothesize that subjective adjectives can more freely occur under both *fact* and *opinion* than can objective adjectives. In this case, the difference between the naturalness ratings for the congruent [opinion+subjective] and incongruent [fact+subjective] conditions is predicted to be *smaller* (and perhaps even absent), compared to the difference between the congruent [fact+objective] and the incongruent [opinion+objective] conditions.

5.5. Results.

5.5.1. Data processing and analysis. We fit a mixed effect model in R (R core team 2018) with fixed effects of noun type (contrast-coded, fact = 0.5, opinion = -0.5) and adjective type (contrast-coded, objective = 0.5, subjective = -0.5), as well as a noun type x adjective type interaction. We report our cumulative link mixed model using raw average ratings fitted with the Laplace approximation in R. Analyses based on both the raw rating scores and z-scores (calculated by targets only) showed similar patterns; we report statistics for the raw scores below. The highest-performing cumulative link mixed model contained random effects of correlated random intercepts for subjects and items (Baayen et al. 2008). Random effects started out fully crossed and fully specified with by-subject and by-item effects of adjective type, noun type, and the interaction. They were reduced (starting with by-item effects) via model comparison. Only random effects that contributed significantly to the model (p's < .05) were included (Baayen et al. 2008).

5.5.2. Naturalness rating results. Figure 3 shows the average naturalness ratings for each condition, on a scale of 1 (very unnatural) to 7 (very natural). Statistical analysis shows a significant effect of both predictors:[1] noun type (β = 0.427, SE = 0.111, z = 3.840, p < 0.001) and adjective type (β = -0.342, SE = 0.111, z = -3.089, p = 0.002). Crucially, the model also found an interaction between noun type and adjective type (β = 0.747, SE = 0.223, z = 3.360, p < 0.001). As can be seen in Figure 3, the [opinion+objective] condition yielded lower naturalness ratings than the other three conditions. In other words, participants found *the opinion that* embedding objective material to be the least natural of the four conditions.

Additional planned comparisons confirm that the [fact+subjective] and [opinion+subjective] conditions do not differ significantly from each other (β = 0.0433, SE = 0.159, z = 0.272, p = 0.785), while the [fact+objective] condition is rated significantly more natural than the [opinion+objective] condition (β = 0.754, SE = 0.159, z = 4.736, p < 0.001).

These results support the Asymmetrical Congruency Hypothesis, as we find asymmetrical congruency patterns: while objective information is rated significantly less natural when embedded under *the opinion that* than under *the fact that*, subjective information is rated equally natural with *opinion* and *fact*.

Given these results, one might be tempted to conclude that *the fact that* is not sensitive at all to the subjective vs. objective distinction, and that subjective information can equally felicitously occur under both *the fact that* and *the opinion that*. However, Experiment 1 leaves open the possibility that we are dealing with a weaker effect, perhaps one that cannot be detected with offline naturalness rating data. To address this question, in Experiment 2 we use an online method, self-paced reading, to get more fine-grained information about real-time processing. In

[1] Clmm model equation: Rating ~ nountype * adjectivetype + (1|participant) + (1|item).

self-paced reading, participants see sentences word-by-word, and reading times for each word are analyzed. Thus, this method can reveal transient differences in processing load (as reflected in reading times), which may not be detectable in end-of-sentence acceptability judgments.

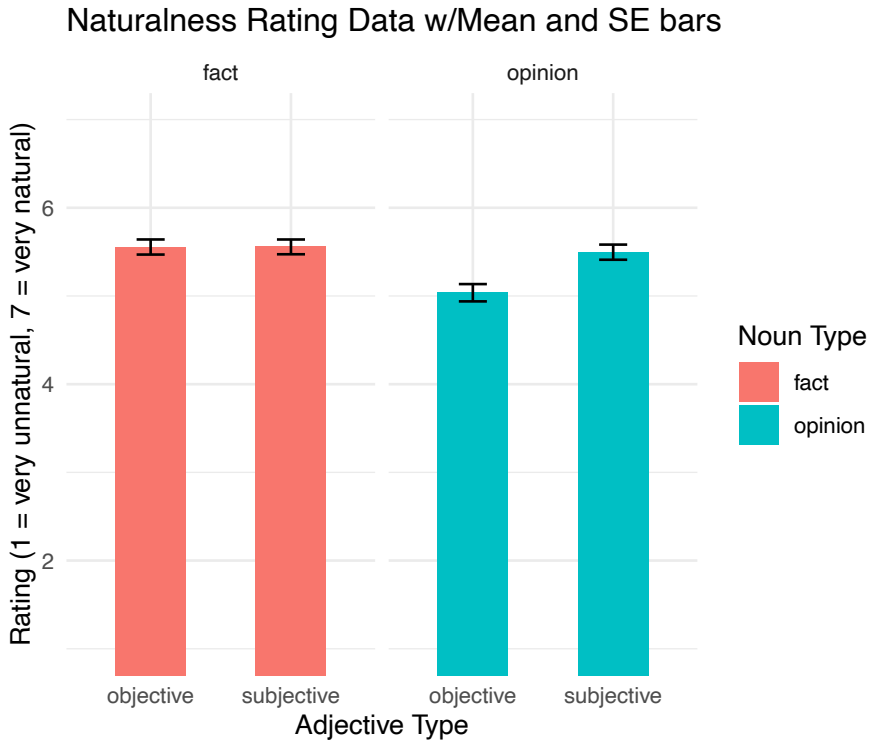

Figure 3. Mean values for naturalness ratings (1 = very unnatural, 7 = very natural), split by noun type (*fact* vs. *opinion*) and adjective type (objective vs. subjective). Error bars show +/- 1 SE

6. Experiment 2: Self-paced reading.

6.1. PARTICIPANTS. We report data for 61 adult native U.S. Native English speakers recruited via Prolific who completed the study on PCIbex (Zehr & Schwarz 2018) and received $6. Participants were instructed to use full-screen mode or maximized window setting. No participant from Experiment 1 took part in Experiment 2. Participants were excluded if they were not born in the United States, did not speak English as their first language, or did not have normal/corrected-to-normal vision (based on self-report, 13 participants excluded). Participants were also excluded if they made more than one error on four unambiguous practice items or scored below 80% accuracy on comprehension questions (8 participants excluded). Exclusion criteria were determined prior to data analysis.

6.2. MATERIALS AND DESIGN. The self-paced reading study had the same 2x2 design and used the same targets as Experiment 1. Experiment 2 also contained 32 filler sentences for a total of 56 items. Fillers again included a variety of structures, including target-like sentences (with the N that constructions), garden paths, and other structures. Sentences were presented word-by-word. The key or critical region of interest is the subjective/objective adjective (see Figure 4). After reading *the opinion that* or *the fact that*, how do participants react upon encountering a subjective or objective (congruent or incongruent) adjective?

Beginning of sentence		Critical region	Spillover region			Sentence end
Certainty Marker	**Embedded Sent.**	**Crit. Word**	**Aux.**	**Adv**	**Main Verb**	**Sentence End**
The fact that	*the holiday parties are*	*mandatory*	*is*	*sometimes*	*endorsed by*	*the CEO of the company.*
The fact that	*the holiday parties are*	*boring*	*is*	*sometimes*	*endorsed by*	*the CEO of the company.*
The opinion that	*the holiday parties are*	*boring*	*is*	*sometimes*	*endorsed by*	*the CEO of the company.*
The opinion that	*the holiday parties are*	*mandatory*	*is*	*sometimes*	*endorsed by*	*the CEO of the company.*

Figure 4. All four conditions of a sample target sentence divided into the various regions

6.3. PROCEDURE. The participants' task was to read each sentence one word at a time (pressing the spacebar to move onto the next word, example in Figure 5) and then to answer a comprehension question about the sentence they just read. Each sentence was displayed on the screen individually and designed such that the critical word and spillover region were not located at a line break.

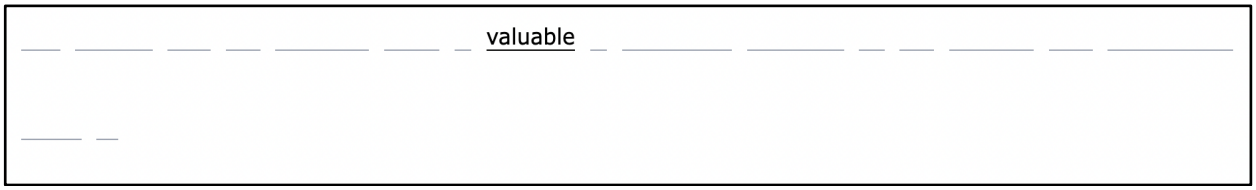

Figure 5. A screenshot from Exp.2 on PCIbex that shows a target sentence

Participants were instructed to have their index fingers on the 'F' ('Yes') and 'J' ('No') keys and their dominant thumb on the spacebar.

6.4. PREDICTIONS. Recall that Experiment 1 found that subjective adjectives are rated equally natural under *the fact that* and *the opinion that*, while objective adjectives are rated more natural under *fact* than *opinion*, supporting the Asymmetrical Congruency Hypothesis. Given this, Experiment 2 allows us to test whether the congruency effect with objective adjectives are replicated in reading times. In other words, do we find longer reading times on the adjective or following spillover regions in the [opinion+objective] condition than in the [fact+objective] condition?

Furthermore, it will be important to see whether reading times, being a more sensitive measure, allow us detect any congruency effects with subjective adjectives as well. In other words, are reading times on the adjective or following spillover regions in the [opinion+subjective] and in the [fact+subjective] condition the same (echoing the naturalness ratings), or can we detect a slowdown in the [fact+subjective] condition? If we find RT differences not only between the [opinion+objective] and [fact+objective] condition but also between the [opinion +subjective] and the [fact+subjective] condition, this would suggest that congruency effects exist in all four configurations, even though they may be stronger in some than in others.

6.5. RESULTS.

6.5.1. DATA PROCESSING AND ANALYSIS. Reaction times that were shorter than 100 ms or longer than 2500 ms were excluded. Additionally, reaction times that were more than 3 SD away from the mean were removed (349 datapoints). This resulted in removal of 5.88% of the original data. We then fitted linear mixed-effects models in R to each word position to analyze the log-transformed RTs with fixed effects of noun type (contrast-coded, fact = 0.5, opinion = -0.5) and

adjective type (contrast-coded, objective = 0.5, subjective = -0.5), as well as a noun type x adjective type interaction. Analyses based on both the raw reaction times and log-transformed reaction times showed similar patterns; we report statistics for the log-transformed reaction times below. The highest-performing linear mixed-effects models contained random effects of correlated random intercepts for subjects and items (Baayen et al. 2008). Random effects started out fully crossed and fully specified with by-subject and by-item effects of adjective type, noun type, and the interaction. They were reduced (starting with by-item effects) via model comparison. Only random effects that contributed significantly to the model (p's < .05) were included (Baayen et al. 2008).

Figure 6 provides an example of a target sentence divided into word positions. The critical region is the adjective (e.g. *braided/irritating*, region labelled Adj Type in Fig.6), and the spillover regions are the first four words immediately following this region (e.g. *is sometimes endorsed by,* auxiliary verb, adverb, main verb and preposition).

Sentence beginning						**Critical region**	**Spillover region**			
D	Noun Type	C	D	N	Aux	Adj Type (ADJ)	Aux	ADV	MV	P
1	2	3	4	5	6	7	8	9	10	11
The	fact/opinion	that	the	wigs	are	braided/irritating	is	sometimes	endorsed	by…

Figure 6. An example showing the word regions

6.5.2. SELF-PACED READING RESULTS. The reading time results are in Figure 7, which shows the mean RT in milliseconds for each condition at each word position. The critical adjective is shown in the box, and the spillover regions are indicated by the dotted box.

For the pre-critical regions, there is a main effect of noun type at **word positions 2 and 3** (*fact/opinion that*) (p's < 0.001).[2] This is expected for reasons of both frequency and length. First, the phrase *the fact that* is more frequent than *the opinion that*. For example, in the Corpus of Contemporary American English (Davies 2008), searching for *the fact that* yields 114,686 occurrences, while searching for *the opinion that* yields only 1,123 occurrences. Second, the word *opinion* is longer than the word *fact*. Thus, slower RTs for *opinion* conditions are expected. Other pre-critical regions showed no effects of noun type. There were no significant effects of adjective type or interactions between noun type and adjective type at any pre-critical regions.

At the **critical region** (the adjective), there are significant effects of noun type (β = -0.0436, SE = 0.0137, t = -3.190, p = 0.0015)[3] and adjective type (β = 0.0324, SE = 0.0137, t = -2.368, p = 0.018). RTs are longer at the critical adjective in the *opinion* conditions than in the *fact* conditions, and longer for objective adjectives than for subjective adjectives. However, the noun type x adjective type interaction is not significant (β = -0.0051, SE = 0.0273, t = -0.185, p = 0.853).

In addition to the critical adjective, we analyzed reading times for the following four words (four spillover regions) as well. The **first spillover region**[4] shows a main effect of adjective type (subjective adjectives are read faster, β = 0.0296, SE = 0.0119, t = 2.482, p = 0.0132), no main effect of noun type (β = -0.0194, SE = 0.0119, t = -1.625, p = 0.1044) and no interaction (β =

[2] Lmer model for word position 2: RTlog ~ nountype * adjectivetype + (1|participant) + (1|item). Lmer model for word position 3: RTlog ~ nountype * adjectivetype + (1|participant) + (1|item).
[3] Lmer model for critical region (adjective): RTlog ~ nountype * adjectivetype + (1|participant) + (1|item).
[4] Lmer model for the first spillover region: RTlog ~ nountype * adjectivetype + (1|participant) + (1|item).

-0.0351, *SE* = 0.0239, *t* = -1.471, *p* = 0.1417). The **second spillover region** shows no main effect of noun type (*β* = -0.0035, *SE* = 0.0117, *t* = -0.300, *p* = 0.765),[5] no effect adjective type (*β* = 0.0057, *SE* = 0.0117, *t* = 0.489, *p* = 0.625), but reveals a significant interaction (*β* = -0.0582, *SE* = 0.0234, *t* = -2.492, *p* = 0.0128). We discuss this pattern more below. The third and fourth spillover regions (the main verb and the preposition) show no significant main effects or interactions.

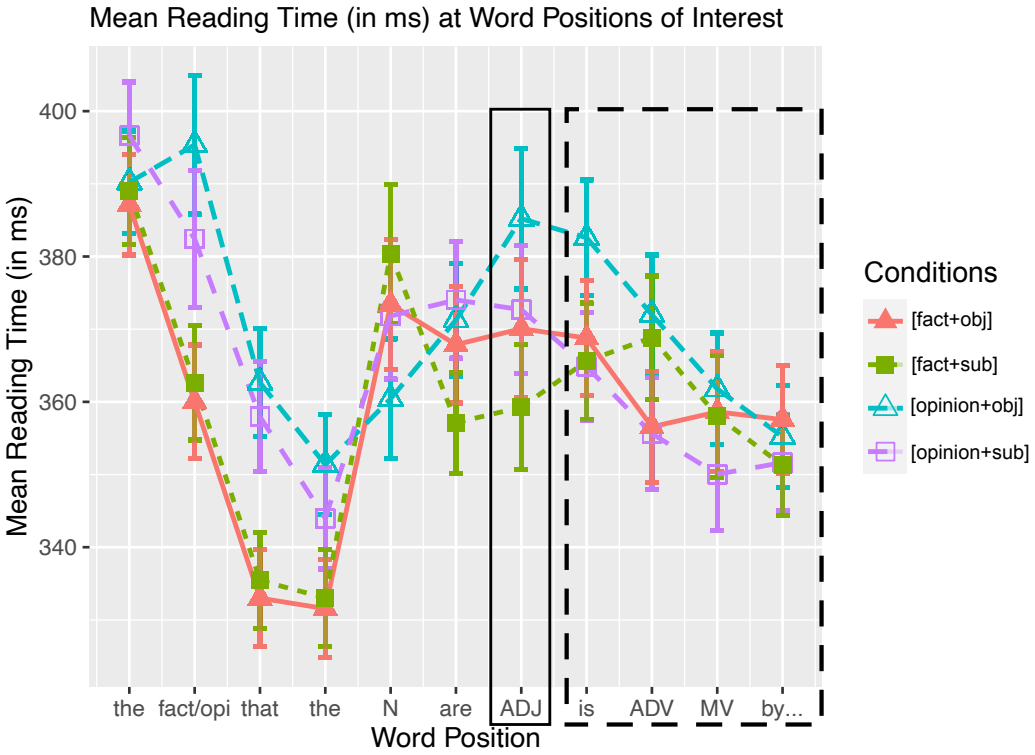

Figure 7. Mean reading times (raw RTs in ms) by condition and region in the target sentences. The critical word is denoted by a solid box. The four spillover regions are inside the dotted box. obj = objective, sub = subjective.

6.5.3. PLANNED COMPARISON. Although we found an interaction between noun type and adjective type only at the second spillover region, we conducted planned comparisons at the critical region (adjective) and the first spillover word, as well as the second spillover word. This is because paired comparisons between the two subjective conditions ([fact+subjective] and [opinion+subjective]) as well as between the two objective conditions ([fact+objective] and [opinion+objective]) are a core part of our predictions and experimental design.

At the **critical adjective**, the 'congruent' [fact+objective] condition was read faster than the 'incongruent' [opinion+objective] condition (the red filled triangle is below the open blue triangle; *β* = -0.048, *SE* = 0.0200, *t* = -2.406, *p* = 0.0164). Furthermore, and perhaps surprisingly, when it comes to the subjective adjectives, we find that the 'incongruent' [fact+subjective] condition (filled green square) was read *faster* than the 'congruent' [opinion+subjective] condition (open purple square) (*β* = -0.0364, *SE* = 0.0183, *t* = -1.99, *p* = 0.0471): The finding that conditions where a subjective adjective occurs after *the fact that* are read *faster* than conditions where a subjective adjective occurs after *the opinion that* indicates that *the fact that* easily

[5] Lmer model for the second spillover region: RTlog ~ nountype * adjectivetype + (1|participant) + (1|item).

embeds subjective information. Thus, although we see the expected congruency effect with objective adjectives, we see an *incongruency effect* with subjective adjectives. If anything, this aligns better with the Asymmetrical Congruency Hypothesis.

At **spillover 1**, the [fact+subjective] and [opinion+subjective] conditions (squares) are read equally fast (β = -0.0017, SE = 0.0169, t = -0.099, p = 0.921). Echoing what we already saw at the critical region, the 'congruent' [fact+objective] (red filled triangle) is again read faster than the 'incongruent' [opinion+objective] (blue open triangle) (β = -0.0368, SE = 0.0168, t = -2.194, p = 0.0286).

At **spillover 2,** conditions with subjective adjectives again do not differ statistically from each other regardless of whether the adjective is embedded under *the fact that* or *the opinion that* ([fact+subjective] vs. [opinion+subjective]: β = 0.0247, SE = 0.017, t = 1.486, p = 0.138). When it comes to objective adjectives, however, we again see an incongruence slowdown: the 'congruent' [fact+objective] was read faster than the 'incongruent' [opinion+objective] (β = -0.032, SE = 0.016, t = -1.979, p = 0.0483).

In sum, sentences with subjective adjectives show no consistent evidence of RT slowdowns when the subjective adjectives are embedded under *the fact that* compared to *the opinion that,* whereas sentences with objective adjectives repeatedly elicit slower RTs in incongruent conditions, i.e., when the objective adjectives are embedded under *the opinion that* compared to *the fact that*. Thus, these results support the Asymmetrical Congruency Hypothesis, as we find asymmetrical congruency patterns: while the processing behaviors with objective information significantly differ when embedded under *the opinion that* vs. under *the fact that*, the processing of subjective information behaves similarly when embedded under *opinion* vs. *fact*.

7. General discussion. The two studies reported in this paper investigated the naturalness and processing ease of embedding subjective and objective adjectives under two kinds of embedding nominals: *the fact that* and *the opinion that*.

Naturalness rating data from Experiment 1 reveals a congruency effect with objective adjectives – objective information is rated significantly less natural when embedded under *the opinion that* than under *the fact that* – but no such congruency effect with subjective adjectives, which are rated equally natural with *opinion* and *fact*. This supports the Asymmetrical Congruency Hypothesis, according to which objective and subjective information differ in how 'flexibly' they can embedded under *fact* and *opinion.*

Experiment 2 used a self-paced reading task to see whether a more sensitive method would reveal congruency effects with subjective adjectives as well. The results echo Experiment 1: objective adjectives are read more slowly when embedded under *the opinion that* compared to under *the fact that*, but subjective adjectives show no slowdowns when embedded under *the fact that* compared to under *the opinion that*.

As a whole, these results favor the Asymmetrical Congruency Hypothesis, i.e., that objective adjectives exhibit stronger congruency effects than subjective adjectives. Our results thus suggest that *the fact that* is quite flexible in what kinds of information it can embed: both subjective and objective adjectives are rated natural when embedded under *the fact that,* and subjective adjectives exhibit no reading time slowdowns when embedded under *the fact that*. We find no evidence to support the claim that *the fact that* requires objective information, nor provides a reliable signal that upcoming information is objective.

Our psycholinguistic results thus align well with prior corpus work showing that *the fact that* can embed subjective information as well (e.g. Granath 2001; Guo 2022), and provide the first experimental investigation of these issues (to the best of our knowledge).

The finding the expression *the fact that* does not signal that the upcoming information is objective suggests that it may be semantically bleached, as proposed by Schmid (2000). Examining the cognitive processing of these types of constructions more in future work might offer insight into whether or not semantic bleaching is occurring for this high-frequency phrase (Schmid 2000). Another open question for future work has to do with other linguistic expressions that, at least at first glance, may seem to be signaling objectivity (e.g. *Objectively,…*, or *Verifiably,…*). Because the present paper focuses only on the minimal pair *the fact that / the opinion that*, other expressions merit closer investigation in future work.

Lastly, it is worth considering why the phrase *the fact that* is used, given the findings that it does not reliably signal that the upcoming information is objective. One corpus study proposes that this phrase provides a useful way to package phrasal content into nominal form and thus acts as a framing device (Granath 2001). However, it could also be that this expression is a rhetorical device that writers use to make statements more convincing. Corpus work by Jiang & Hyland (2015) found that *the fact that* is frequent in academic research articles across many disciplines, especially in marketing, as well as the natural and social sciences. Experimental investigations of whether *the fact that* impacts readers' perception of how convincing a certain claim can be is an important question for future work.

References

Baayen, Rolf H., Douglas J. Davidson & Douglas M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59(4). 390–412. https://doi.org/10.1016/j.jml.2007.12.005.

Davies, Mark. 2008. The Corpus of Contemporary American English (COCA). Online corpus. https://www.english-corpora.org/coca/.

Gentens, Caroline. 2019. The diachrony of *the fact that*-clauses. *English Studies*100(2). 220–239. https://doi.org/10.1080/0013838X.2019.1566853.

Glanzberg, Michael. 2007. Context, content, and relativism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 136(1). 1–29. http://www.jstor.org/stable/40208765.

Granath, Solveig. 2001. Is that a fact? A corpus study of the syntax and semantics of *the fact that*. In Paul Rayson, Andrew Wilson, Tony McEnery, Andrew Hardie & Shereen Khoja (eds.). *Proceedings of the Corpus Linguistics 2001 Conference*, 234–244. Lancaster: University Centre for Computer Corpus Research on Language.

Guo, Man. 2022. A corpus-based study of fact projection in English. *International Journal of English Linguistics* 12(2). 36–46. https://doi.org/10.5539/ijel.v12n2p36.

Kennedy, Christopher & Malte Willer. 2022. Familiarity inferences, subjective attitudes and counterstance contingency: Towards a pragmatic theory of subjective meaning. *Linguistics and Philosophy* 45. 1395–1445. https://doi.org/10.1007/s10988-022-09358-x.

Kiparsky, Paul & Carol Kiparsky. 1970. FACT. In Manfred Bierwisch & Karl Erich Heidolph (eds.), *Progress in linguistics: A collection of papers*, 143–173. Berlin: De Gruyter Mouton. https://doi.org/10.1515/9783111350219.143.

Kölbel, Max. 2003. Faultless disagreement. *Proceedings of the Aristotelian Society* 104. 53–73.

Lasersohn, Peter. 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, John. 2007. Relativism and disagreement. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 132(1). 17–31. http://www.jstor.org/stable/25471842.
Palanab, Stefan & Christian Schitter. 2018. Prolific.ac – A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17. 22–27. https://doi.org/10.1016/j.jbef.2017.12.004.
R Core Team. 2019. R: A language and environment for statistical computing [Computer software manual]. Vienna. https://www.R-project.org.
Rubin, Victoria, Elizabeth Liddy & Noriko Kando (eds.). 2006. *Computing attitude and affect in text: Theory and applications* (The Information Retrieval Series 20). Dordrecht: Springer.
Sæbø, Kjell J. 2009. Judgment ascriptions. *Linguistics and Philosophy* 32(4). 327–352. http://www.jstor.org/stable/27783995.
Salkie, Raphael. 2017. Legal phraseology in contrast: *The fact that* and its German counterparts. In S. Goźdź-Roszkowski, & G. Pontrandolfo (Eds.). Phraseology in legal and institutional settings: A corpus-based interdisciplinary perspective, 126–142. Londonn: Routledge.
Schmid, Hans-Jörg. 2000. *English abstract nouns as conceptual shells: From corpus to cognition* (Topics in English Linguistics 34). Berlin: Mouton de Gruyter.
Stephenson, Tamina. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30. 487–525. https://doi.org/10.1007/s10988-008-9023-4.
Wiebe, Janyce & Ellen Riloff. 2005. Creating subjective and objective sentence classifiers from unannotated texts. In Alexander Gelbukh (ed.), *Computational linguistics and intelligent text processing* (Lecture Notes in Computer Science 3406). Berlin: Springer. https://doi.org/10.1007/978-3-540-30586-6_53.