

Knowledge and loose talk

Alexander Dinges

alexander.dinges@fau.de

forthcoming in *Skeptical Invariantism Reconsidered*, ed. by C. Kyriacou and K. Wallbridge,
Routledge

Abstract. Skeptical invariantists maintain that the expression “knows” invariably expresses an epistemically extremely demanding relation. This leads to an immediate challenge. The knowledge relation will hardly if ever be satisfied. Consequently, we can rarely if ever apply “knows” truly. The present paper assesses a prominent strategy for skeptical invariantists to respond to this challenge, which appeals to loose talk. Based on recent developments in the theory of loose talk, I argue that such appeals to loose talk fail. I go on to present a closely related response strategy, which appeals to conversational exculpation, a phenomenon recently studied by Hoek (2018, 2019). A promising account of knowledge ascriptions results, which combines the virtues of invariantism and elaborate versions of contextualism, positing a unique knowledge relation while still making precise predictions on when knowledge ascriptions ordinarily count as appropriate.

10,040 words

1 Introduction

Skeptical invariantists maintain that the expression “knows” invariably expresses an epistemically extremely demanding relation. This leads to an immediate challenge. The knowledge relation will hardly if ever be satisfied. Consequently, we can rarely if ever apply “knows” truly. This seems troublesome because we make knowledge ascriptions all the time and because, intuitively, this is perfectly all right.¹

A prominent response strategy on behalf of skeptical invariantism appeals to loose talk (see e.g. Conee, 2005: 52–53; Davis, 2007; Fantl and McGrath, 2009: 194; Bonjour, 2010: 73; Dinges, 2016; Kyriacou, 2017, 2019a). Consider the following sentences.

- (1) Hannah arrived at 3pm.
- (2) Hannah is 6 foot 2.

¹ Davis (2007) and Kyriacou (2017, 2018), for instance, defend *moderate* skeptical invariantism, where knowledge is not so rare after all. It is still rarer than ordinary knowledge ascriptions, so the worry above still arises and my discussion of this worry remains relevant.

(3) The table is flat.

Utterances of (1) and (2) can be proper even if Hannah didn't arrive at *precisely* 3pm and isn't *exactly* 6 foot 2, and utterances of (3) can be proper even if the table isn't *perfectly* flat. Skeptical invariantists can argue that knowledge ascriptions work just like that. We accept knowledge ascriptions when we don't *strictly* know things for just the same reasons for which we accept e.g. ascriptions of flatness to objects that are less than *perfectly* flat. We speak loosely in each case.

This paper argues that such appeals to loose talk on behalf of skeptical invariantism fail. It also presents a closely related, more promising response to the challenge above. Skeptical invariantism comes out as an account of knowledge ascriptions that combines the virtues of invariantism and contextualism. It posits a stable, context-independent knowledge relation that can unify the theoretical roles that knowledge plays in philosophy, psychology, legal theory, ethics, etc. At the same time, it accommodates the flexible, relatively weak demands on knowledge ascriptions in ordinary discourse. It achieves all that in a semantically and pragmatically precise, yet parsimonious framework that employs only familiar mechanisms that we have to posit for independent purposes anyway.

For starters, I introduce some formalism and terminology (§2). Then I look at some striking parallels between loose talk and knowledge ascriptions. They motivate the appeal to loose talk on behalf of skeptical invariantism, and they indicate some desiderata for any plausible account of knowledge ascriptions (§3). I then turn to familiar theories of loose talk and their application to the case of "knows." Numerous accounts of loose talk have been suggested in the literature. On many of these accounts, loose talk is a semantic phenomenon, where the semantic value of e.g. "3pm" varies with the context of utterance, denoting smaller or larger ranges of points in time (e.g. Krifka, 2007; DeRose, 2012: 714–717; Solt, 2014). If any of these accounts is correct, the appeal to loose talk on behalf of skeptical invariantism is self-refuting. After all, skeptical invariantists are invariantists and hence they cannot countenance semantic variability in "knows."² Meanwhile, I show that familiar pragmatic accounts of loose talk won't help the skeptical invariantist either, because they make false predictions when applied to "knows." The upshot is that if any extant account of loose talk is on the right track, loose talk is of no avail to skeptical invariantists (§4). Hoek (2018, 2019) specifically explains loose talk as a type of what he calls conversational exculpation. Following the results from the previous section, knowledge ascriptions don't instantiate this type of conversational exculpation, but they may still instantiate the general phenomenon. Indeed, I argue that such an approach looks

² On Carter's (2019) dynamic semantic account of loose talk, expressions like "3pm" have a fixed *static* semantic value and a variable *dynamic* semantic value. I assume that skeptical invariantists want to hold both these semantic values fixed so that Carter's account is unavailable to them. See also Dinges (MS) for independent concerns with Carter's account.

promising if we adopt independently motivated assumptions about the dynamics of pragmatic pre-suppositions familiar from Blome-Tillmann's (2014) works on epistemic contextualism (§5).

2 A little bit of formalism

In presenting the subsequent material, it will be useful to work in a regimented system, and I will assume a simple intensional semantics for this purpose. Thus, I will assume that sentences express propositions represented as sets of possible worlds, which result from composing semantic values of atomic expressions. Consider, for instance, the following semantic clauses for "knows" derived from familiar attempts to analyze knowledge.

- $$\llbracket \text{knows} \rrbracket = \{ \langle S, p, w \rangle \mid \text{At } w,$$
- i. S has a true belief with the content p, and
 - ii. This belief is justified/safe/reliably formed}

In the assumed semantic framework, these semantic clauses determine the propositions expressed by a sentence of the form "S knows that p" as follows.

$$\llbracket \text{S knows that p} \rrbracket = \{ w \mid \langle \llbracket \text{S} \rrbracket, \llbracket \text{that p} \rrbracket, w \rangle \in \llbracket \text{knows} \rrbracket \}$$

In the given system, we can clarify skeptical invariantism. This view is compatible in spirit with any of the semantic clauses just proposed for "knows." We just have to interpret the respective second condition as being invariably maximally demanding on the respective dimension. In what follows, I will understand skeptical invariantism in this way, i.e., as the view that the semantic value of "knows" invariably entails the maximum of some pertinent epistemic quantity, like justification, safety, reliability, etc.

It will be useful in what follows to have a term that, by definition, expresses the demanding relation that "knows" expresses according to skeptical invariantism, and I will use "knows_s" for this purpose. Here are some possible semantic clauses for "knows_s."

- $$\llbracket \text{knows}_s \rrbracket = \{ \langle S, p, w \rangle \mid \text{At } w,$$
- i. S has a true belief with the content p, and
 - ii. This belief is perfectly justified/safe across every possible world/100% reliably formed}

It will also be useful to have a term that stipulatively denotes the kind of relation that normally holds between a subject and a proposition when a knowledge ascription seems appropriate. This relation is much less demanding than knowledges_s, and it corresponds to the semantic clauses outlined for "knows" before. Let's use "knows_o" for this purpose.

- $$\llbracket \text{knows}_o \rrbracket = \{ \langle S, p, w \rangle \mid \text{At } w,$$
- i. S has a true belief with the content p, and
 - ii. This belief is justified/safe/reliably formed}

Properly speaking, there isn't just one such relation, because the precise requirements for proper knowledge attribution vary from context to context. But for present purposes, we can simplify matters a bit by ignoring this variation. Later on, I will consider how the context-sensitivity in the interpretation of "knows" can be accounted for on my proposal.

To fix some further terminology, I will somewhat tendentiously refer to the proposition expressed by a knowledge ascription when "knows" is read as "knows_s" as the *strict reading* of that knowledge ascription, and I will refer to the proposition expressed by a knowledge ascription when "knows" is read as "knows_o" as the *loose reading* of that knowledge ascription.

Similar distinctions can be drawn in obvious ways for sentences (1) to (3) above. We can take e.g. "3pm" to denote a precise point in time to arrive at the *strict reading* of (1), while we get the *loose reading* when we take "3pm" to denote, say, a range of points in time surrounding this precise point in time.

3 Parallels

This section presents a number of striking parallels between, on the one hand, cases of loose talk as in (1) to (3) and knowledge ascriptions on the other.

Semantic contradictions. Even in contexts where the proper use of sentences like (1) to (3) would require only the truth of their loose reading, it is normally infelicitous to conjoin them with a sentence that contradicts only the strict reading. Thus, sentences of the following form sound odd (e.g. Lasnik, 1999: 535; Lauer, 2012: 390; Solt, 2014: 521; Hoek, 2019: 171; Carter, 2019: 8; Moss, 2019: 263).

(4) # Hannah arrived at 3pm, but she didn't arrive before 3:02.

Exactly parallel phenomena are familiar from the debate on so-called concessive knowledge attributions. Here too it seems that even in contexts where the proper use of a knowledge ascription would merely require the truth of its loose reading, it is infelicitous to directly contradict the strict reading.

(5) # Hannah knows the bank will be open tomorrow, but she cannot rule out that it burns down over night.

Embeddings. The loose reading of (1) to (3) typically embeds under various types of logical operations (Davis, 2007: 408–409; Carter, 2019: 5–6; Hoek, 2019: 172).

- (6) a. Hannah isn't 6 foot 2.
b. If Hannah is 6 foot 2, then she needs the larger suit.
c. Is Hannah 6 foot 2?

For instance, the more or less trivial fact that Hannah isn't *exactly* 6 foot 2 normally doesn't suffice for the proper use of (6.a). Intuitively, this requires that Hannah isn't even *around* 6 foot 2.

The same goes for knowledge ascriptions, where the loose reading also embeds (Kindermann, 2019).

- (7) a. Hannah doesn't know the bank will be open.
b. If Hanna knows the bank will be open, we can go home now.
c. Does Hannah know the bank will open?

For instance, a proper use of (7.a) seems to require that Hannah fails to know_o that the bank will be open. The mere failure to know_s doesn't suffice.

Pedantry. While challenging only the strict reading of sentences like (1) to (3) normally sounds pedantic, it still has the effect that the proper use of these sentences subsequently requires the truth of the strict reading (e.g. Lewis, 1979: 351–354; Grice, 1989: 45; Klecha, 2018: 92–93).

- (8) *A:* Hannah will arrive at 3pm.
B: In fact, she will arrive at 2:58pm.
A: True. Still, she will arrive at 3pm.

In the case of knowledge ascriptions, so-called salient alternative effects bring out parallel phenomena.

- (9) *A:* Hannah knows the bank will be open.
B: But she cannot rule out that a meteor destroys it overnight.
A: True. Still, she knows the bank will open.

B's utterance arguably challenges only the strict reading of the initial knowledge ascription. Correspondingly, it sounds somewhat pedantic, while still raising the demands for subsequent uses of knowledge sentences.³

Relevant differences. Depending on the context, the loose reading of sentences like (1) to (3) can be more or less demanding, allowing for more or less slack. It's a widely shared assumption that the amount of slack allowed in any given context is determined at least in part by what differences count as relevant (Lasersohn, 1999: 526; Klecha, 2018: 97; Carter, 2019: 19; Hoek, 2018: 173). Take sentence (1). In a context where Hannah is dropping by for a coffee, differences of a couple of minutes don't matter and this sentence will be used with a corresponding amount of slack. The admissible range of points in time surrounding 3pm is relatively wide. Meanwhile, in a context where we are planning our next diamond heist, much smaller differences may matter, and this reduces this range.

This phenomenon seems reflected in familiar stakes effects on knowledge attributions. Consider the bank cases (DeRose, 1992). Hannah wants to decide whether she should wait in line at the bank to deposit her paycheck or return on the next day, a Saturday, when the lines are shorter. She has been to the bank two weeks before on a Saturday, and it was open then. Based on this, she says, "I know

³ For pertinent results from experimental philosophy, see e.g. Gerken et al., 2020 and the references therein.

the bank will be open tomorrow.” Intuitively, the requirements for the proper use of this sentence shift with what is at stake. If it is very important to deposit the paycheck before Monday, we apply stricter standards. For instance, we require safety across a wider range of possible worlds.⁴

This phenomenon is easily re-described to show that relevant differences affect the admissible amount of slack in the case of knowledge ascriptions. In the bank cases, Hannah wants to decide whether she should wait in line or come back the other day. If the stakes are low, it doesn’t matter whether e.g. her belief is safe across every possible world or just the nearby ones, because she will head home either way. Meanwhile, if the stakes are high, this difference becomes relevant, because Hannah may decide to wait in line already when her belief fails to be safe across relatively distant possible worlds. The requirements for proper knowledge attributions vary accordingly. In this way, relevant differences affect the admissible amount of slack.

Slack-regulators. The target expressions in sentences like (1) to (3) all come with an associated set of so-called “slack regulators” (Lasersohn, 1999: 525; Klecha, 2018: 95; Hoek, 2019: 177–178; Carter, 2019: 4–5; Moss, 2019: 262). These expressions reduce the admissible amount of slack, so that the requirements for proper use come closer to the respective strict reading. Consider the following pairs.

(10) a. Hannah arrived at 3pm. / Hannah is 6 foot 2. / The table is flat.

b. Hannah arrived at *exactly* 3pm. / Hannah is *exactly* 6 foot 2. / The table is *perfectly* flat.

In the b-sentences, “exactly” and “perfectly” feature as a slack regulator that raise the standards for proper usage. Slack regulators show an interesting kind of unidirectionality (Carter, 2019: 5–6).

(11) a. Hannah is 6 foot 2, but she isn’t exactly 6 foot 2.

b. ? Hannah isn’t exactly 6 foot 2, but she is 6 foot 2.

The first sentence in this pair sounds perfectly fine, while the reversed, second sentence sounds comparatively odd.

Here the parallels to knowledge ascriptions are least clear. In particular, familiar slack regulators sound off when applied to “knows.”

(12) # Hannah perfectly/exactly knows that the bank will be open.

Still, there are candidates for slack regulators. Moss (2017: 113), for instance, suggests “for sure” as a potential slack regulator (see also Ludlow, 2005: 19–20 and Kyriacou, 2019b: 16 for other candidates).

⁴ The existence of these effects is somewhat controversial, but see e.g. Dinges and Zakkou, forthcoming and the references therein for supportive findings.

- (13) a. Hannah knows the bank will be open.
b. Hannah knows *for sure* that the bank will be open.

Adding “for sure” arguably raises the standard, such that the second sentence comes out as intuitively more demanding than the first. The indicated asymmetry may be felt too, though the intuitions are shaky (one reviewer denies them entirely).

- (14) a. Hannah knows that the bank will be open, but she doesn’t know for sure that it will be open.
b. ? Hannah doesn’t know for sure that the bank will be open, but she knows that it will be open.

In summary, there are striking parallels between knowledge ascriptions and cases of loose talk with respect to a whole range of data points that have been crucial to each debate. This lends strong prima facie credibility to the idea that knowledge ascriptions are cases of loose talk.

4 Pragmatic accounts of loose talk

Despite these parallels, appeals to loose talk are of no avail to skeptical invariantism. As indicated, many authors propose semantic treatments of loose talk, which are, by definition, incompatible with skeptical invariantism. Familiar pragmatic accounts are problematic too, because they clash with the data outlined when applied to knowledge ascriptions, as I subsequently show.

4.1 Lasersohn’s pragmatic halos

Lasersohn (1999) presents the first substantial pragmatic account of loose talk. On his view, sentences (1) to (3) express their respective strict reading, which is normally obviously false or at least obviously unjustified. We normally don’t possess the kinds of measurement instruments that would be required to ascertain e.g. that Hannah arrived at *exactly* 3pm. To make this compatible with the observation that we can properly use these sentences nonetheless, Lasersohn proposes to modify the norms of assertion. A commonly accepted norm of assertion ties the appropriateness of literal assertions to the truth of the proposition expressed.

A literal assertion of S in context C is proper only if the proposition expressed by S in C is true.

On this norm, utterances of sentences (1) to (3) come out as improper most of the time, assuming they are literal and assuming Lasersohn’s semantic commitments. He thus proposes the following replacement norm.

A literal assertion of S in context C is proper only if one proposition in the pragmatic halo of S in C is true.

This norm predicts that utterances of sentences (1) to (3) are often proper, despite the fact that they express obviously false or unjustified propositions. To see this, an account of the pragmatic halo of a sentence in context is required, and Lasersohn provides such an account along the following lines.⁵

The starting idea is to associate each atomic expression in a given sentence with a pragmatic halo. The pragmatic halo of a given atomic expression comprises a set of objects of the same logical type as the original denotation, such that the difference between these objects and the original denotation is “pragmatically ignorable” (526). Take “3pm” in sentence (1). Its halo comprises points in time surrounding the precise point in time 3pm, such that each of these points in time differs from 3pm in pragmatically ignorable ways. What counts as pragmatically ignorable depends on context. If, for instance, we only care about Hannah’s rough time of arrival, the halo of “3pm” might comprise all points in time in the interval of $3\text{pm} \pm 5$ minutes. Where H_C gives you the halo of an expression in such a context C , we can write:

$$H_C(3\text{pm}) = \{t \mid t \in [2:55\text{pm} \dots 3:05\text{pm}]\}$$

The halo of entire sentences is compositionally derived from the halos of the atomic expressions, by deriving one proposition for each combination of elements in these halos.

$$\begin{aligned} H_C(\text{Hannah arrived at } 3\text{pm}) \\ = \{ \{w \mid \langle x, y, w \rangle \in R \} \mid x \in H_C(\text{Hannah}); y \in H_C(3\text{pm}); R \in H_C(\text{arrived at}) \} \end{aligned}$$

In the case at hand, this gives us propositions to the effect that Hannah arrived at t , for each point in time t in the interval $[2:55\text{pm} \dots 3:05\text{pm}]$.

We can now see that utterances of e.g. sentence (1) can be proper on the suggested norm of assertion. Given this norm, they are proper only if one of the propositions in the halo is true, i.e. only if Hannah arrived within the interval of $3\text{pm} \pm 5$ minutes. This condition is easily satisfied in everyday contexts.

Promisingly, ordinary knowledge ascriptions come out as proper too, even assuming skeptical invariantism. For concreteness, consider a skeptical invariantist who adopts the following, reliabilist clause for “knows.”

$$\begin{aligned} \llbracket \text{knows} \rrbracket &= \{ \langle S, p, w \rangle \mid \text{At } w, \\ &\quad \text{i. } S \text{ has a true belief with the content } p, \text{ and} \\ &\quad \text{ii. This belief is } 100\% \text{ reliably formed} \} \end{aligned}$$

⁵ As I understand Lasersohn, loosely used sentences do not assert (or convey) anything other than what they express, and thus they are literal. They correspondingly aim to add obviously false or unjustified propositions to the common ground. This is permissible on the given norm, and we tolerate such propositions in the common ground, as long as their falsity is harmless, by way of a truth in the pragmatic halo. For a competing interpretation on which Lasersohn proposes a kind of non-literality, see Carter, 2019: 7. My worries apply regardless.

Depending on the context, the difference between 100% reliability and, say, 80% reliability may not matter. If we face low stakes, for instance, one’s decisions may be the same either way. We thus obtain a halo for “knows” along the following lines, where N_c is the lowest percentage that differs irrelevantly from 100% in C .

$$H_c(\text{knows}) = \{ \langle S, p, w \rangle \mid \text{At } w, \\ \text{i. } S \text{ has a true belief with the content } p, \text{ and} \\ \text{ii. This belief is } N\% \text{ reliably formed} \mid N \geq N_c \}$$

In our low stakes bank case, for instance, N_c may be 80, such that the halo of “knows” comprises one relation for any N in the interval $[80 \dots 100]$. Given the indicated norm of assertion, a knowledge ascription is thus proper when the subject satisfies one of these relations, i.e., when she has a pertinent true belief that has been formed on a basis that is at least 80% reliable. This condition is satisfiable, and hence ordinary knowledge ascriptions come out proper.

This promising outlook notwithstanding, skeptical invariantists end up with false predictions about embedded knowledge ascriptions. Consider “ S doesn’t know that p .” According to Lasersohn, the halo of this sentence is derived by composing the elements of the halos of its atomic expressions. As Lasersohn (1999: 548) himself points out, the halo of every atomic expression always comprises its literal denotation. This is because the “difference” between the literal denotation and itself is always pragmatically ignorable. For instance, the halo of “knows” always comprises the literal denotation. The halo of “ S doesn’t know that p ” therefore comprises the proposition it literally expresses, which results from composing the literal denotations of its atomic expressions. On standard assumptions about negation, this is the following proposition.

$$\{ w \mid \langle \llbracket S \rrbracket, \llbracket \text{that } p \rrbracket, w \rangle \notin \llbracket \text{knows} \rrbracket \}$$

Assuming the above skeptical invariantist semantics, this proposition is obviously true when it comes to most everyday propositions, for we hardly if ever use 100% reliable belief formation methods. This means that there is an obviously true proposition in the halo of most knowledge denials. Consequently, knowledge denials should be obviously proper most of the time, which is not what we want. When people have a decent amount of evidence, it is normally improper to deny them knowledge, as indicated above.⁶

⁶ Similar problems arise when we apply Lasersohn’s account to embeddings of (1) to (3); see e.g. Carter, 2019: 7 and Hoek, 2019: 172. The relevance theoretic account of loose talk in Sperber and Wilson, 1985 faces similar concerns because, on this account, loose talk can only weaken the proposition expressed. See Hoek, 2018: 159.

4.2 Klecha’s optimality account

Klecha (2018) proposes a different pragmatic account of loose talk, which also starts from the assumption that sentences (1) to (3) express their obviously false or unjustified strict reading. Unlike Lasersohn, he retains the familiar norm of assertion, which ties the appropriateness of literal assertions to the truth of the proposition expressed. But he argues that we use sentences like (1) to (3) non-literally in ordinary contexts. The non-literal message they convey corresponds to their loose reading, which is often true. According to Klecha, respective utterances are proper for this reason (see also Davis, 2007, who appeals to conversational implicature).

Klecha proposes an optimality theoretic account to spell out how people derive the non-literal loose reading (see also Krifka, 2002). According to optimality theory, hearers generally use the following strategy to derive what speakers want to convey with their utterances. They ask, which belief state do I have to attribute to the speaker such that the sentence she uttered comes out as an optimal choice? Once they have figured this out, they take the content of the respective belief state as the message conveyed. Klecha offers an account of optimality that supposedly allows one to derive loose readings in this way. Here is a simplified statement of Klecha’s proposal.

The central notion in Klecha’s account is that of the coarsening of a proposition. It relies on a certain way to partition the set of possible worlds with which I will begin. In any given context, people have things they want to achieve, they have so-called “domain goals” (Roberts, 2012). Depending on what the world is like, one or another means will be conducive to these goals, but not every aspect of the world matters. Suppose Hannah wants to drop by for a coffee, and I want to make preparations. In this context, it doesn’t matter to me whether Hannah arrives at 3pm or 3:01pm, maybe because I will set up the table some minutes before 3pm anyway. We can therefore use the domain goals to induce a partition on the set of possible worlds, where each partition cell contains a set of worlds such that whichever of these worlds we occupy, the recommended means to achieve our domain goals are the same. Let’s call the partition so-induced in a given context C , π_C .

We can now define the coarsening, K_C , of a proposition p in a context C . Intuitively, coarsening makes propositions less precise, so that they speak only to the domain goals. Formally, coarsening gives you the union of the partition cells in π_C that overlap with p (Figure 1).

$$K_C(p) = \bigcup \{p' \in \pi_C \mid p' \cap p \neq \emptyset\}$$

To illustrate, the coarsening of the strict reading of (1) no longer tells us that we find ourselves in a world where Hannah arrived at *exactly* 3pm. Instead, it tells us that we find ourselves in world from a partition cell containing worlds where this is so. Thus, we could still be in a world where Hannah arrived at 3:01pm as long as such a world is clustered together with a world where she arrived at 3:00pm.




Figure 1. Schematic depiction of the coarsening of two different propositions. Lines demarcate partition-cells induced by π_C . The grey area on the left covers the worlds in which the proposition is true. The result of coarsening is shown in grey on the right.

With this idea of coarsening at hand, Klecha suggests that the proposition conveyed by a sentence S in context C is determined as follows (simplifying a bit; see Klecha, 2018: 108). The subsequent principle explicates the optimality theoretic idea that the conveyed message corresponds to the content of the belief state that make the chosen sentence optimal.

The proposition conveyed by S in $C = \cup\{K_C(B) \mid B \in \mathcal{B}_C \text{ and } S \in \text{opt}(\mathcal{S}_C|B)\}$

\mathcal{B}_C stands for the set of belief states the speaker is in, for all the audience knows (see Klecha, 2018: 107). A belief state is represented as the set of possible worlds that are consistent with it. \mathcal{S}_C denotes the set of “utterance alternatives,” i.e., roughly, the set of sentences that the speaker can relevantly utter in the context at hand, including the actually uttered sentence, S (see Klecha, 2018: 108). The set $\text{opt}(\mathcal{S}_C|B)$ is the set of optimal sentences among \mathcal{S}_C on the assumption that the speaker is in belief state B . The proposition conveyed thus corresponds to the union of the coarsenings of the possible belief states on behalf of the speaker relative to which the actually uttered sentence is optimal compared to its alternatives.

What is missing is a definition of the set $\text{opt}(\mathcal{S}_C|B)$, the set of optimal sentences relative to a belief state. Klecha offers three constraints that together determine this set. One constraint is Faithfulness, where sentences rank highly roughly when the speaker believes the coarsening of the proposition expressed (see Klecha, 2018: 109). Another constraint is Manner, where sentences rank highly when they are perspicuous in a relevant sense, e.g., when they are brief and use “round” expressions such as “3pm” instead of “3:01pm” (see Klecha, 2018: 98–99, 109–110; see also Krifka, 2002). For our purposes, we can focus exclusively on the third constraint called Informativity, which ranks sentences as follows.

$\text{INFORMATIVITY}(S) \geq \text{INFORMATIVITY}(S')$ iff $K_C(\llbracket S \rrbracket) \subseteq K_C(\llbracket S' \rrbracket)$

A sentence is more informative than another, and thus pro tanto more optimal, when the coarsening of the proposition expressed by the former sentence is stronger than the coarsening of the proposition expressed by the latter.

This principle gets Klecha in trouble when it comes to knowledge denials. To see this, consider a context like the bank cases where we only care about whether Hannah knows₀ that the bank will be open, because this will settle our decision whether to wait in line. In such a context, π_C partitions the set of possible worlds into worlds where Hannah at least knows₀ that the bank will be open and worlds where she doesn't at least know₀ that.

$$\pi_C = \left\{ \begin{array}{l} \{w \mid \text{At } w, \text{ Hannah at least knows}_0 \text{ that the bank will be open}\}, \\ \{w \mid \text{At } w, \text{ Hannah doesn't at least know}_0 \text{ that the bank will be open}\} \end{array} \right\}$$

Now look at the result of coarsening the proposition expressed by “Hannah doesn't know that the bank will be open.” On a skeptical invariantist semantics, coarsening this proposition leads to a complete triviality, the entire set of possible worlds. This is because lacking knowledge_s is compatible with both lacking and having knowledge₀. The intersection with every partition cell is non-empty (Figure 2).




Figure 2. Schematic depiction of the coarsening of a knowledge denial in a bank case context. Lines demarcate partition-cells in π_C . The grey area on the left covers the worlds in which the knowledge denial is true on a skeptical invariant semantics. The result of coarsening is shown in grey on the right.

The knowledge denial thus comes out as maximally uninformative. It's coarsening is entailed by the coarsening of every sentence, and hence no sentence is less informative. Such utterly uninformative sentences shouldn't be optimal relative to any belief state we are normally prepared to ascribe.⁷ This means that, when we derive the proposition conveyed by the principle above, i.e., by unioning the belief states relative to which the uttered sentence is optimal, we end up with the empty set. Klecha's theory predicts that knowledge denials typically convey the equivalent of a contradiction, which should make them trivially improper.⁸ This is evidently no improvement on Lasersohn's theory, which made knowledge denials trivially proper.

4.3 Hoek's conversational exculpatures

Hoek (2018, 2019) provides another pragmatic account of loose talk. Like Klecha, he denies the literalness of typical utterances of (1) to (3) to square their propriety with standard norms of assertion.

⁷ The inference here is a bit compressed and the details depend on exactly which constraints on optimality we adopt and how they are ranked. I hope the inference is intuitive enough as it stands. See Dinges, MS for details.

⁸ Klecha (2018: 108–109) points out a rescue strategy for when the conveyed proposition is the empty set. But this rescue strategy is designed to deal with the specific phenomenon of unidirectionality, and it is of no avail in the present context. See Dinges, MS for details.

Instead of appealing to optimality theory, though, he derives the conveyed message through a mechanism that he dubs *conversational exculpature*.

Conversational exculpature is a type of pragmatic repair strategy that hearers apply under certain conditions. Hearers apply this repair strategy when the proposition literally expressed by an utterance fails to be wholly relevant to the question under discussion (henceforth, QUD). The notion of being wholly relevant is defined against the background of the following, familiar understanding of a question.

A **question** Q is a partition of the space W of possible worlds. (2019: 173)

Here, each partition cell stands for one complete answer to the question Q . Take the question *how many cows are there on this ranch*. This question induces one partition cell for each answer of the form *there are exactly n cows on this ranch*, containing just those worlds where this answer is true. Against this background, we can define what it means to be wholly relevant.

A proposition $p \subseteq W$ is **wholly relevant** to a question Q if and only if no partition cell of Q contains worlds where p is false and also worlds where p is true. (2019: 173)

To illustrate, the responses *there are exactly three cows on this ranch* and *there are between three and five cows on this ranch* are both wholly relevant to the question above. Meanwhile, the response *there are exactly three cows and a donkey on this ranch* is not because it is true in only some of the worlds in the three-cows-partition cell, namely, those where there is also a donkey (notice that coarsening in Klecha's sense leads to a proposition that is wholly relevant to the question/partition π_C).

Here is how the so-triggered repair strategy works. The basic idea is to replace the proposition expressed with another proposition, which is then taken as being conveyed. Which proposition replaces the proposition expressed? In making utterances, speakers often *speak as if* certain assumptions were true, while in fact, they are knowingly false or uncertain (Hoek, 2018: 160–161). These *pragmatic presuppositions*, together with the QUD, determine the proposition conveyed by sentence S in context C , as follows.

The proposition conveyed by S in C = the unique proposition p , if there is any, such that

- (i) p is wholly relevant to the QUD, and
- (ii) p is equivalent to the proposition expressed conditional on the pragmatic presuppositions (Hoek, 2019: 174).⁹

⁹ To be precise, Hoek (2018: 160) appeals to contextual presuppositions (following Simons, 2002) rather than pragmatic presuppositions. I use the latter notion because it is more familiar and seems to do the job as well. See below for more on pragmatic presuppositions.

Consider a case where, according to Hoek, this mechanism is at work independently of loose talk. Suppose the QUD is *what type of hat does Hannah wear*, which induces one partition cell for each type of hat there is. You respond, “She wears the same type of hat Sherlock Holmes wears.” This response isn’t wholly relevant. Take e.g. the partition cell induced by the response that Hannah wears a straw-hat. The given response is true in only some of the worlds in this partition cell, namely, those where both Hannah and Sherlock Holmes wear a straw-hat. This triggers conversational exculpation, and the proposition conveyed is determined as follows. Naturally enough, the speaker speaks as if the Sherlock Holmes stories are true by “referring” to Sherlock Holmes. Given this presupposition, there is a unique proposition that is (i) wholly relevant to the QUD and (ii) conditionally equivalent to the proposition expressed. This is the proposition that Hannah wears a deerstalker. This proposition is evidently wholly relevant. It is conditionally equivalent to proposition expressed for the following reason. If Hannah wears the same hat as Sherlock Holmes does while the Holmes stories are true, then she wears a deerstalker. After all, Holmes wears a deerstalker. Conversely, if Hannah wears a deerstalker while the Holmes stories are true, then she wears the same hat as Sherlock Holmes does. Other wholly relevant propositions—such as the proposition that Hannah wears a straw-hat—aren’t conditionally equivalent to the proposition expressed. The proposition that Hannah wears a deerstalker is therefore conveyed (which is desirable, because if this proposition is conveyed, the speaker is not committed to the existence of Sherlock Holmes).

Let’s turn to loose talk. According to Hoek, loose talk is a special case of conversational exculpation, where the contextual presupposition involved takes the specific shape of a *scale presupposition*. Scale presuppositions are presuppositions induced when speakers select an expression from a conventionally restricted *scale* of alternative expressions. Their content is, roughly, that the true description of the object in question involves one of these expressions. For instance, speakers of English conventionally use expressions from the following scale to describe personal heights, at least in contexts where precision is not required.

{... “5 foot 11,” “6 feet,” “6 foot 1”, “6 foot 2” ...}

By uttering, “Hannah is 6 foot 2,” you select an expression from this scale, thereby indicating that you are following the indicated convention. By following this convention, you speak as if Hannah’s height was located on the indicated scale. Hence, you (falsely) presuppose that Hannah is an exact number of inches tall. This is a scale presupposition.

Let’s see how this scale presupposition figures in a specific application of conversational exculpation to loose talk. Suppose we want to know Hannah’s height to the nearest inch, a question that induces partition cells containing worlds where Hannah is $n \pm \frac{1}{2}$ inches tall. You respond, “Hannah is 6 foot 2.” Given Hoek’s semantic commitments, this response isn’t wholly relevant because it is true in only some of the worlds in the $6 \text{ foot } 2 \pm \frac{1}{2}$ -inches-cell. This triggers conversational exculpation. By the

reasoning above, your use of “6 foot 2” induces the scale presupposition that Hannah is an exact number of inches tall. Given that, there is a unique proposition that is (i) wholly relevant to the QUD and (ii) conditionally equivalent to the proposition expressed. This is the proposition that Hannah is 6 foot 2 \pm ½ inches. Evidently, this proposition is wholly relevant. It is conditionally equivalent to the proposition expressed for the following reason. If Hannah is 6 foot 2, then Hannah is 6 foot 2 \pm ½ inches independently of any additional assumptions. Meanwhile, if Hannah is 6 foot 2 \pm ½ inches and we assume, via the scale presupposition, that Hannah is an exact number of inches tall, then she must be 6 foot 2. The proposition that Hannah is 6 foot 2 \pm ½ inches is therefore conveyed, which is as desired.

This account doesn’t have specific problems with embedded knowledge ascriptions (see below). Unfortunately, though, it doesn’t even make the right predictions when it comes to unembedded ones. On the suggested account, loose talk arises against the background of a scale presupposition, which, in turn, depends on a scale of expressions, which encodes conventional ways to describe e.g. the height of a person. To apply Hoek’s account to knowledge ascriptions, we need such a scale for “knows,” containing expressions conventionally used to describe each other’s epistemic states. No scale suits our purposes, however.

Consider the range of expressions that we conventionally use to describe each other’s epistemic states. Sentence like “I know/heard/read/saw/recall that p” are all used regularly in this context. Therefore, we should end up with a scale like this.

{“know,” “hear,” “read,” “see,” “recall,” ...}

But this scale doesn’t give us what we want. The proposition conveyed with a knowledge ascription should be some proposition to the effect that S at least knows₀ that p. This would deliver the correct predictions about when knowledge ascriptions are appropriate. So, if Hoek’s account of loose talk is correct, and knowledge ascriptions are a case of loose talk, this proposition should be conditionally equivalent to the proposition expressed, i.e., the proposition that S knows_S that p. But that’s not the case. Even given the assumption that we are located on the indicated scale, the proposition that S at least knows₀ that p doesn’t entail that S knows_S that p. For instance, S may know₀ that p because S has read that p. In that case, S is in one of the states on the scale and still lacks knowledges.

One could instead propose the following scale, which is familiar from discussions of so-called Horn-scales (e.g. Levinson, 2000: 87).

{“believe,” “know”}

But again, the proposition that S at least knows₀ that p doesn’t entail that S knows_S that p, even when it is stipulated that we are located on this scale. Indeed, the proposition that S at least knows₀ that p

entails the scale presupposition because it entails that S believes that p. As such, the scale presupposition cannot generate any novel entailments.

One might suggest the following scale instead.¹⁰

{“merely believe,” “know”}

To assess this scale, we first have to get clear on what “merely believing” is. One idea would be that merely believing p is believing p without any evidence. This would lead to correct predictions e.g. in a bank case context, where we only care about whether Hannah at least knows₀ that the bank will be open (henceforth, B) and where the QUD thus entails the bi-partite partition familiar from the discussion of Klecha’s account. According to skeptical invariantism, the proposition expressed by “Hannah knows B” is the proposition that Hannah knows_s B, which isn’t wholly relevant. This triggers conversational exculpation. On the assumption that Hannah is located on the above scale interpreted as described, the proposition that Hannah knows_s B is uniquely conditionally equivalent to the wholly relevant proposition that she at least knows₀ B. For if Hannah at least knows₀ B, while being located on the above scale, then the only option is that she knows_s B. Meanwhile, if Hannah knows_s B then she at least knows₀ B independently of any additional assumptions.

Unfortunately, it just seems false that, whenever we make appropriate knowledge ascriptions in ordinary contexts, we speak as if we are located on the indicated scale, i.e., as if we either know_s something or lack evidence entirely. Suppose I say, “Hannah doesn’t know who perpetrated the crime. She found John’s weapon, but someone might have placed it there. Meanwhile, Sarah knows who did it, because she talked to the neighbor who saw everything.” I am acknowledging a whole range of possible epistemic positions here, and if I invoke any scale at all, it is the rich scale initially provided, which doesn’t yield the right interpretation.

“Merely believing” something can be interpreted in other ways. Two natural options would be as believing without knowings_s and as believing without knowing₀. On either interpretation, however, we face by now familiar problems. On the interpretation as believing without knowings_s, we cannot establish the conditional equivalence between knowings_s and at least knowing₀. This is because our scale presupposition now says that we either believe p without knowings_s p or know_s p. Given that knowledges entails belief, this just means that we believe p. Even granting this assumption, knowing₀ p doesn’t entail knowings_s p. On the interpretation as believing without knowing₀, we can derive the required equivalence (readers can verify this for themselves). But once more, it seems implausible that we typically presuppose that our epistemic position is either weaker than knowledges₀ or as strong as knowledges. In the crime example above, the speaker acknowledges many epistemic states

¹⁰ Thanks to Wayne Davis.

(testimonial and perceptual belief, for instance) that often amount to knowledge₀ while being insufficient for knowledge_s.¹¹

5 A novel approach

The appeal to loose talk on behalf of skeptical invariantism fails if we assume that current theories of loose talk are on the right track. This puts strong pressure on proponents of skeptical invariantism who defend their view by appeal to loose talk. In what follows, I explore a closely related defense of skeptical invariantism, which appeals to the more general phenomenon of conversational exculpation instead. This approach is more promising. To make this clear, I begin with a brief presentation of epistemic contextualism as proposed by Lewis (1996) and spelled out by Blome-Tillmann (2014). This will serve as a helpful background later on.

5.1 Epistemic contextualism

Lewis (1996) presents a familiar contextualist semantics for “knows” along the following lines.

$$\llbracket \text{knows} \rrbracket_C = \left\{ \langle S, p, w \rangle \mid \begin{array}{l} \text{At } w, S\text{'s evidence eliminates all } \neg p\text{-worlds,} \\ \text{except for those that are properly ignored in } C \end{array} \right\}$$

On this view, “knows” expresses differently demanding epistemic relations depending on which possible worlds are “properly ignored” in the context hand. Blome-Tillmann (2014) elaborates on this proposal, offering specific rules for when a possible world counts as properly ignored. He largely accepts most of Lewis’ rules, specifically, the Rules of Actuality, Resemblance, Belief, Reliability, Method and Conservatism (31). Crucially, though, he replaces Lewis’ Rule of Attention by the Rule of Presupposition (20; I rephrased this rule for presentational purposes).

Rule of Presupposition

A possible world w is properly ignored in C only if the pragmatic presuppositions in C entail that w doesn’t obtain.

For instance, we must presuppose that the bank hasn’t changed its hours in order to properly ignore worlds where it has. Following Stalnaker (1978), Blome-Tillmann (2014: 30) defines the relevant notion of pragmatic presupposition as follows.

x pragmatically presupposes p in $C \leftrightarrow x$ is disposed to behave, in her use of language, as if she believed p to be common ground in C .

¹¹ Another potential scale would be {“know,” “know for sure,” “know with absolute certainty”}. Thanks to Christos Kyriacou. But the corresponding scale presupposition entails that S knows_s that p because, according to skeptical invariantism, “know” already has this demanding interpretation. It seems implausible that speakers typically presuppose that all available epistemic states are at least that strong; see the crime case above.

Importantly, on this conception of pragmatic presupposition, one can pragmatically presuppose p even if one doesn't believe p . One must only behave *as if* one took p to be common ground. This aligns with Hoek's characterization of pragmatic presupposition also in terms of "speaking as if." For instance, we pragmatically presuppose that the bank hasn't changed its hours once we speak as if this was common ground, independently of whether this is actually so.

To accommodate data such as those outlined in §3, Blome-Tillmann argues that pragmatic presuppositions vary e.g. with salient alternatives and stakes (36–37). In the low stakes bank cases, for instance, we pragmatically presuppose that the bank hasn't changed its hours, behaving as if this was common ground. When someone makes this possibility salient, by mentioning it, we stop presupposing this. Similarly, high stakes force us to take this possibility seriously, which again leads us to abandon the respective presupposition. This variation in pragmatic presuppositions leads to variations in which worlds our evidence needs to eliminate before we can properly ascribe "knowledge," via the Rule of Presupposition. And this explains, e.g., why knowledge ascriptions are sensitive to stakes and salient error-possibilities.

5.2 Knowledge ascriptions and conversational exculpation

Now consider a skeptical invariantist who adopts the following, maximally demanding, invariant variation on Lewis' semantics for "knows."

$$\llbracket \text{knows}_S \rrbracket = \llbracket \text{knows} \rrbracket = \{ \langle S, p, w \rangle \mid \text{At } w, S\text{'s evidence eliminates all } \neg p\text{-worlds} \}$$

Such a skeptical invariantist faces the familiar challenge to explain why we properly ascribe knowledge all the time, given that our evidence hardly ever eliminates every $\neg p$ -world. Based on Blome-Tillmann's assumptions about pragmatic presuppositions, conversational exculpation yields a promising response.

Let's begin with a stipulatively defined notion of *properly treating a possible world as ruled out*. Properly treating a possible world as ruled out is just like properly ignoring it, except that the following rule replaces the Rule of Presupposition.

*Rule of Presupposition**

A world w is properly treated as ruled out in C only if the pragmatic presuppositions in C entail that the target subject's evidence eliminates w .

Properly treating a world as ruled out, unlike properly ignoring it, depends on presuppositions about our evidence, particularly, presuppositions about the evidence available to the target subject, i.e., the subject to whom knowledge is ascribed. Properly treating a world w as ruled out requires the presupposition that this subject can rule w out.

I submit that we properly treat many possible worlds as ruled out in ordinary contexts. More specifically, I submit that the worlds we properly treat as ruled out are more or less the same worlds as the worlds that, according to Blome-Tillmann, we properly ignore (see below for some important divergences). For instance, in the low stakes bank case, we properly ignore worlds where the bank has changed its hours, where a meteor destroys it overnight, where an evil demon deceives us, etc. This is because we presuppose that the bank hasn't changed its hours, etc. Blome-Tillmann at least finds this assumption "intuitive" (21), and I agree. It is likewise intuitive, I think, that we properly treat the indicated worlds as ruled out. For instance, we plausibly presuppose that Hannah can rule out that the bank has changed its hours, that she can rule out that a meteor destroys it overnight, that she can rule out that an evil demon deceives us, etc. We plausibly treat all of these worlds as if it was *settled* that they don't obtain.

Importantly, in the case of both properly ignoring and properly treating as ruled out, the respective presuppositions need not align with our beliefs or our knowledge. We may believe or know neither that the bank hasn't changed its hours nor that Hannah can rule this out. We make the relevant presuppositions nevertheless by *speaking as if* they were true. The situation here is analogous to the cases of conversational exculpation discussed above, which also featured knowingly false presuppositions such as scale presuppositions or presuppositions about Sherlock Holmes.

With this in mind, let's turn to how conversational exculpation applies to knowledge ascriptions. First, consider the QUD and the pragmatic presuppositions in e.g. the bank cases. As for the QUD, we are interested in Hannah's evidential position, for this bears on our decision to wait in line at the bank or to come back the other day. However, not every aspect of her evidential position matters. Consider, for instance, possible worlds that we properly treat as ruled out. By the Rule of Presupposition*, we presuppose that Hannah's evidence eliminates those worlds and so we presumably don't care about them. A natural QUD would be, "Which possible worlds *that aren't properly treated as ruled out* does Hannah's evidence eliminate?" As for the pragmatic presuppositions, consider once more the possible worlds we properly treat as ruled out. By the Rule of Presupposition*, we pragmatically presuppose that Hannah's evidence eliminates them, i.e., we pragmatically presuppose that Hannah's evidence eliminates all possible worlds that are properly treated as ruled out.

With these assumptions about the QUD and the pragmatic presuppositions in place, conversational exculpation makes promising predictions about knowledge ascriptions. Consider an utterance of "Hannah knows that the bank will be open (B)" in the bank cases. According to the above version of skeptical invariantism, this utterance expresses the proposition that Hannah's evidence eliminates *all* $\neg B$ -worlds. This proposition isn't wholly relevant to the QUD, which triggers conversational exculpation. Conversational exculpation then predicts that the knowledge ascription conveys that Hannah's evidence eliminates all $\neg B$ -worlds *that aren't properly treated as ruled out*. Here is why. The latter proposition is evidently wholly relevant to the QUD. It is unique in also being conditionally

equivalent to the proposition expressed, i.e., the proposition that Hannah’s evidence eliminates *all* $\neg B$ -worlds. If Hannah’s evidence eliminates *all* $\neg B$ -worlds, then her evidence eliminates all $\neg B$ -worlds *that aren’t properly treated as ruled out*, independently of any additional assumptions. Meanwhile, if her evidence eliminates all $\neg B$ -worlds *that aren’t properly treated as ruled out*, then her evidence eliminates *all* $\neg B$ -worlds, given the pragmatic presupposition that her evidence eliminates all worlds *that are properly treated as ruled out*.

Knowledge ascriptions thus convey that the subject’s evidence eliminates all $\neg p$ -worlds that aren’t properly treated as ruled out. More precisely, they convey that the subject knows_o that p , where the exact requirements for knowledge_o shift with the context as follows.

$$\llbracket \text{knows}_o \rrbracket_c = \left\{ \langle S, p, w \rangle \left| \begin{array}{l} \text{At } w, S\text{'s evidence eliminates all } \neg p\text{-worlds,} \\ \text{except for those that are properly treated as ruled out in } C \end{array} \right. \right\}$$

Given that we properly treat as ruled out more or less the same worlds as the worlds that we properly ignore, as argued above, it follows that ordinary knowledge ascriptions are often proper, even assuming skeptical invariantism. For ruling out the remaining worlds is typically feasible (that’s at least what contextualists like Lewis and Blome-Tillmann assume).

I can also make good on the promissory note above that I will be able to explain why the propriety conditions for ordinary knowledge ascriptions vary with the context. On my view, this is due to variations in which possible worlds are properly treated as ruled out, where these variations in turn are due to variations in our pragmatic presuppositions. Such variations shift the proposition conveyed with knowledge ascriptions and thereby the conditions under which they are appropriate. My account thus mimics the precise and detailed predictions of Blome-Tillmann’s contextualism, while sticking with an invariantist semantics.

Unlike loose use accounts, I can also explain the data from §3. Let’s begin with knowledge denials, which spelled doom for Lasersohn’s and Klecha’s theory of loose talk when applied to “knows.” Conversational exculpation easily explains why knowledge denials are neither trivially proper nor trivially improper. Take an utterance of “Hannah doesn’t know B” in the bank cases. According to skeptical invariantism, this utterance expresses the proposition that Hannah’s evidence doesn’t eliminate *all* $\neg B$ -worlds (i.e., that Hannah doesn’t know_s B). Once more, this proposition isn’t wholly relevant to the QUD, which triggers conversational exculpation. Conversational exculpation then predicts that the knowledge denial conveys that Hannah’s evidence doesn’t eliminate all $\neg B$ -worlds *that aren’t properly treated as ruled out* (i.e., that Hannah doesn’t know_o B). This proposition is unique in being wholly relevant and conditionally equivalent to the proposition expressed. It is evidently wholly relevant given the QUD. It is conditionally equivalent to the proposition expressed—that Hannah’s evidence doesn’t eliminate *all* $\neg B$ -worlds—for the following reason. If Hannah’s evidence doesn’t eliminate all $\neg B$ -worlds *that aren’t properly treated as ruled out*, then her evidence doesn’t eliminate *all*

\neg B-worlds, independently of any additional assumptions. Meanwhile, if her evidence doesn't eliminate *all* \neg B-worlds, then her evidence doesn't eliminate all \neg B-worlds *that aren't properly treated as ruled out*, given the presupposition that her evidence eliminates all worlds *that are properly treated as ruled out*.

When it comes to salient alternative and stakes effects, I can once more mimic Blome-Tillmann's contextualist account. Following Blome-Tillmann, salient error-possibilities and high stakes, shift our presuppositions and thereby what we properly treat as ruled out. This affects the proposition conveyed by knowledge ascriptions via conversational exculpation and thereby the conditions under which these utterances are appropriate. For instance, in the low stakes bank case, we properly treat worlds as ruled out where the bank has changed its hours, by presupposing that Hannah's evidence eliminates them. Meanwhile, in the high stakes bank case, we give up this presupposition.

Consider, finally, the assumed slack regulator "for sure." Knowledge ascriptions are no longer cases of loose talk, and correspondingly, we can no longer apply Hoek's suggested account of slack regulators, which is tied to loose talk and doesn't carry over to conversational exculpation in general. He suggests that familiar slack regulators shift the governing scale to a more fine-grained scale (2019: 177–178). Since I have abandoned scales, this can't be what's going on. As indicated at the outset, though, the parallel between knowledge ascriptions and loose talk is feeblest when it comes to slack regulators, and familiar slack regulators don't apply to "knows." So this outcome seems just what we want.

At the same time, we only have to modify Hoek's account slightly to offer a plausible account of "for sure." According to Hoek, slack regulators express some kind of interpretation instruction, namely, an instruction to use a finer-grained scale than one would have used otherwise. Our target expression "for sure" doesn't express just that instruction but something similar, namely, an instruction to abandon at least some of the pragmatic presuppositions that one would otherwise have used to interpret the target utterance. Consider an utterance of "Hannah knows for sure that the bank will be open" as made in the low stakes bank case. One would have presupposed e.g. that Hannah's evidence eliminates worlds where the bank changes its hours, but "for sure" instructs one to abandon such presuppositions. Once these presuppositions have been abandoned, conversational exculpation yields a more demanding interpretation, namely, an interpretation on which Hannah's evidence eliminates e.g. the possibility of changed opening hours.

5.3 Comparison

Given the close parallel to Blome-Tillmann's contextualism, let me briefly explore whether the suggested view has any distinctive advantages. Here are two initial candidates.

First, Blome-Tillmann arguably has problems with certain kinds of third-personal knowledge ascriptions. Emily says, “Joe is directing plays again,” thereby pragmatically presupposing that Joe directed plays before (henceforth, P). She subsequently says, “Nathan knows P.” This knowledge ascription threatens to be trivially true on Blome-Tillmann’s contextualism. It is true on this account iff Nathan’s evidence eliminates all $\neg P$ -worlds, except for those that are properly ignored. But why couldn’t Emily properly ignore just every $\neg P$ -world? The Rule of Presupposition, for instance, tells us that we can properly ignore a $\neg P$ -world only if our pragmatic presuppositions entail that it doesn’t obtain. Since Emily pragmatically presupposes P, this condition is satisfied for every $\neg P$ -world. Maybe Lewis’ other rules help to explain why Emily cannot properly ignore some $\neg P$ -worlds, but Ichikawa (2015) suggests that we can construct the case in such a way that they don’t.¹² This problem doesn’t arise on my exculpation account. On this account, the knowledge ascription conveys that Nathan’s evidence eliminates every $\neg P$ -world, except for those that are properly treated as ruled out. The Rule of Presupposition* tells us that we can properly treat a $\neg P$ -world as ruled out only if we presuppose that, in this case, Nathan’s evidence eliminates it. But we don’t presuppose that in the case at hand. We presuppose that $\neg P$ -worlds don’t obtain. It doesn’t follow that we presuppose that Nathan’s evidence eliminates them. As indicated, such presuppositions often go together. In third-person cases, however, they tend to come apart, in a way that favors the exculpation account.

Second, Blome-Tillmann himself grants that his version of contextualism posits an “entirely novel and previously unheard of type of context-sensitivity” (124) by tying *semantic* context-sensitivity to *pragmatic* presuppositions. While he plausibly suggests that this need not be a problem, it still seems that, other things being equal, it would be nice to avoid such commitments. Skeptical invariantists can avoid such commitments if they see knowledge ascriptions as just another instantiation of conversational exculpation. Relatedly, and even if Blome-Tillmann’s form of context-sensitivity was preceded, Grice’s “Modified Occam’s Razor” favors skeptical invariantism as outlined. Senses are not to be multiplied beyond necessity. In at least one sense of “sense,” Blome-Tillmann posits many senses where skeptical invariantists posit just one (see Hazlett, 2007 for a pertinent construal of the Razor).

6 Conclusion

Skeptical invariantists must explain how ordinary knowledge ascriptions can be appropriate. Appeals to loose talk are of no avail if we assume that extant accounts of loose talk are on the right track. Skeptical invariantists can still appeal to the general phenomenon of conversational exculpation to explain proper knowledge ascription once they take on board familiar assumptions about the dynamics of pragmatic presuppositions. The resulting view combines the virtues of invariantism and

¹² See Blome-Tillmann (2015) and Ichikawa (2017: 23n25) for further discussion.

contextualism. It posits a unique knowledge relation, and it offers an account of the context-sensitivity of ordinary knowledge ascriptions, which is on a par, in terms of detail and predictive power, with the most elaborate versions of contextualism. This is achieved while relying solely on the independently motivated pragmatic machinery of conversational exculpation. Overall, skeptical invariance comes out as a highly attractive position, not the last resort that it is often taken to be.

Acknowledgements I am grateful to Wayne Davis, Katharina Felka, Daniel Hoek, Miguel Hoeltje, Dirk Kindermann, Christos Kyriacou, Alex Steinberg, Emanuel Viebahn, Julia Zakkou and an anonymous reviewer for very helpful comments on earlier versions of this paper.

Bibliography

- Blome-Tillmann, M. (2014), *Knowledge and Presuppositions* (Oxford: Oxford University Press).
- (2015), "Ignorance, Presuppositions, and the Simple View", *Mind*, 124/496: 1221–1230.
- BonJour, L. (2010), "The myth of knowledge", *Philosophical Perspectives*, 24/1: 57–83.
- Carter, S. (2019), "The dynamics of loose talk", *Noûs*, Early View: 1–28.
- Conee, E. (2005), "Contextualism contested", in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Malden, MA: Blackwell Publishing), 47–56.
- Davis, W. A. (2007), "Knowledge claims and context. Loose use", *Philosophical Studies*, 132/3: 395–438.
- DeRose, K. (1992), "Contextualism and knowledge attributions", *Philosophy and Phenomenological Research*, 52/4: 913–929.
- (2012), "Replies to Nagel, Ludlow, and Fantl and McGrath", *Philosophy and Phenomenological Research*, 84/3: 703–721.
- Dinges, A. (MS), "Generalizing loose talk".
- (2016), "Skeptical pragmatic invariantism: Good, but not good enough", *Synthese*, 193/8: 2577–2593.
- Dinges, A. and J. Zakkou (forthcoming), "Much at stake in knowledge", *Mind & Language*.
- Fantl, J. and M. McGrath (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).
- Gerken, M; J. Alexander; C. Gonnerman et al. (2020), "Salient alternatives in perspective", *Australasian Journal of Philosophy*, Latest Articles: 1–19.
- Grice, H. P. (1989), *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).
- Hazlett, A. (2007), "Grice's razor", *Metaphilosophy*, 38/5: 669–690.
- Hoek, D. (2018), "Conversational exculpature", *The Philosophical Review*, 127/2: 151–196.
- (2019), "Loose talk, scale presuppositions and QUD", *Proceedings of the 22nd Amsterdam Colloquium*, 171–180.
- Ichikawa, J. J. (2015), "Ignorance and Presuppositions", *Mind*, 124/496: 1207–1219.
- (2017), *Contextualising Knowledge: Epistemology and Semantics* (Oxford: Oxford University Press).
- Kindermann, D. (2019), "Knowledge embedded", *Synthese*, Online First.
- Klecha, P. (2018), "On unidirectionality in precisification", *Linguistics and Philosophy*, 41/1: 87–124.
- Krifka, M. (2002), "Be Brief and Vague! And How Bidirectional Optimality Theory Allows for Verbosity and Precision", in D. Restle and D. Zaefferer (eds.), *Sounds and Systems. Studies in Structure and Change. A Festschrift for Theo Vennemann* (Berlin: De Gruyter Mouton), 429–448.

- (2007), “Approximate interpretations of number words. A case for strategic communication”, in G. Bouma; I. Krämer and J. Zwarts (eds.), *Cognitive Foundations of Interpretation* (Amsterdam: Koninklijke Nederlandse Akademie van Wetenschappen), 111–126.
- Kyriacou, C. (2017), “Bifurcated sceptical invariantism”, *Journal of Philosophical Research*, 42: 27–44.
- (2018), “From moral fixed points to epistemic fixed points”, in C. Kyriacou and R. McKenna (eds.), *Metaepistemology. Realism and Anti-Realism* (Cham: Springer International Publishing), 71–95.
- (2019a), “Assertion and practical reasoning, fallibilism and pragmatic skepticism”, *Acta Analytica*.
- (2019b), “Semantic awareness for skeptical pragmatic invariantism”, *Episteme*, First View: 1–19.
- Lasersohn, P. (1999), “Pragmatic halos”, *Language*, 75/3: 522–551.
- Lauer, S. (2012), “On the pragmatics of pragmatic slack”, *Proceedings of Sinn und Bedeutung 16*: 389–402.
- Levinson, S. C. (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (Cambridge, MA: MIT Press).
- Lewis, D. (1979), “Scorekeeping in a language game”, *Journal of Philosophical Logic*, 8/1: 339–359.
- (1996), “Elusive Knowledge”, *Australasian Journal of Philosophy*, 74/4: 549–567.
- Ludlow, P. (2005), “Contextualism and the New Linguistic Turn in Epistemology”, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 11–50.
- Moss, S. (2017), *Probabilistic Knowledge* (Oxford: Oxford University Press).
- (2019), “Full belief and loose speech”, *Philosophical Review*, 128/3: 255–291.
- Roberts, C. (2012), “Information structure in discourse. Towards an integrated formal theory of pragmatics”, *Semantics and Pragmatics*, 5/6: 1–69.
- Simons, M. (2002), “Presupposition and relevance”, in Z. G. Szabó (ed.), *Semantics vs. Pragmatics* (Oxford: Clarendon Press), 329–355.
- Solt, S. (2014), “An alternative theory of imprecision”, *Proceedings of SALT 24*: 514–533.
- Sperber, D. and D. Wilson (1985), “Loose Talk”, *Proceedings of the Aristotelian Society*, 86: 153–171.
- Stalnaker, R. C. (1978), “Assertion”, in P. Cole (ed.), *Pragmatics* (New York, NY: Academic Press), 315–332.