RESPECT

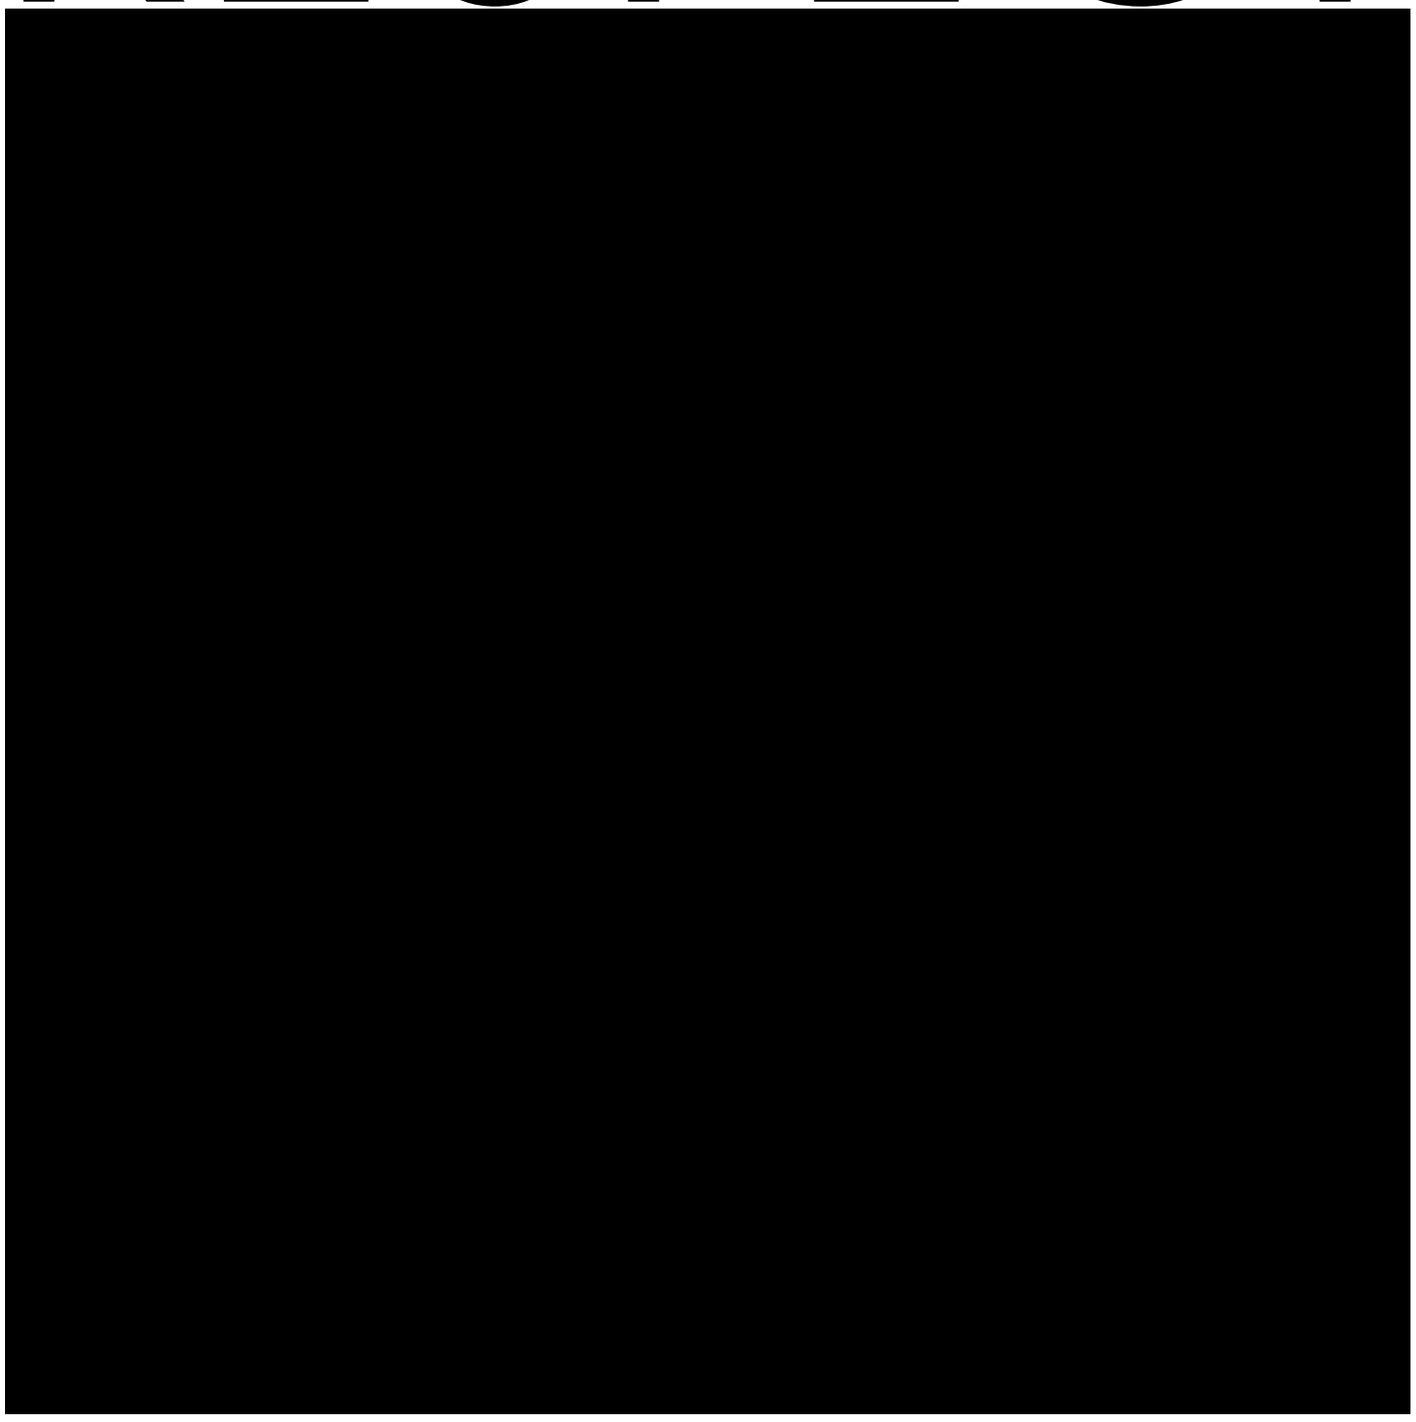

Daniel W. Harris, Hunter College

Points of Agreement

The semantics–pragmatics distinction is about the two parts of an explanation of how we successfully interpret speech.

Points of Agreement

Semantic explanations appeal to hearers' language-specific competencies.

Points of Agreement

Pragmatic explanations appeal to hearer's language-independent competencies.

Points of Agreement

Griceans have gotten some details wrong by missing out on semantic mechanisms they didn’t know about.

Points of Disagreement

How much of interpretation relies on pragmatic, language-independent mechanisms?

Ernie & Matt:
Way less than everyone thinks!

Me:
Way more than Ernie & Matt admit!

Points of Disagreement

Do pragmatic mechanisms aim at a single, correct interpretation?

Ernie & Matt:

Generally not.

"Any conclusions the audience…discovers [via pragmatic reasoning] are implicit and tentative suggestions, rather than transparent and public contributions" (39).

Points of Disagreement

Do pragmatic mechanisms aim at a single, correct interpretation?

Me:

All the damn time!

Points of Disagreement

Are there conversational implicatures whose interpretation we should explain in a roughly Gricean way?

Ernie & Matt:

No.

"We have no use for a category of conversational implicatures, as traditionally and currently understood" (6).

Points of Disagreement

Are there conversational implicatures whose interpretation we should explain in a roughly Gricean way?

Me:

Most definitely.

Points of Disagreement

How best to characterize the semantics–pragmatics boundary?

Ernie & Matt:

Interpretation mostly relies on conventional, linguistic mechanisms. Once in a while, it falls back on intention recognition.

Points of Disagreement

How best to characterize the semantics–pragmatics boundary?

Ernie & Matt:

"Explicit reasoning about what the speaker might have wanted or believed comes in, if at all, only when things go wrong" (13).

Points of Disagreement

How best to characterize the semantics–pragmatics boundary?

Me:

Interpretation is always intention recognition. Semantic mechanisms make the inferences easier by cutting down on the interpretive options.

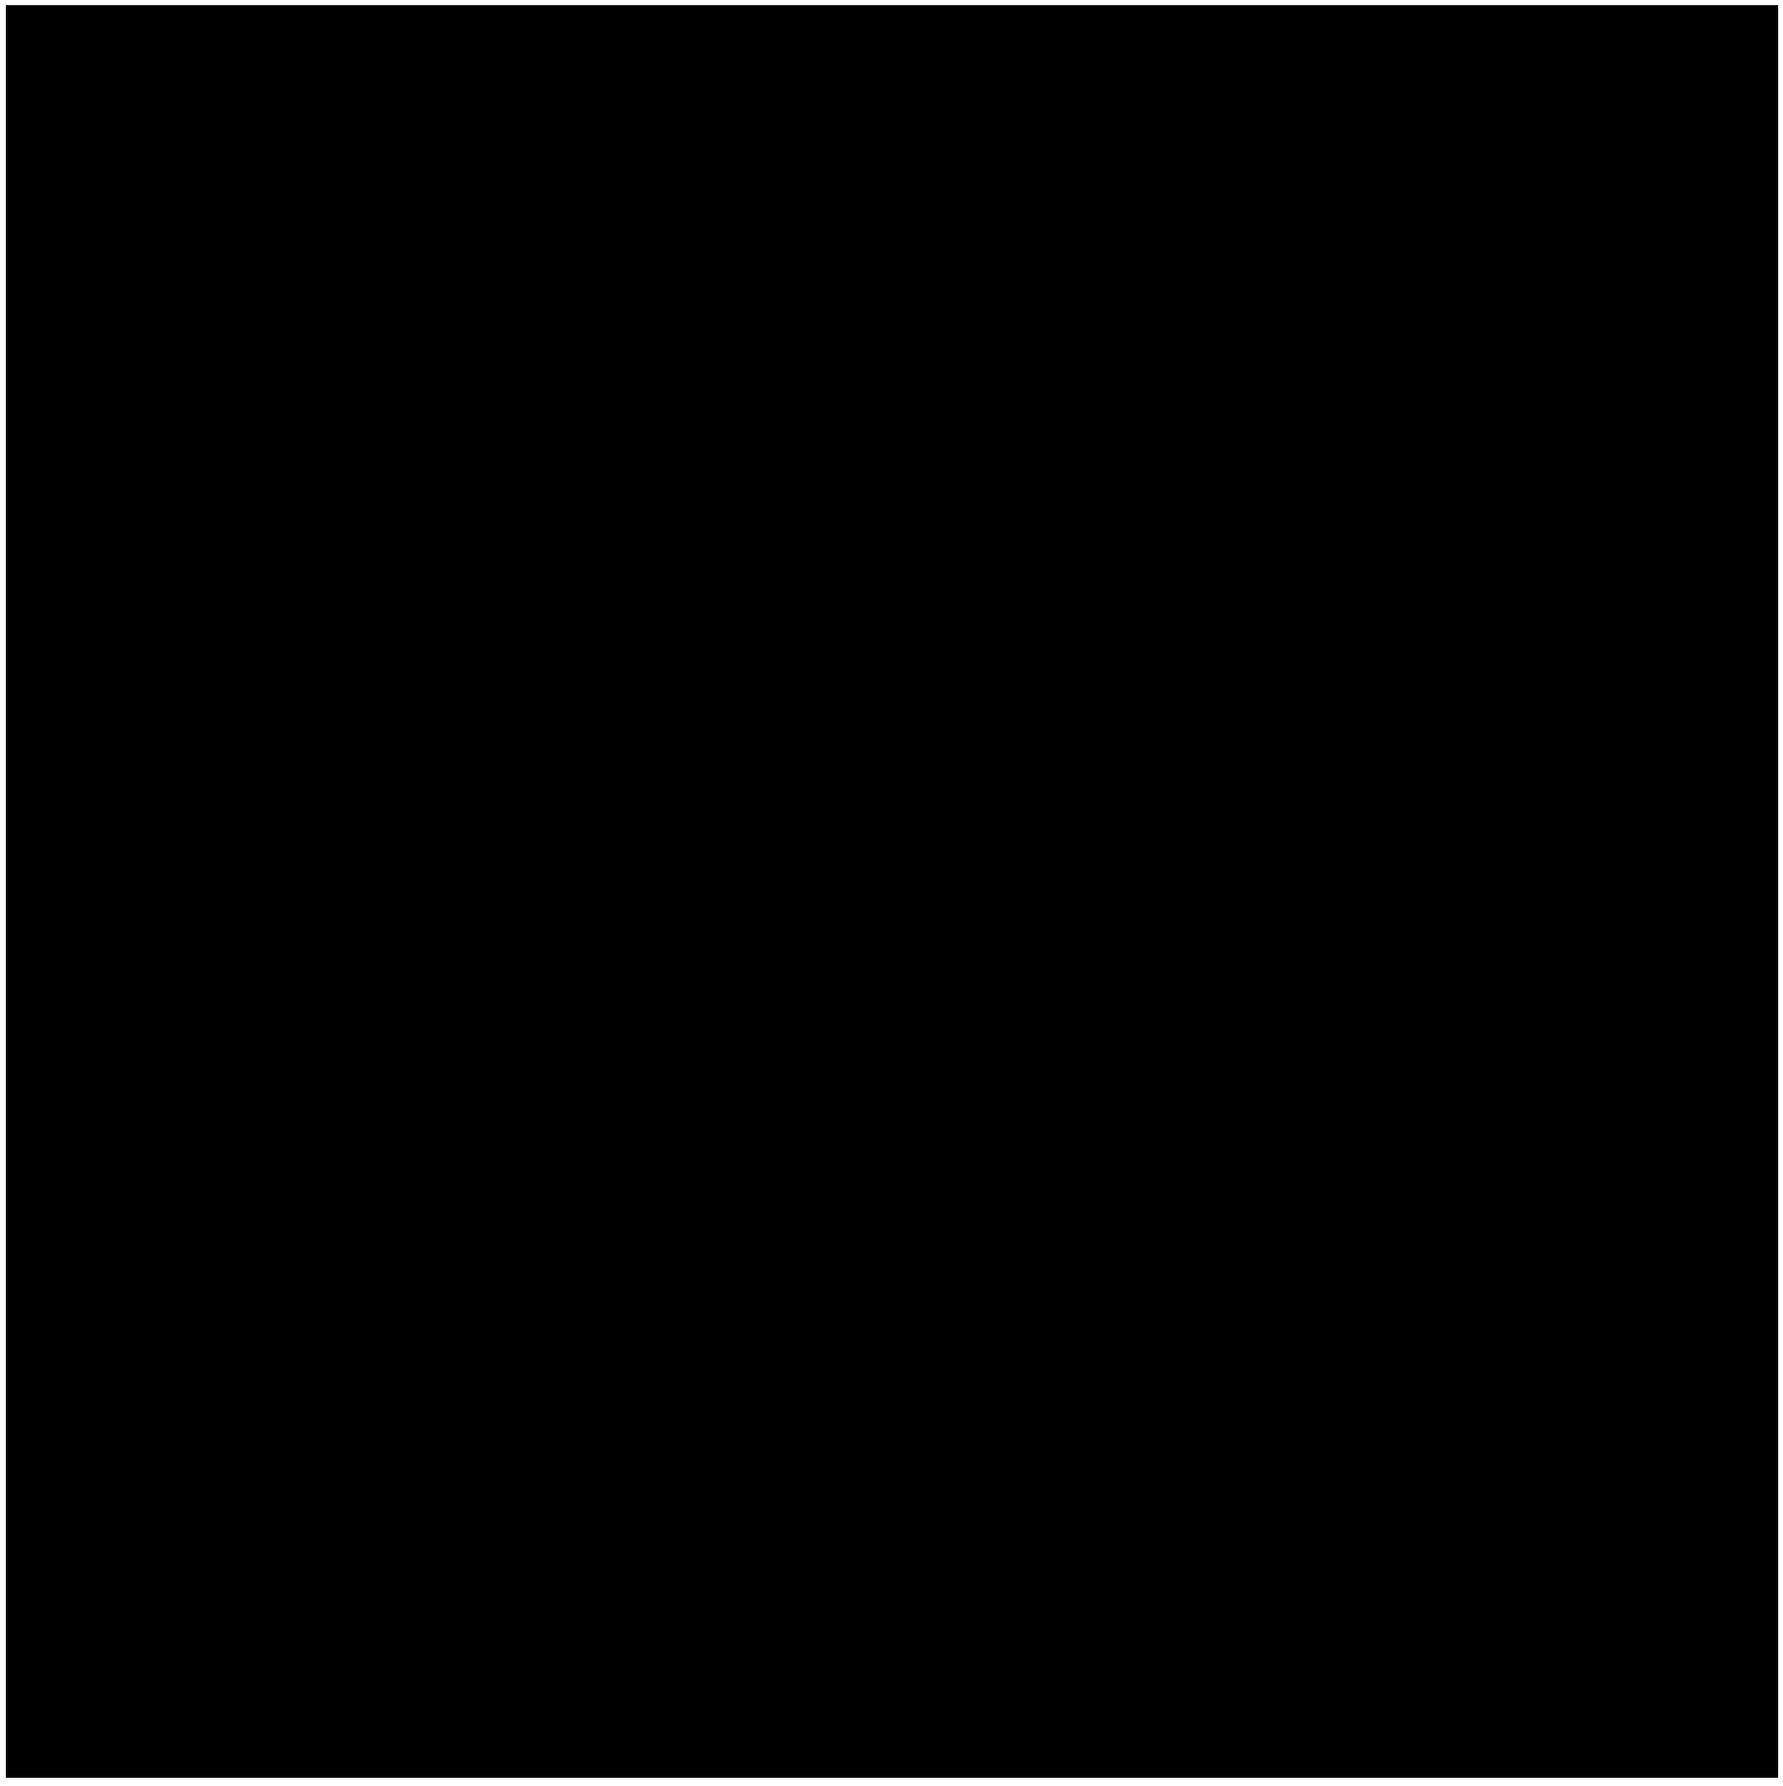

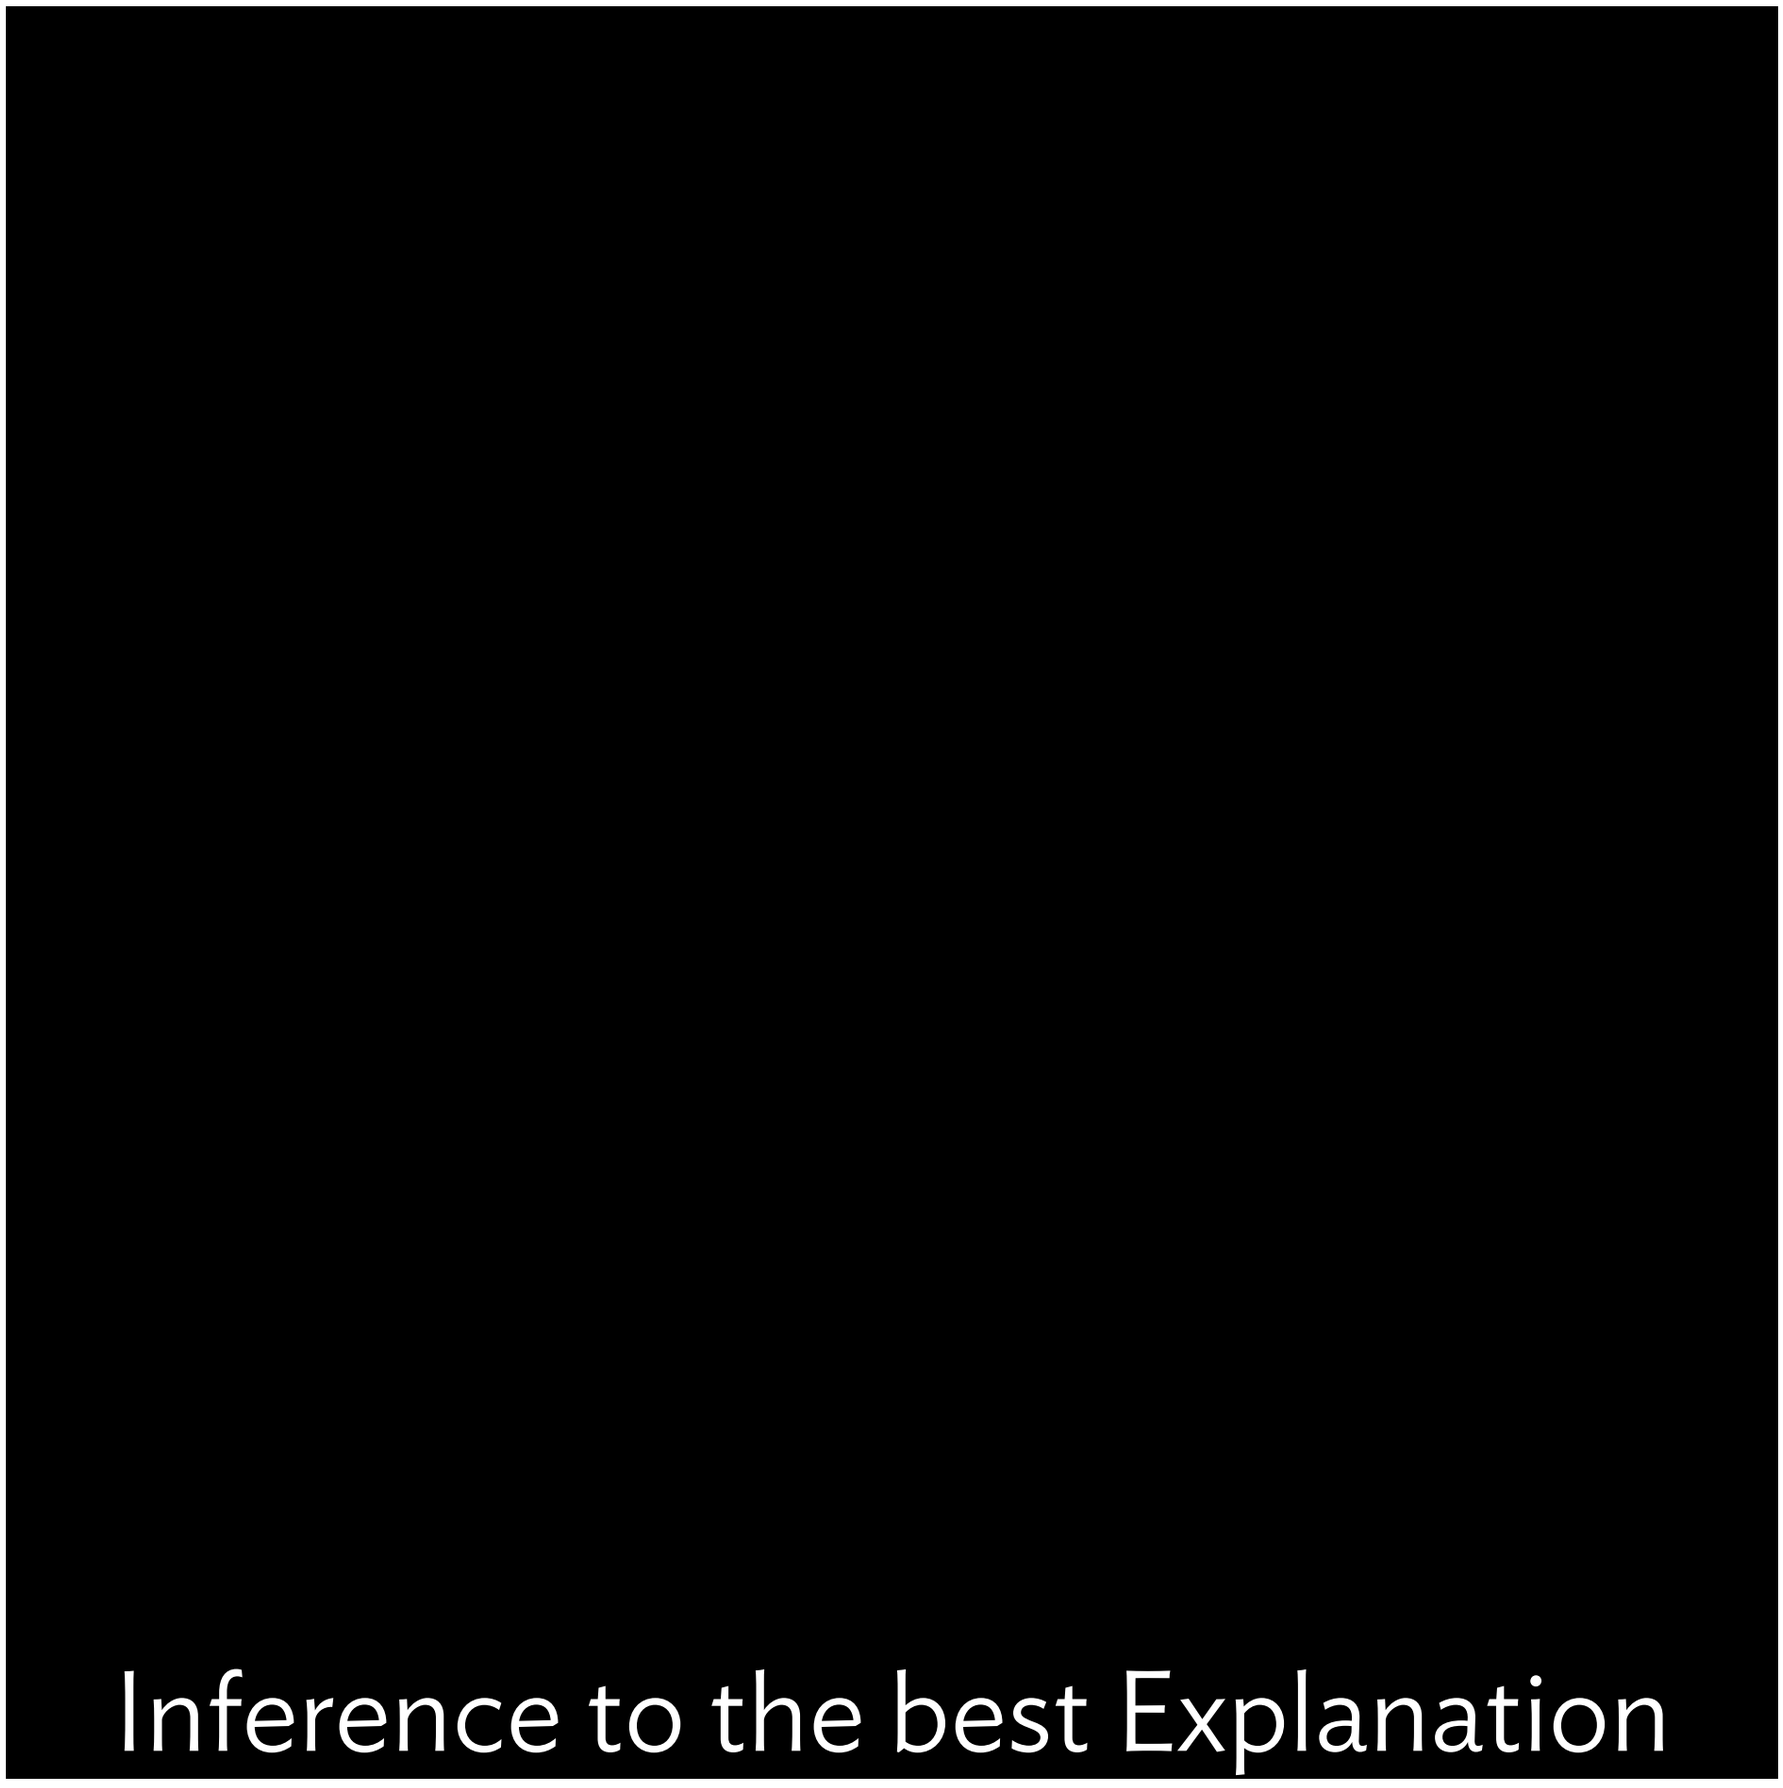

Inference to the best Explanation

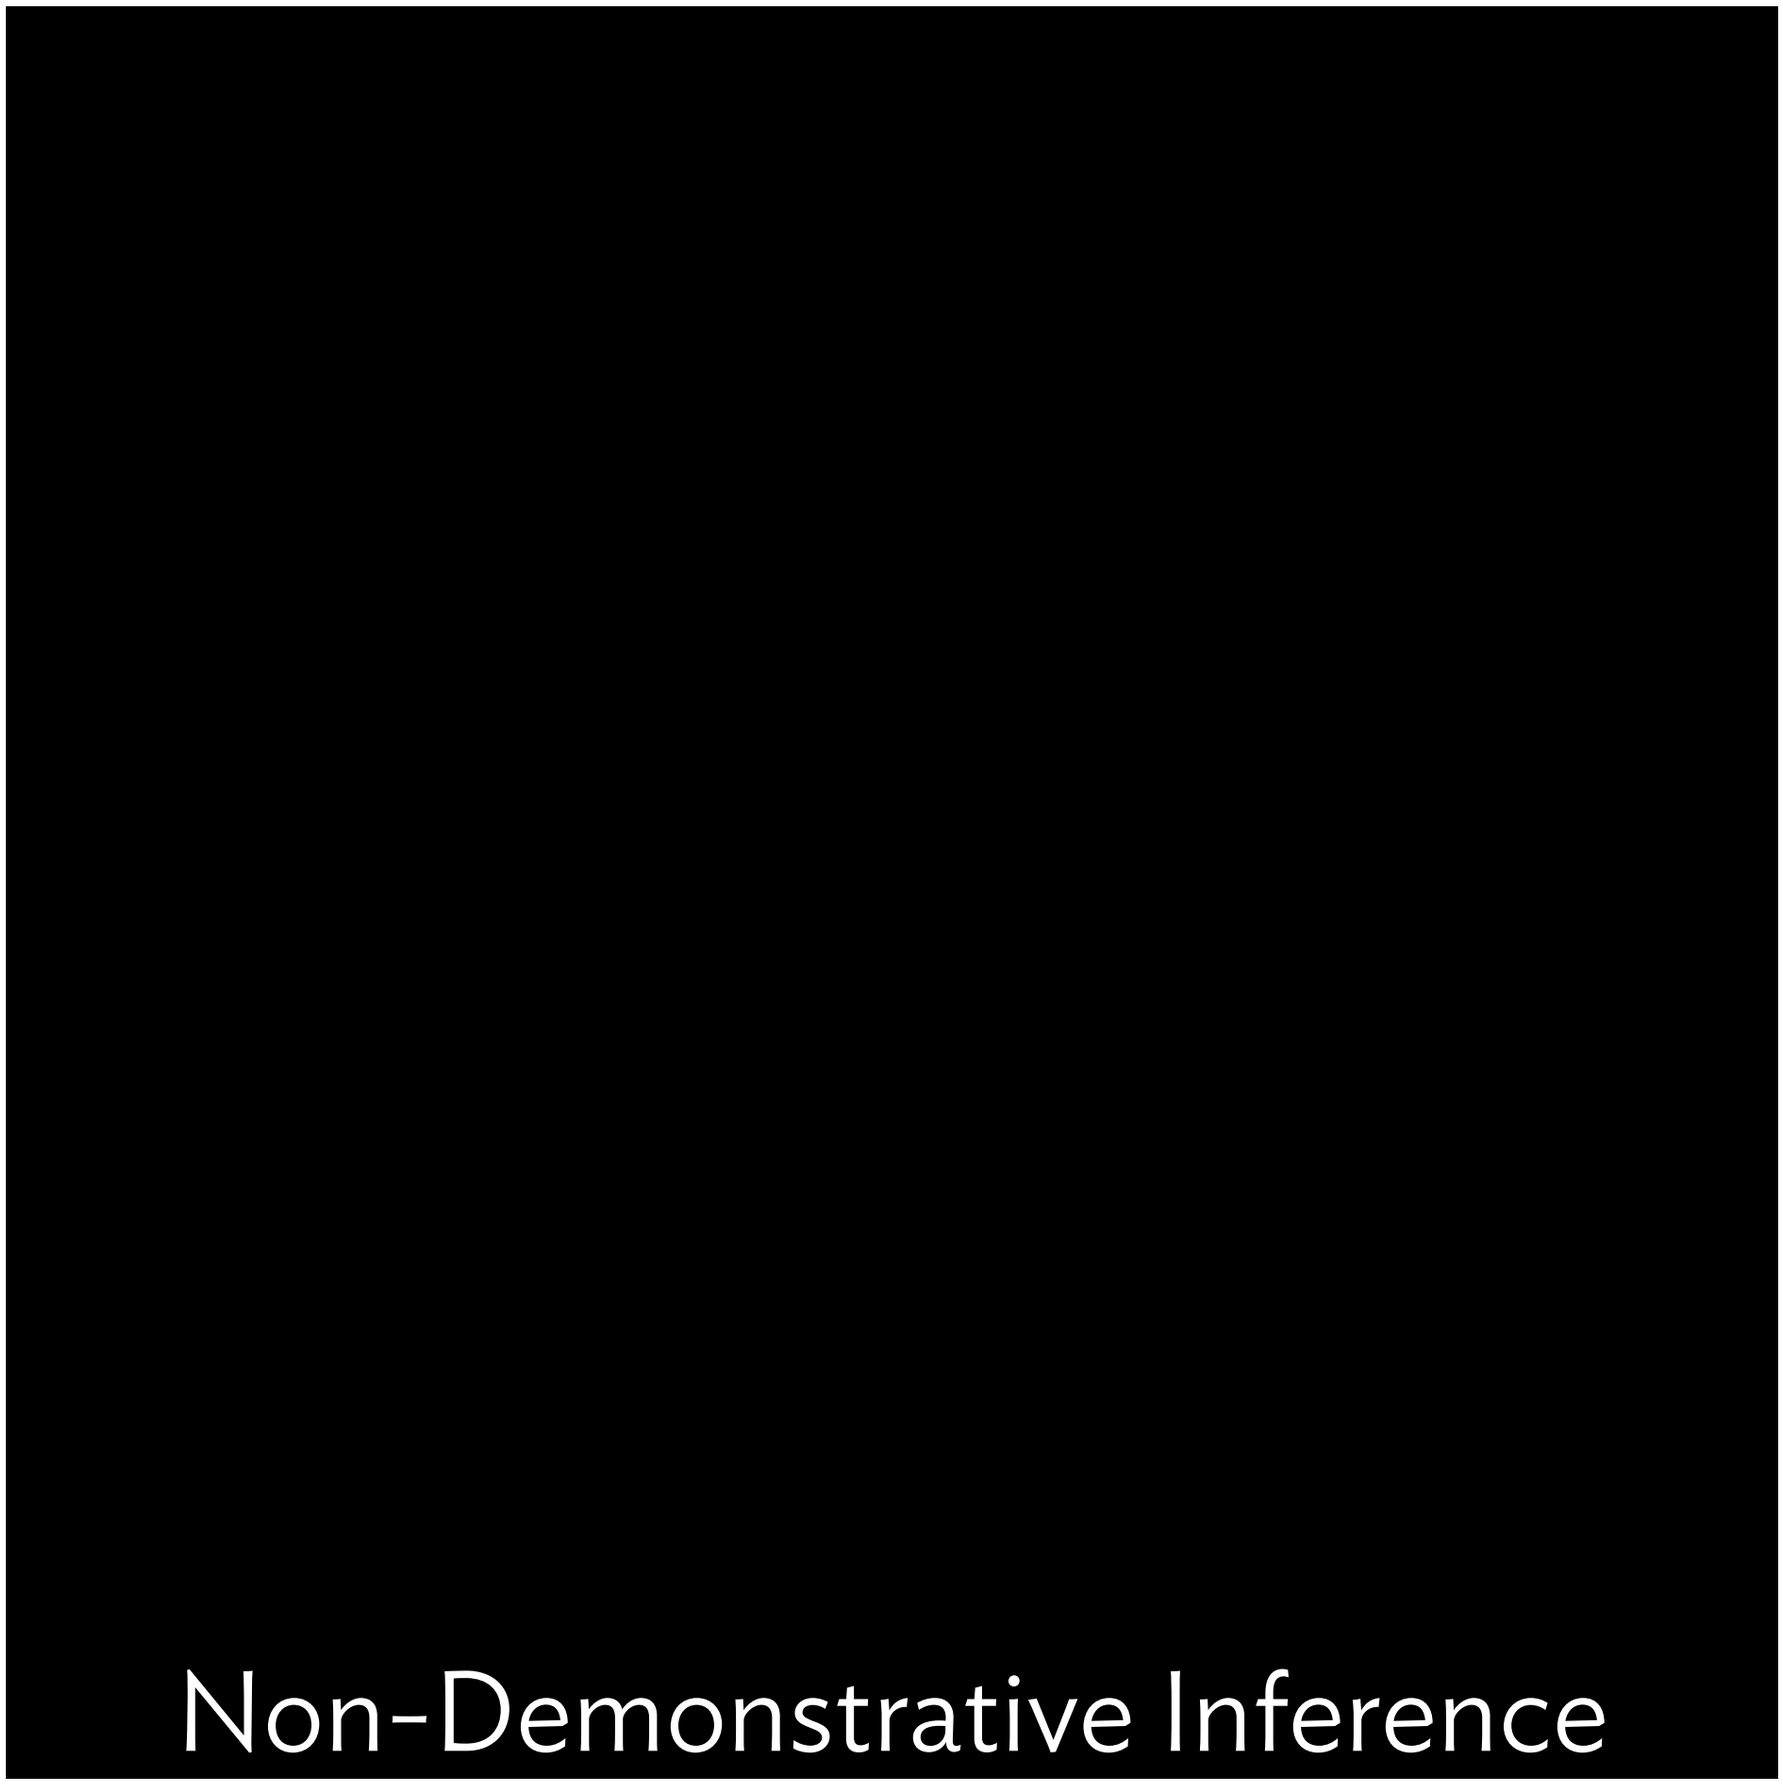

Non-Demonstrative Inference

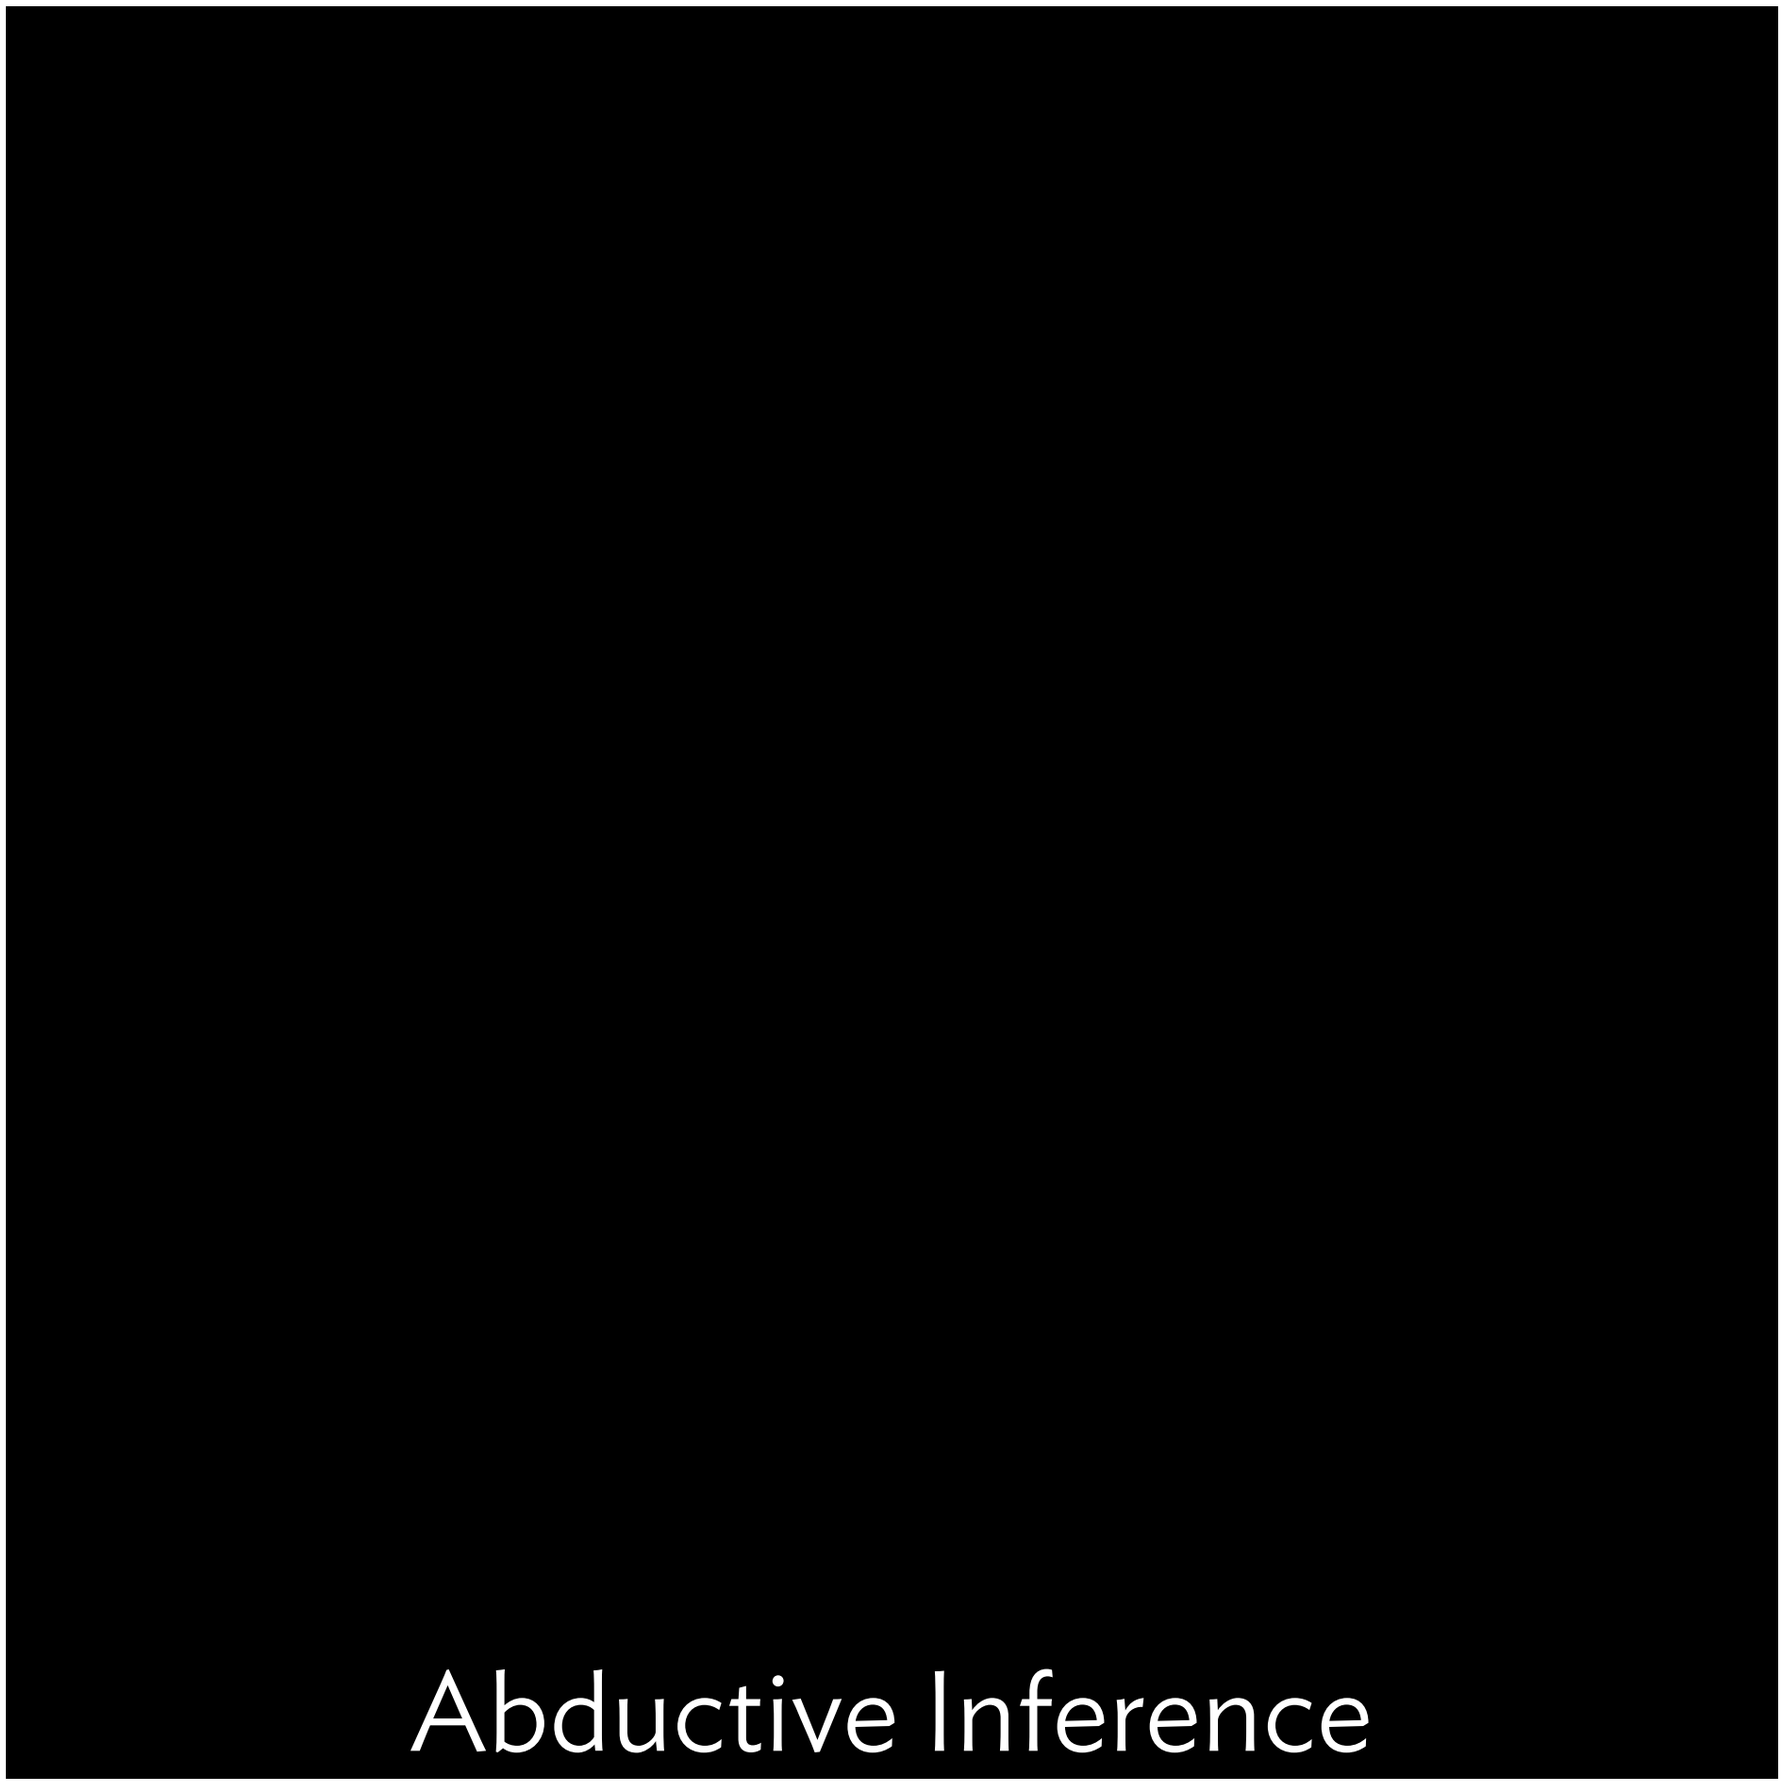

Abductive Inference

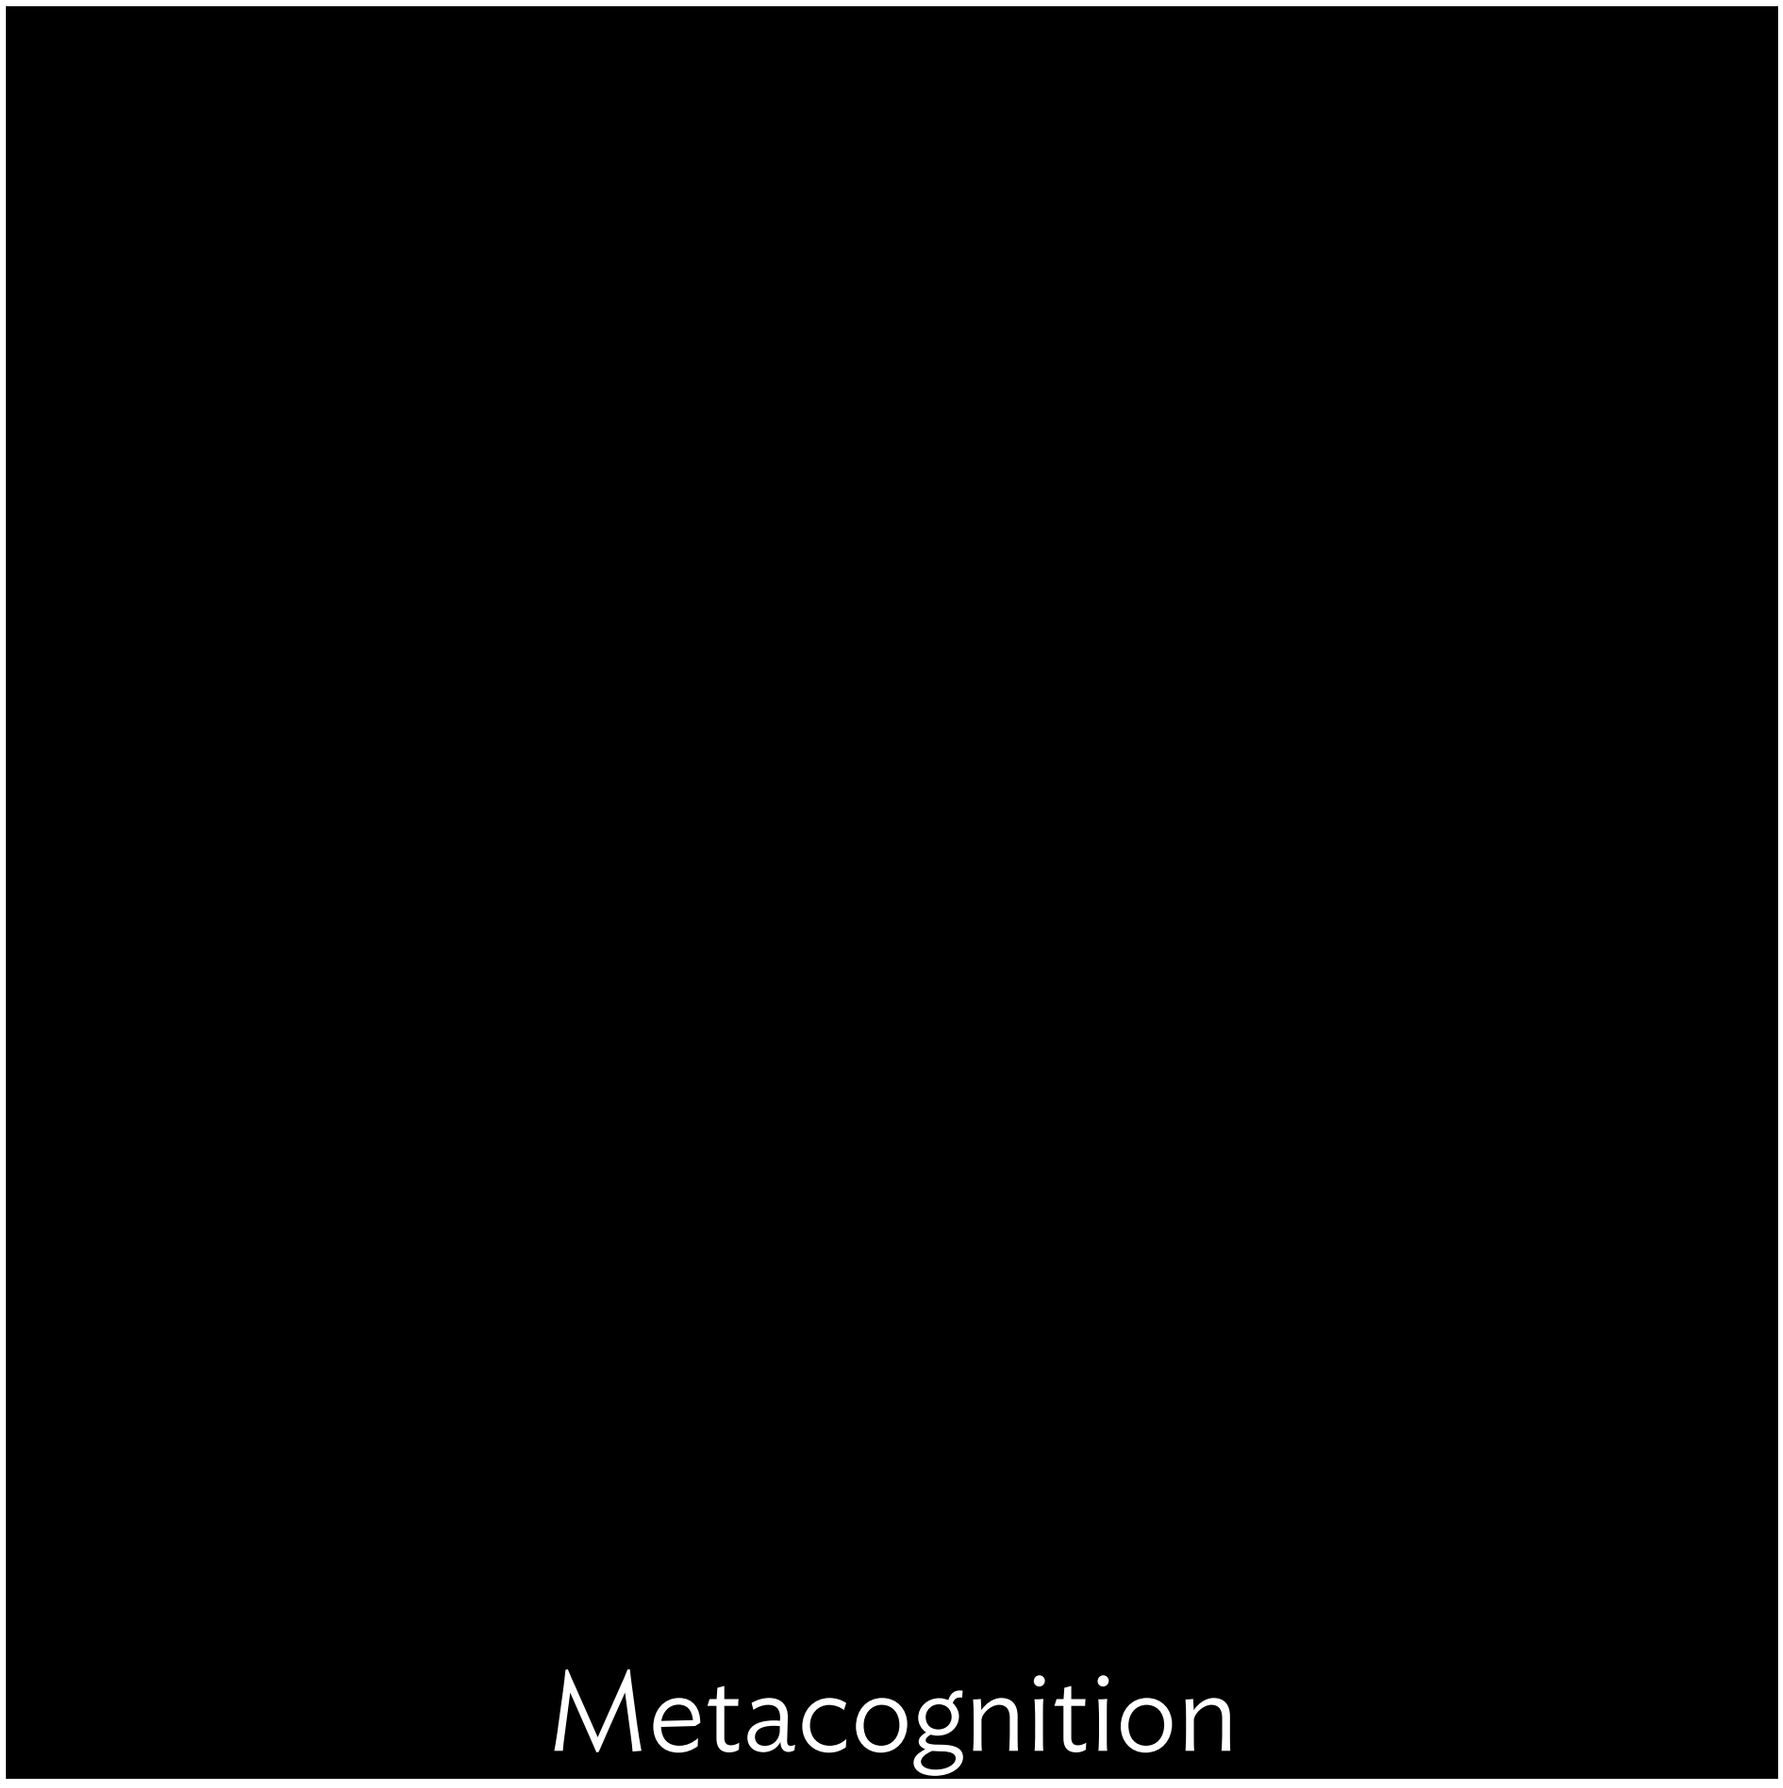

Metacognition

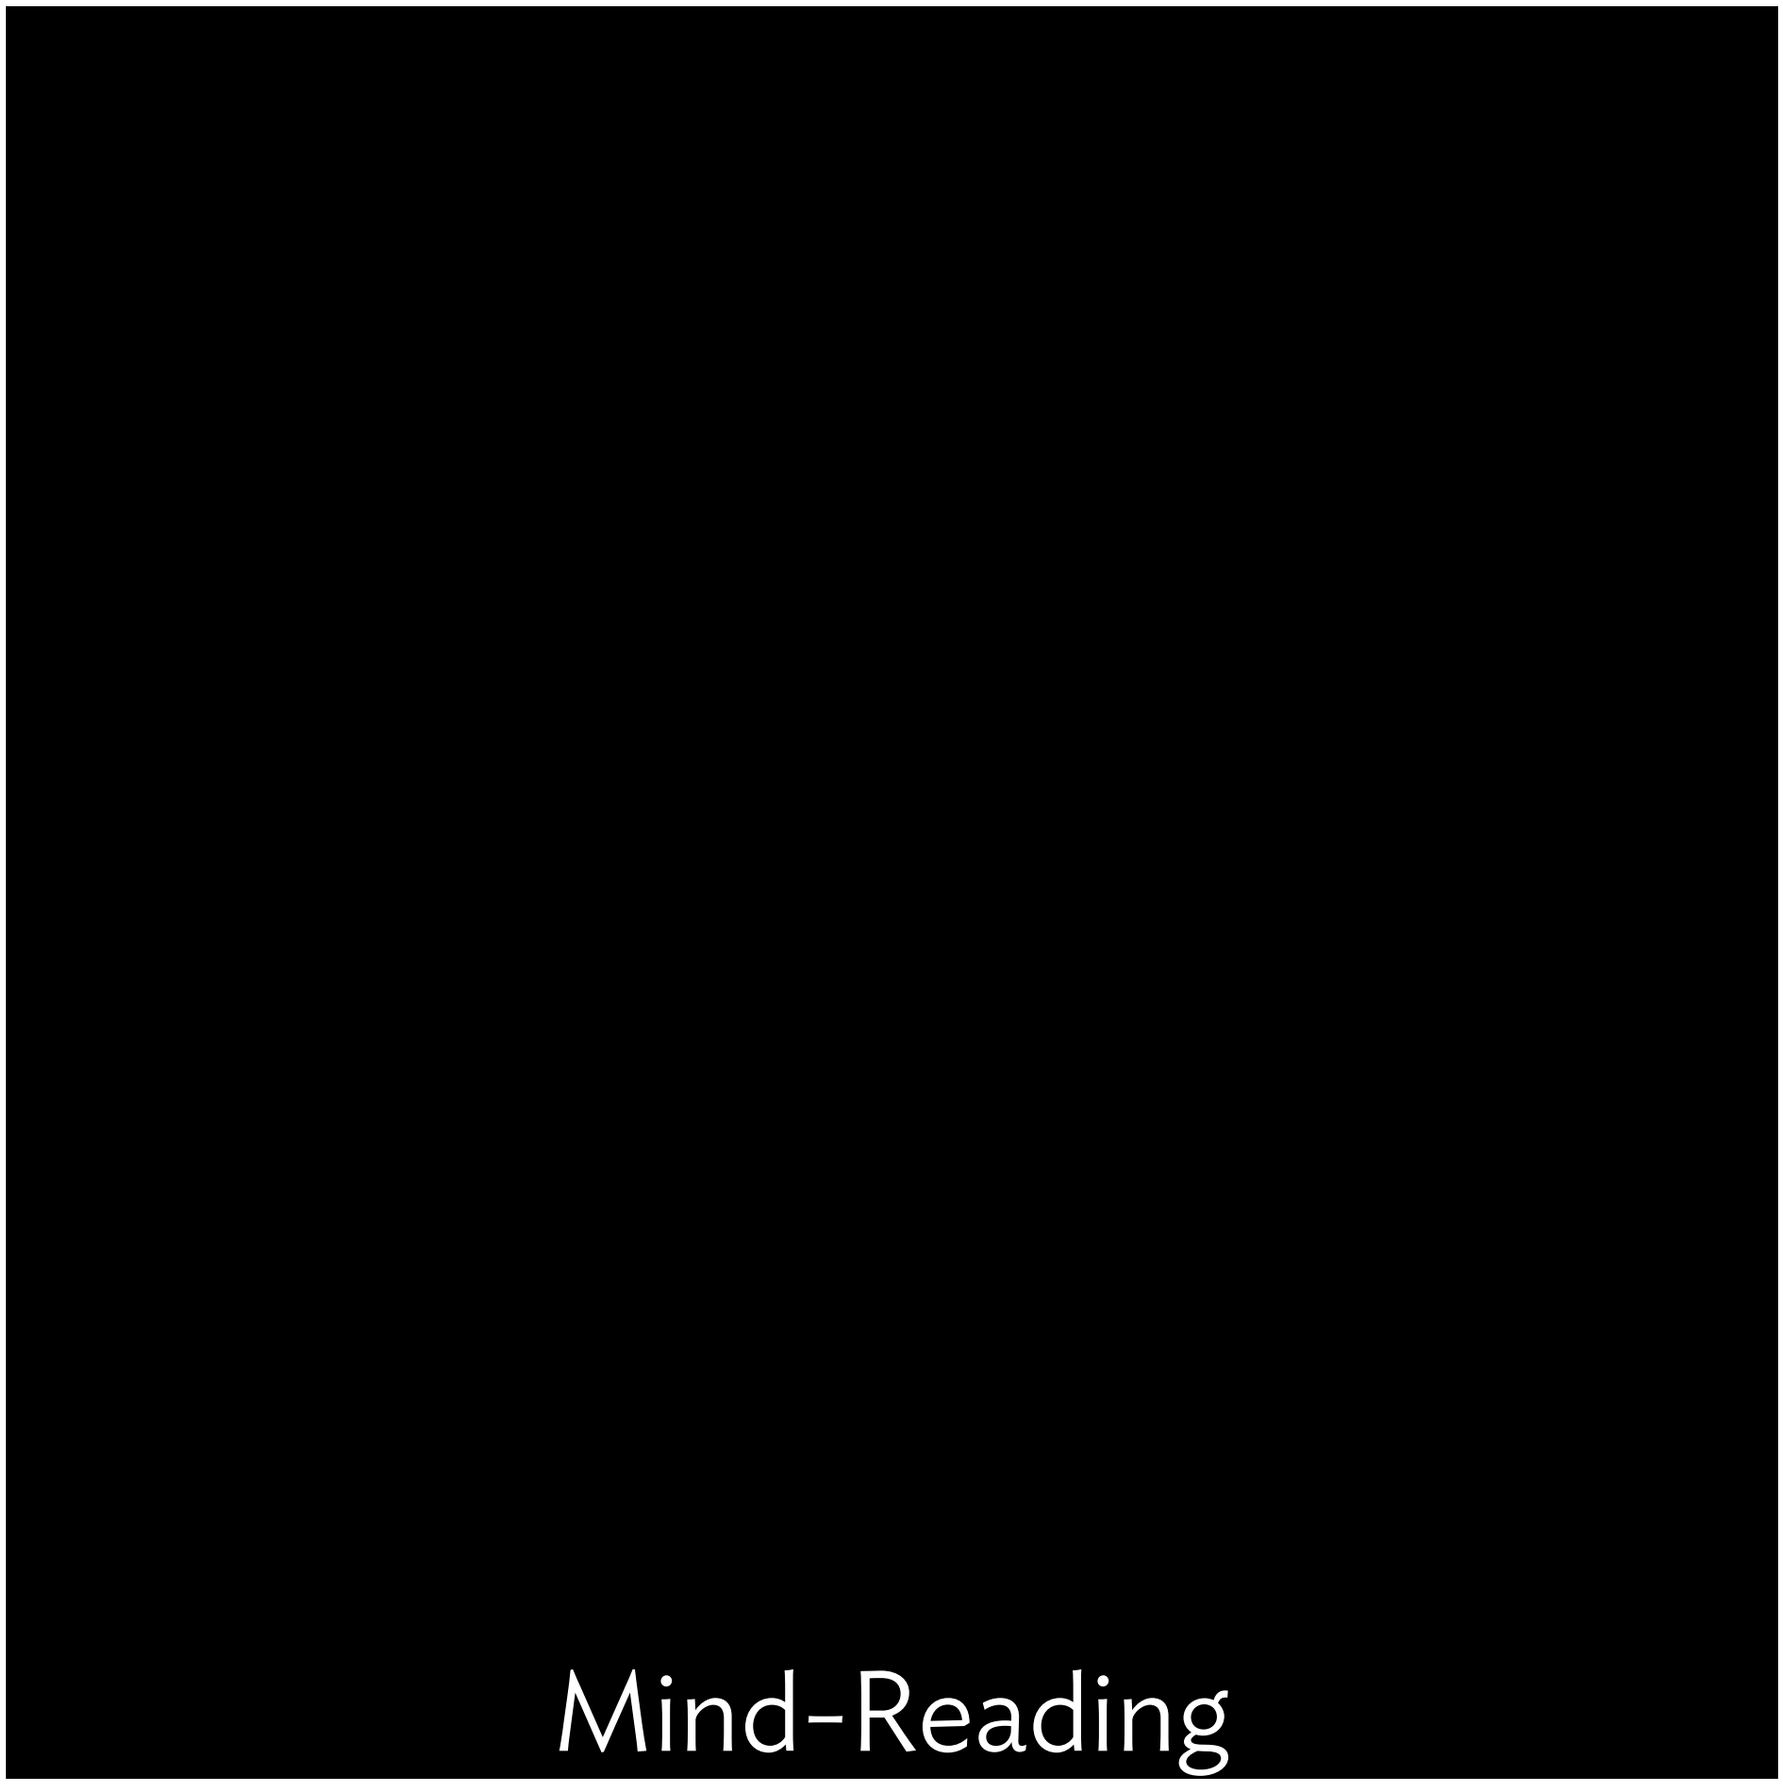

Mind-Reading

Fodor's First Law of the Nonexistence of Cognitive Science

"...the more global (e.g. the more isotropic) a cognitive process is, the less anybody understands it."

—Fodor 1983: 107

Corollary

Pragmatics is hard.

To interpret is to provide an *explanation*, and the concept of interpretation makes no sense in the absence of a *problem* to be solved. We reflexively generate *hypotheses* about the things we perceive. Nowhere is this more in evidence than when we perceive one another's *actions*. We act out of *reasons*. To interpret an action is to form a hypothesis about the *intentions* behind it, the intentions that *explain* it. Interpreting a speech act is a special case of this.

—Neale, 'Pragmatism and Binding', 179

How the audience finds coherence is a flexible and eclectic process. There is no limit to the background knowledge, the familiar patterns, or the explanatory assumptions that the audience can invoke in explaining the speaker's utterance.

—Lepore & Stone, p.90

Is Pragmatics TOO Hard?

Is Pragmatics TOO Hard?

Yes; yes it is.

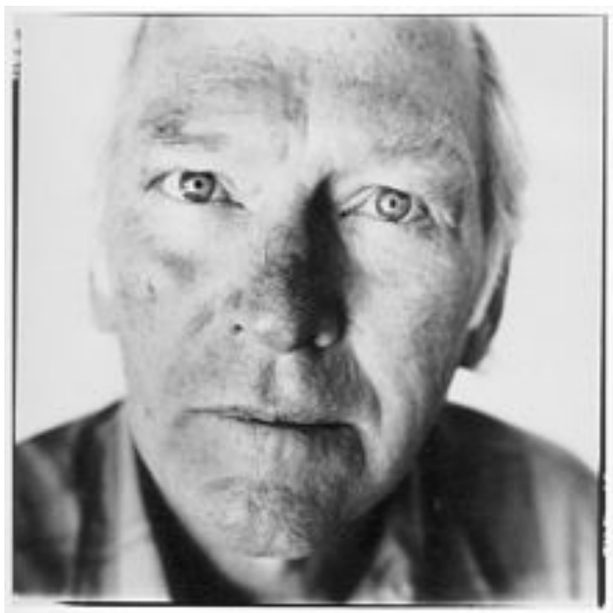

Three Pragmatic Strategies

Idealize

Shrink

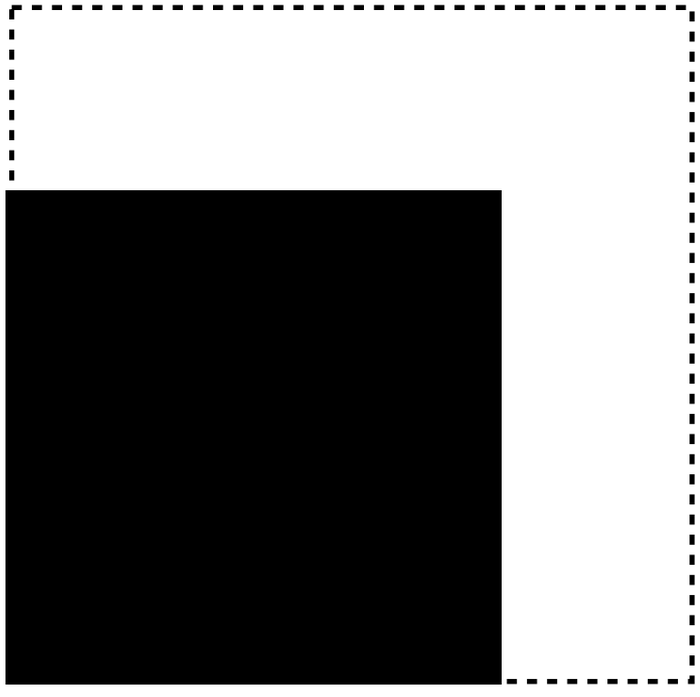

Love

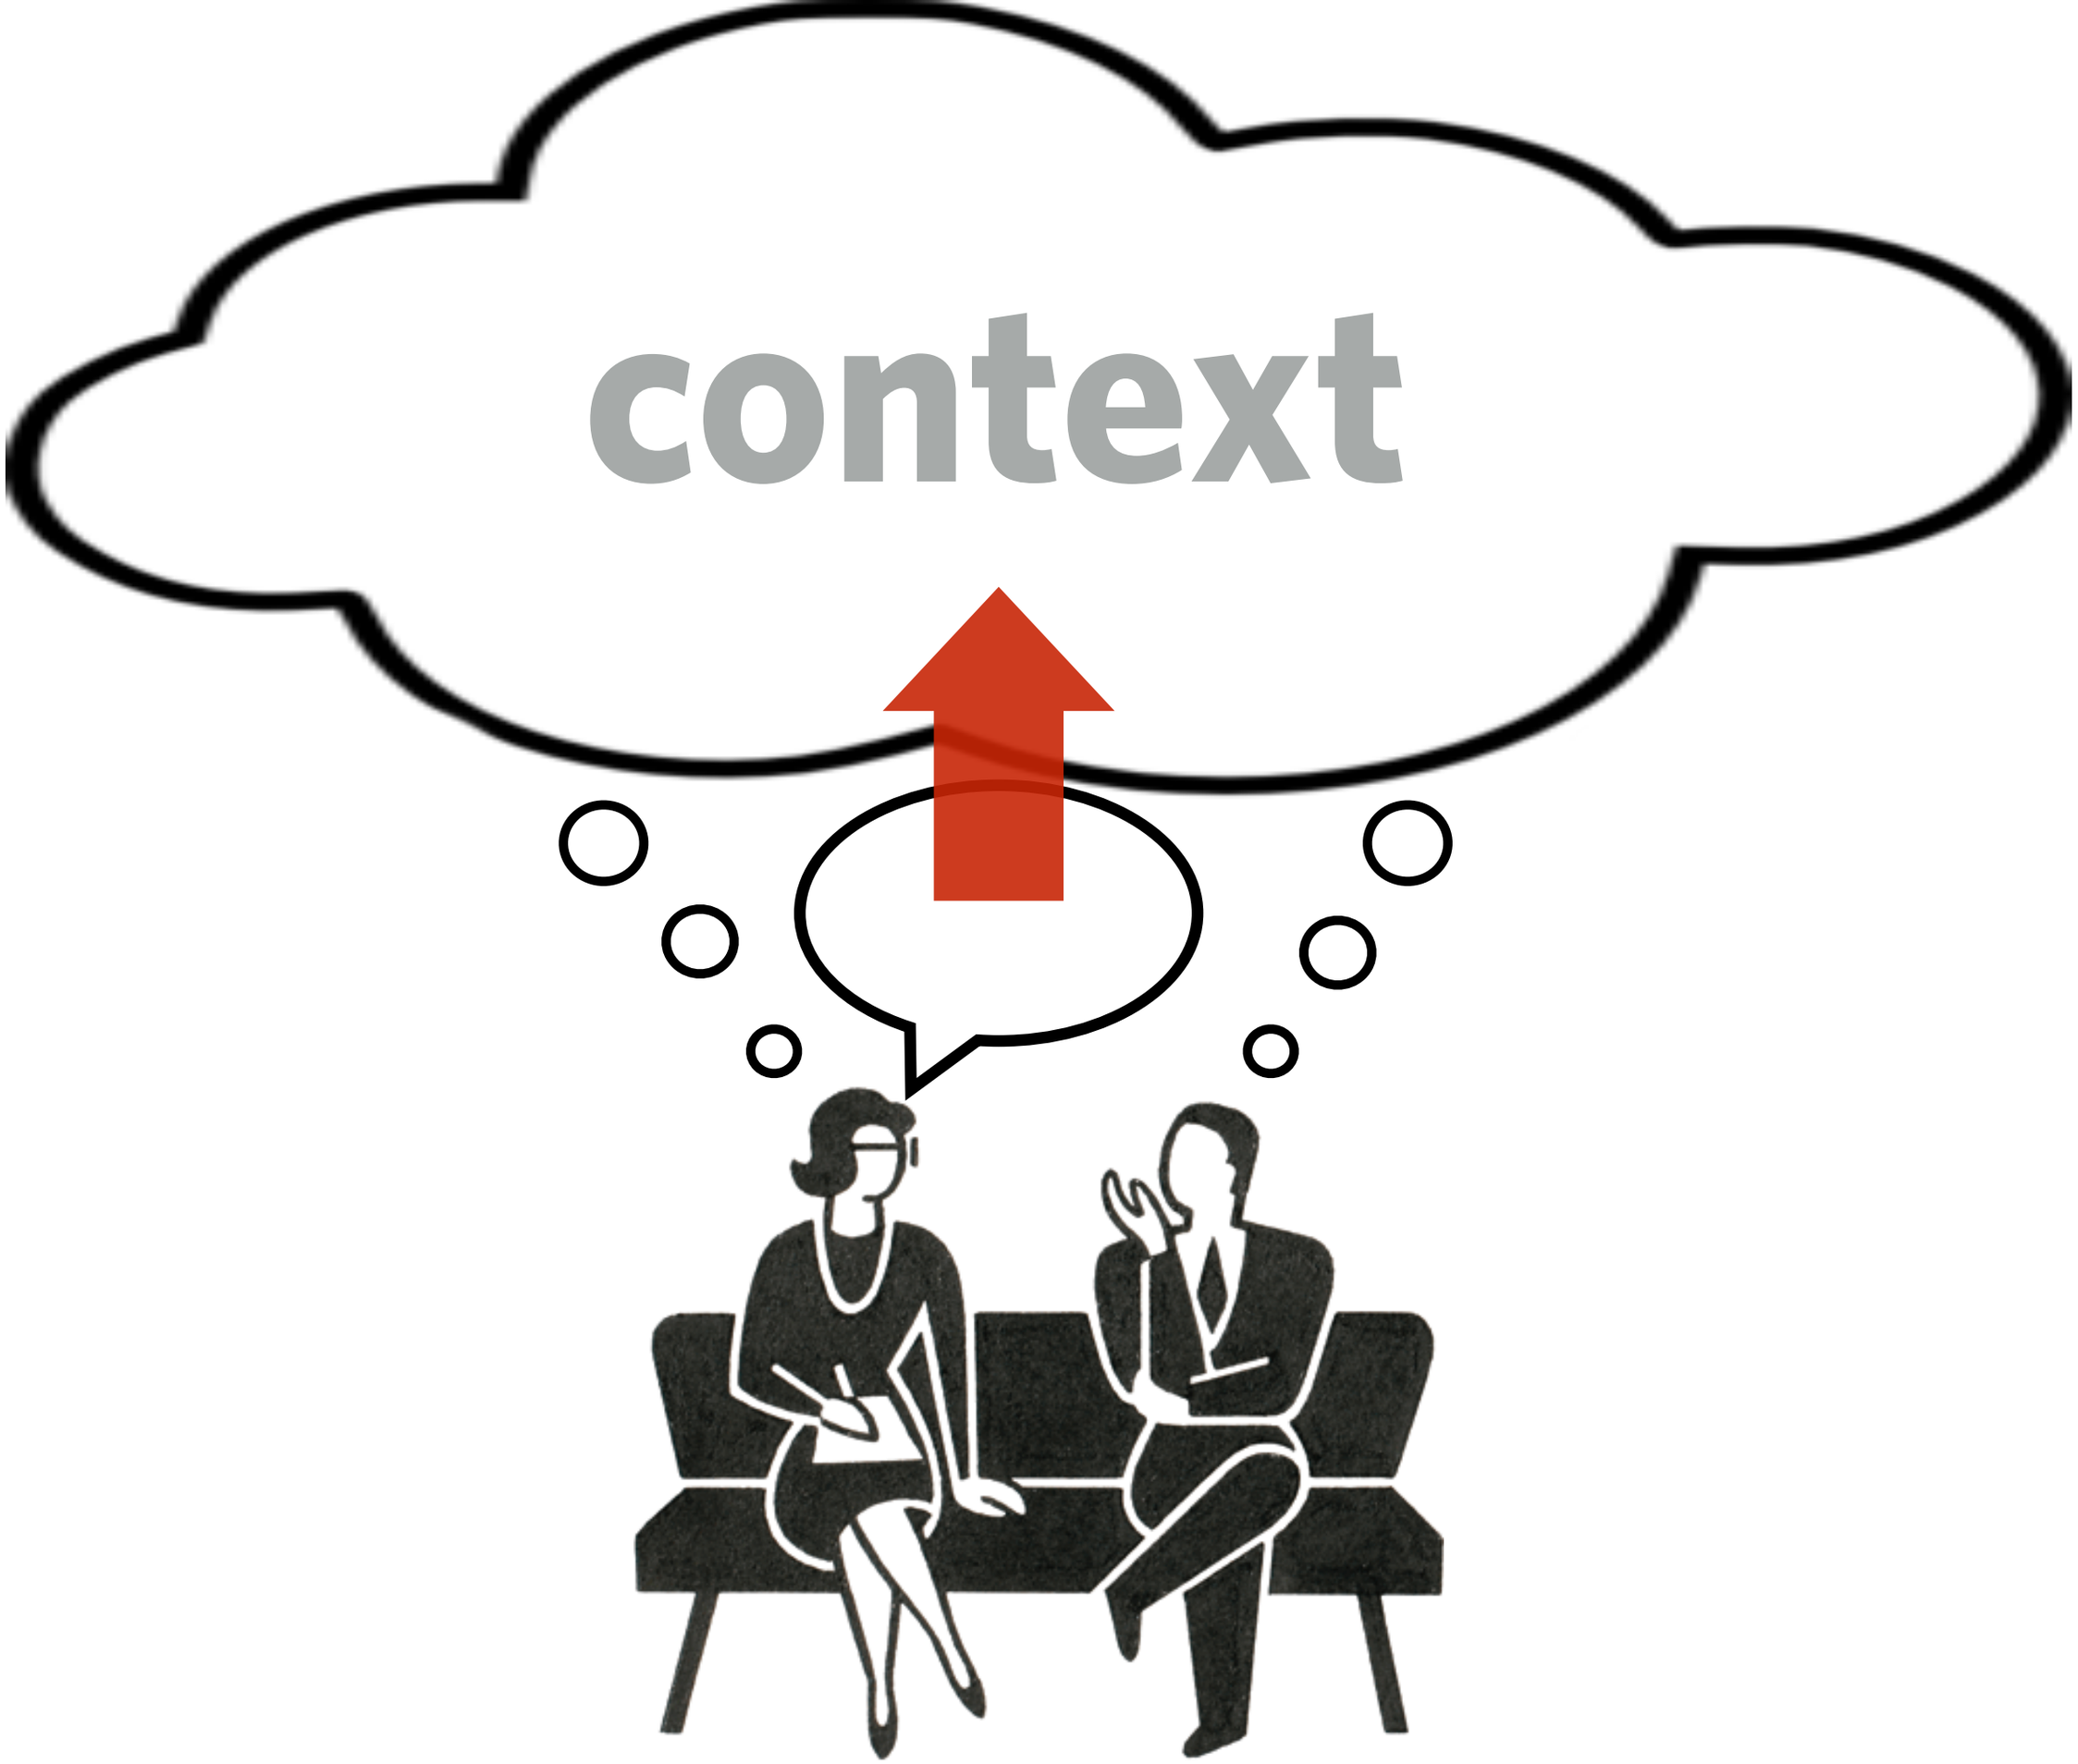
context

p2

p3

p1

p4

p5

p2

p3

p1

p4

~p1

In some dynamic systems, sentence meanings are CCPs—mappings from contexts to contexts.

In order to account for recalcitrant speech acts, this entails building a complete theory of belief and preference update into our semantics.

But: "…it would be a mistake to augment the theory of assertion with the fruits of the epistemological literature on belief-revision."

N. Charlow (2013: §5.6.2)

(archival photo)

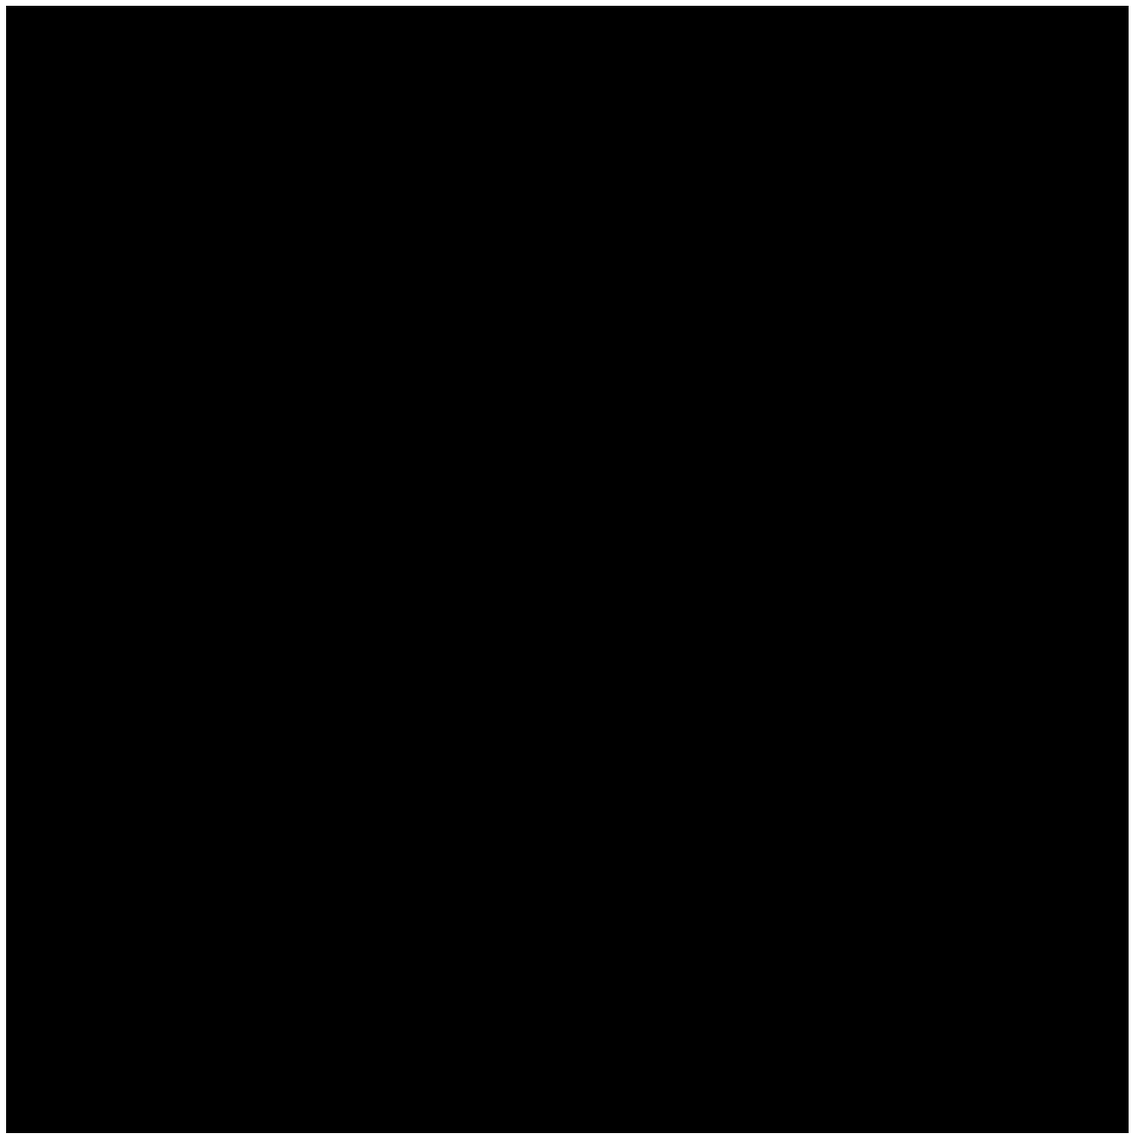

We should explain (most) referential uses of definite and indefinite descriptions by positing a second, referential meaning.

Devitt (2004, etc.)

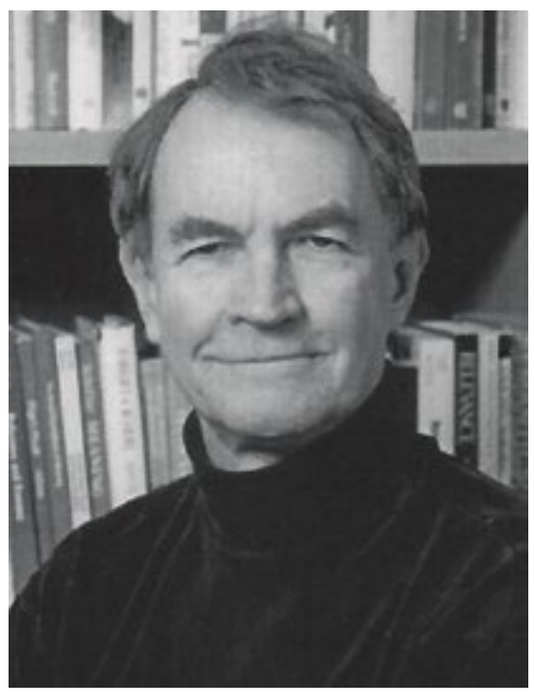

Much more of interpretation than one might have expected is actually encoded in the rules of language.

—Lepore & Stone, p.6

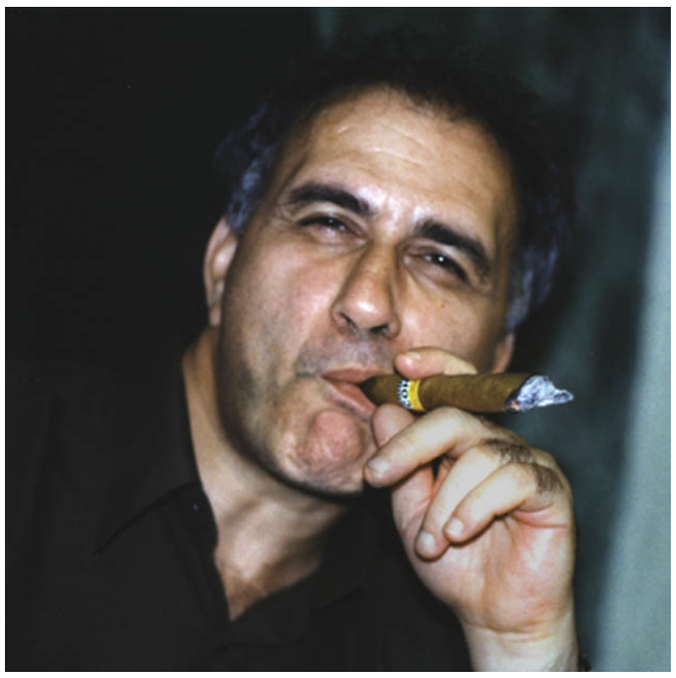

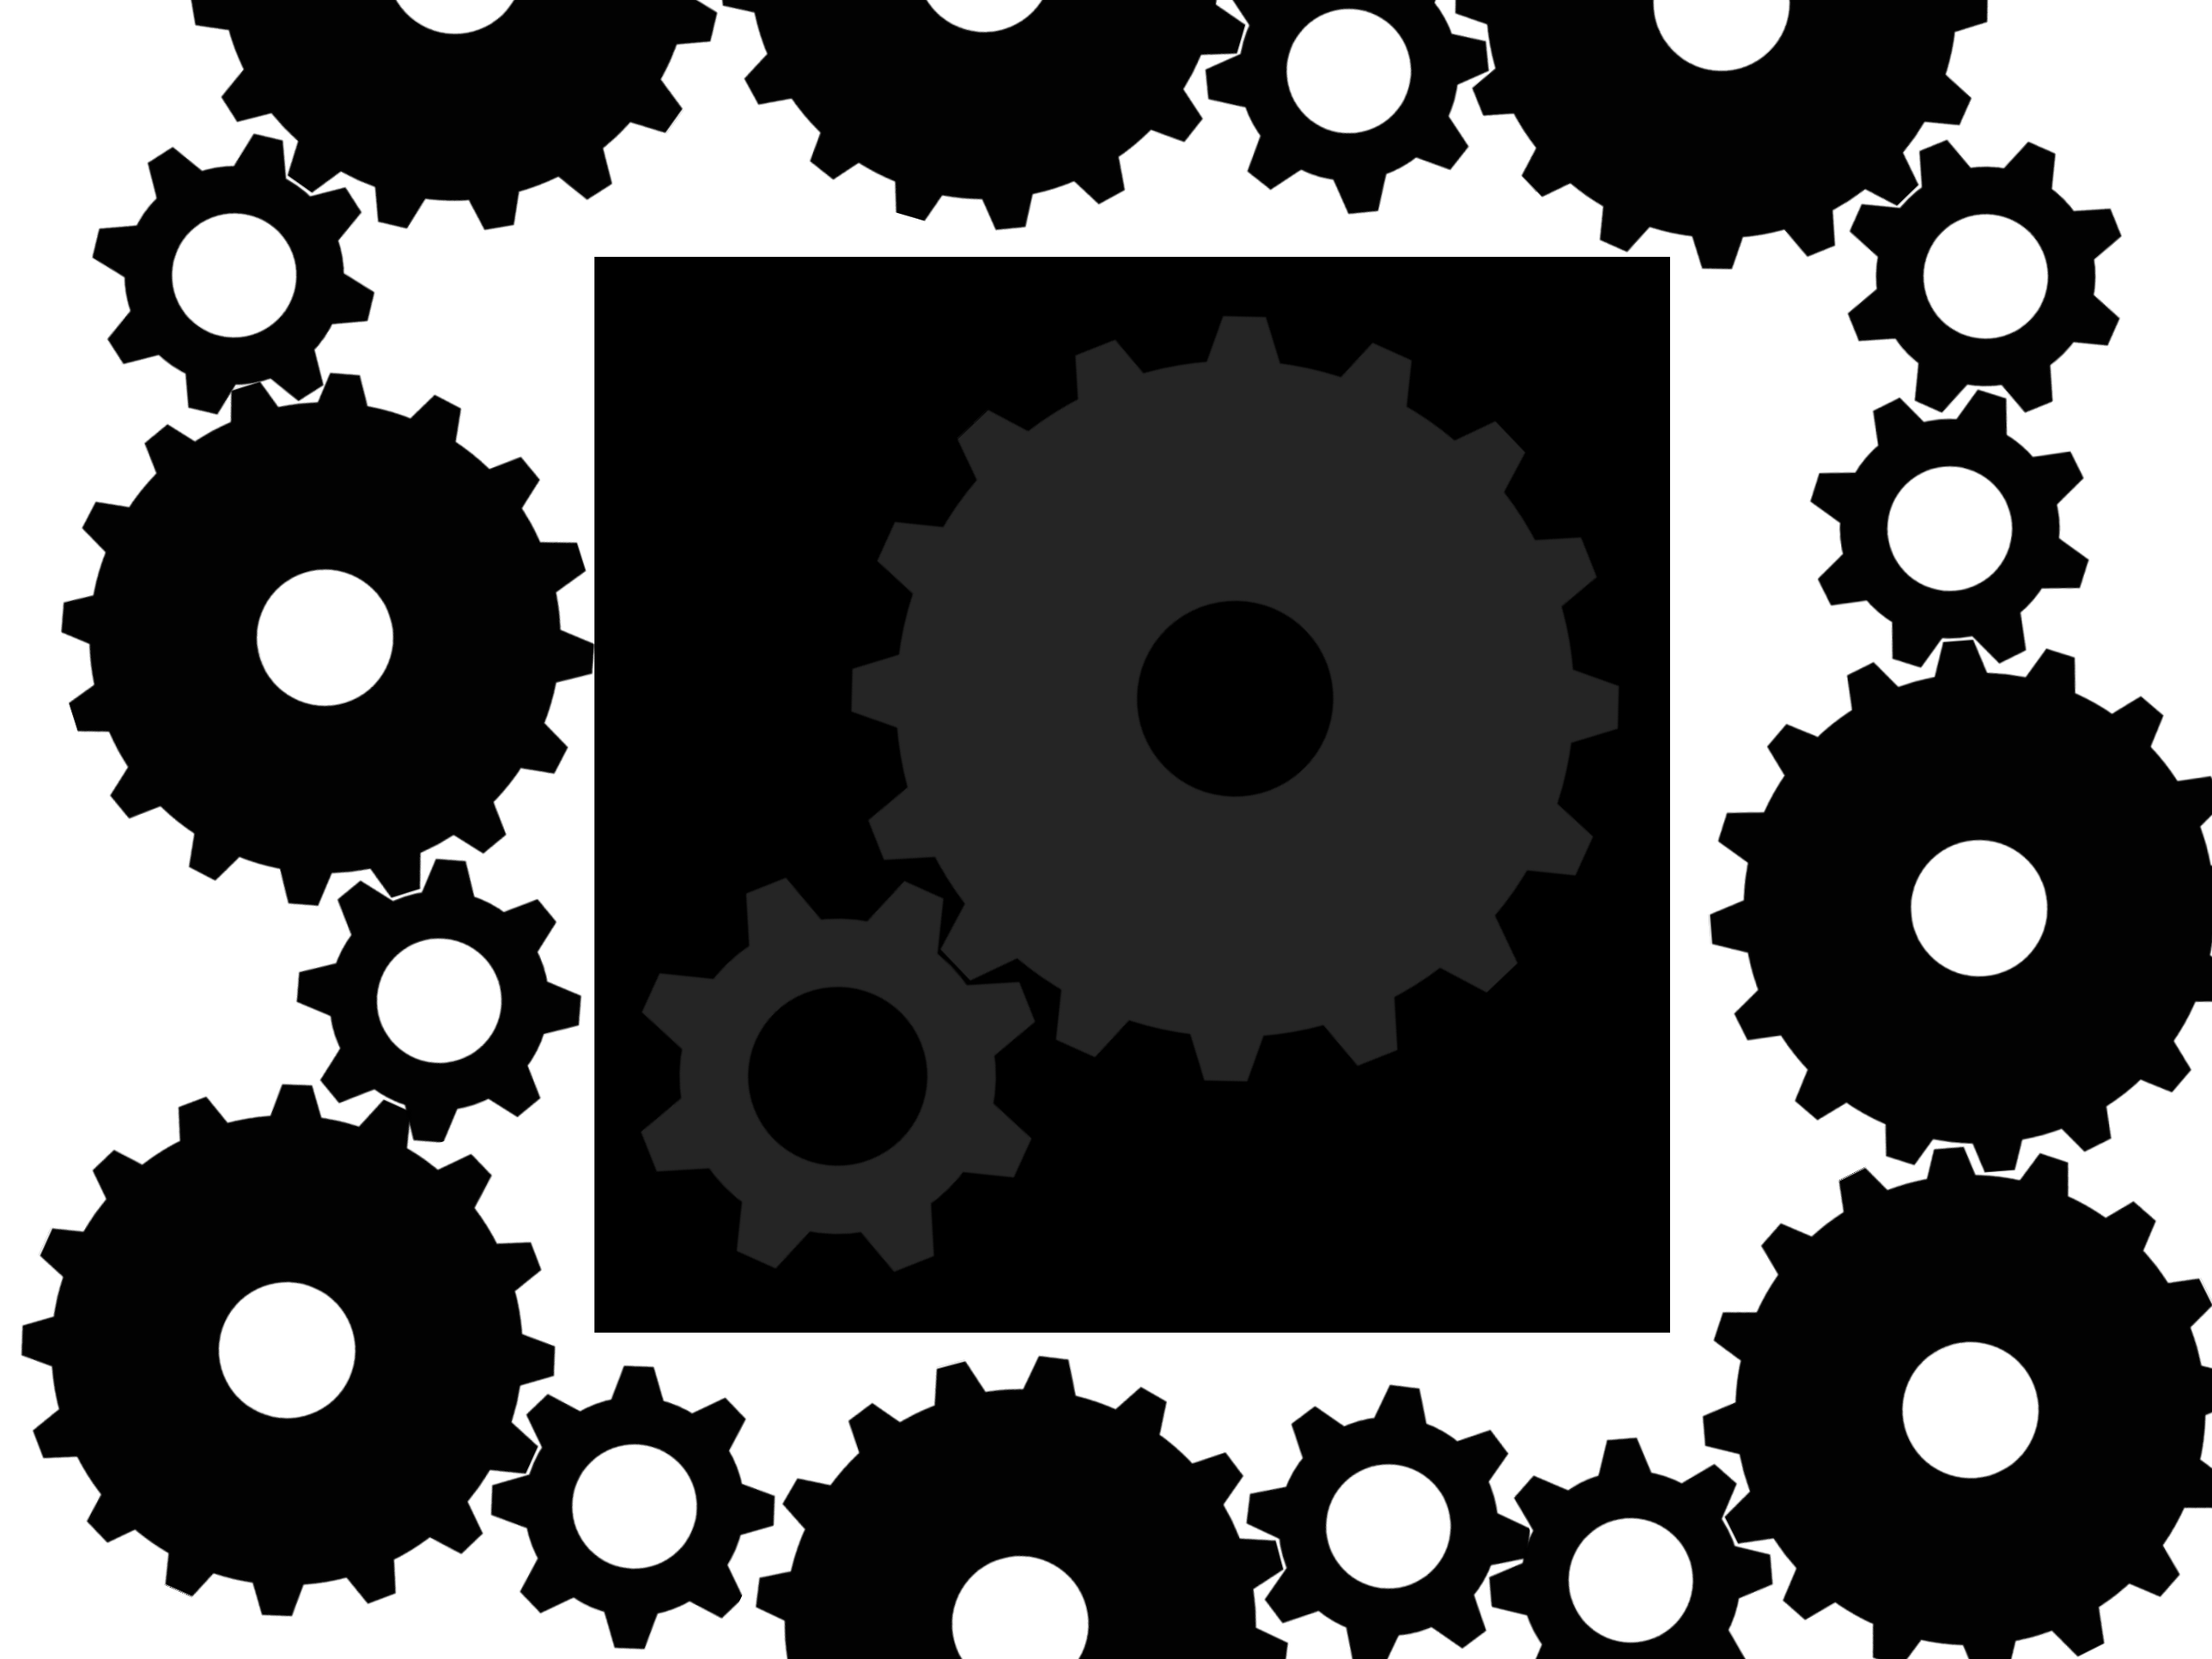

parser

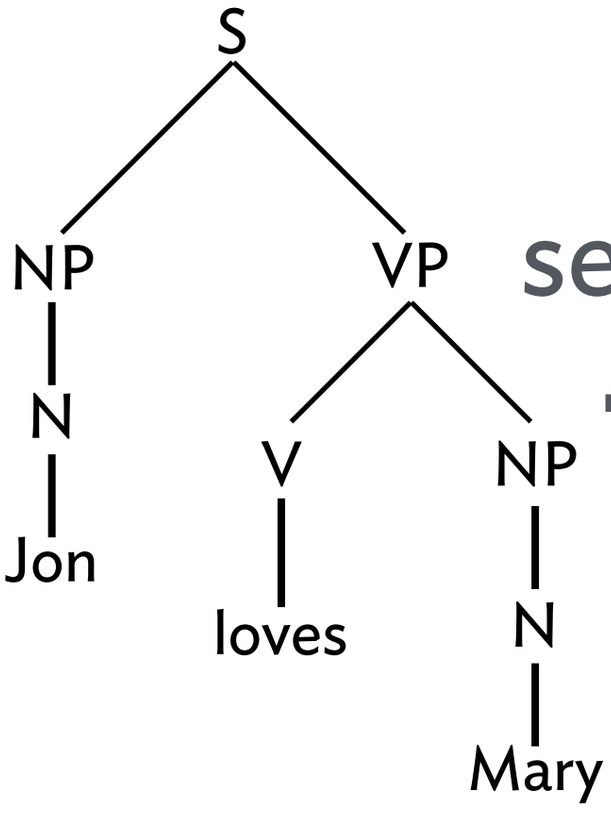

lexical semantics

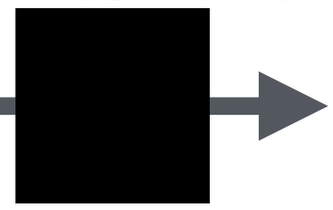

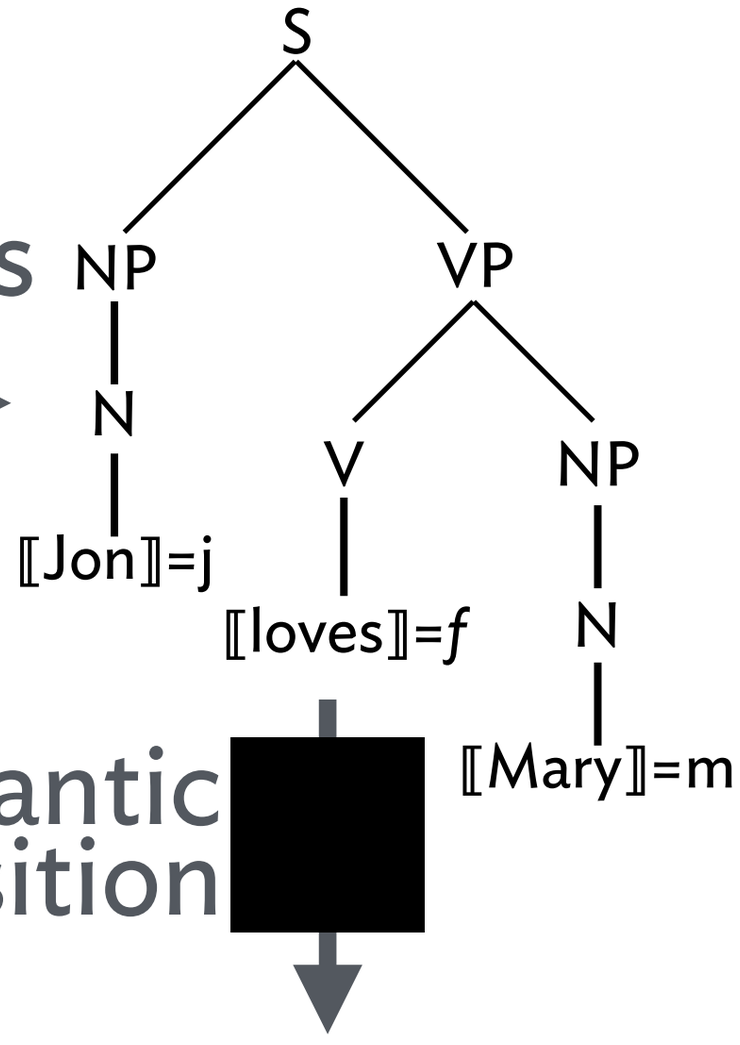

semantic composition

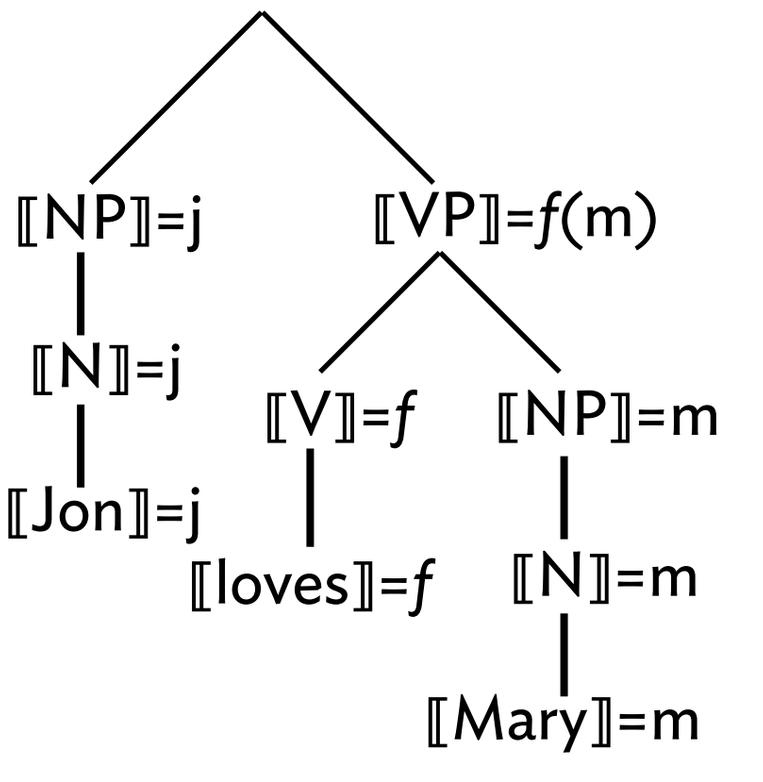

Speaker said (or made as if to say) p

Speaker meant q

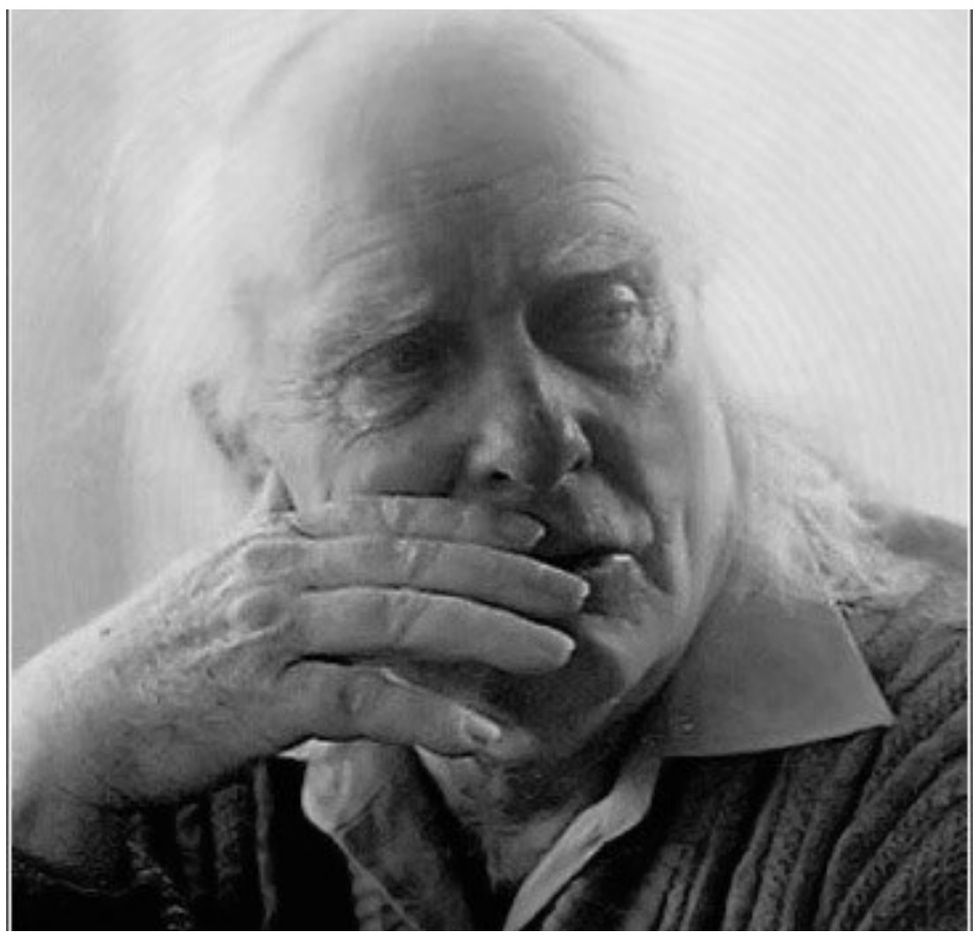

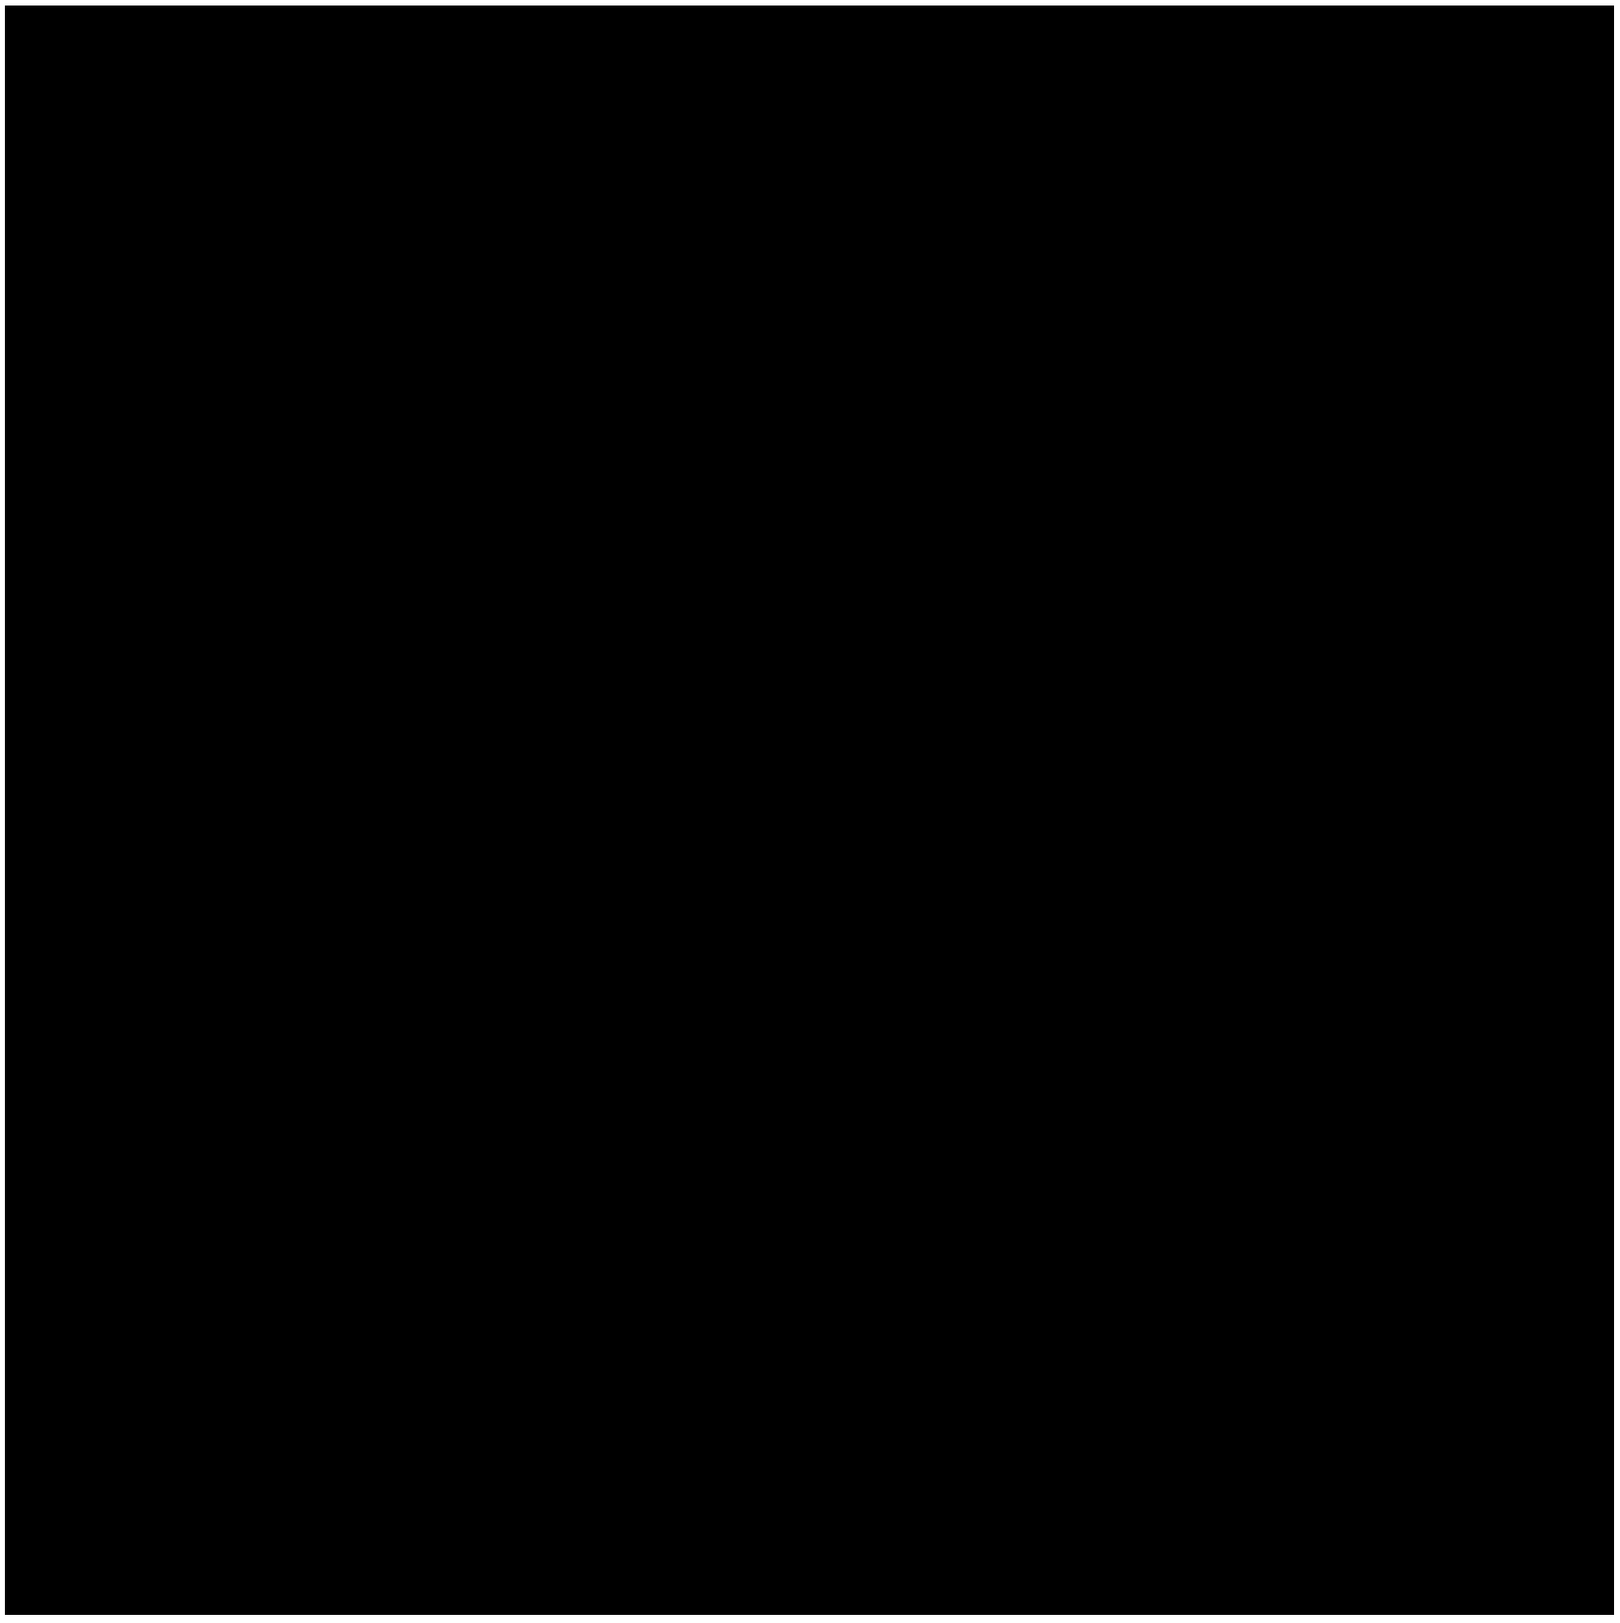

in Grice & Co.

Two Points of Grice Exegesis

1. Grice nowhere insists that pragmatics kicks in only after we work out what is said.

2. Calculability does not mean algorithmic derivability.

Calculability

> ...the presence of a conversational implicature must be capable of being worked out; for even it if can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature; it will be a conventional implicature.

—Grice, 'Logic & Conversation', 50

We interpret these passages [about calculability] rather more stringently than other commentators sometimes do, but that's because we take so seriously the bridge between semantics and pragmatics that they seem to establish.

—Lepore & Stone, p.22

Calculability

A general pattern for the working out of a conversational implicature might be given as follows: "He has said that *p*, there is no reason to suppose that he is not observing the maxims, or at least the cooperative principle; he could not be doing this unless he thought that *q*; he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that *q* is required; he has done nothing to stop me thinking that *q*; and so he has implicated that *q*.

—Grice, 'Logic & Conversation', 50

Calculability

A general pattern for the working out of a conversational implicature might be given as follows: "He has said that *p*, there is no reason to suppose that he is not observing the maxims, or at least the cooperative principle; he could not be doing this unless he thought that *q*; he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that *q* is required; he has done nothing to stop me thinking that *q*; and so he has implicated that *q*.

—Grice, 'Logic & Conversation', 50)

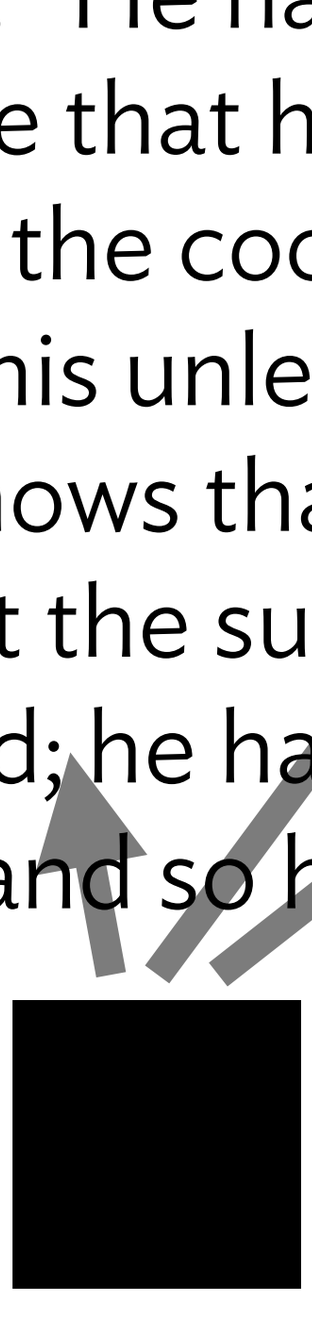

David Marr's classic statement distinguished three levels of description of a computational psychological process (Marr 1982, p. 24 ff.). The first level states which function is computed and why; the second states which algorithm computes the function; and the third states how the algorithm is realized in hardware. The level with which I will be concerned does not lie within one of Marr's levels. It lies between his first and second levels; for this reason I will dub it 'level 1.5'.Level 1.5 **states the information on which the algorithm draws**.

—Peacocke, 'Explanation in Computational Psychology', Mind & Language, 1986, p.101

In [Sperber & Wilson's] hands, Grice's rational reconstruction of the type of argument an audience should be capable of constructing to verify the presence and content of a conversational implicature is at best something that points in the direction of a theory of non-demonstrative inferences that need to be characterized in terms of brute facts about human cognition.

—Neale, 'Implicit Reference', 49

[redacted]

in Communication

(1) Data

(A) ⋮

S: α

⋮

(B) ⋮

S: β

⋮

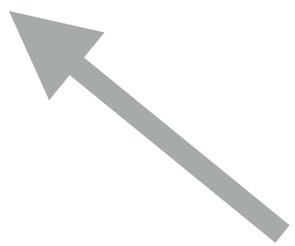

Dialogues (A) and (B) differ only in that *S* utters α in (A) and β in (B).

(2) Different Readings

(A) ⋮

S: α ← In uttering α, S performs an indirect speech act σ.

⋮

BUT!

(B) ⋮

S: β ← In uttering β, S **does not** perform σ.

⋮

(3) Theoretical Similarity

(A) ⋮

S: α ←

⋮

The presence of σ in (A) is usually explained by Gricean mechanisms that turn on α's syntactic/semantic properties XYZ.

BUT!

(B) ⋮

S: β ←

⋮

β also has properties XYZ, and these do not give rise to a σ-reading.

(4) Anti-Grice Conclusion

(A) ⋮

S: α

⋮

The Gricean explanation of (A)'s σ-reading over-generates.

(B) ⋮

S: β

⋮

We should give some other kind of explanation.

(5) Alternative Explanation

(A) ⋮

S: α

⋮

(B) ⋮

S: β

⋮

Here's an alternative, grammatical explanation of (A)'s σ-reading.

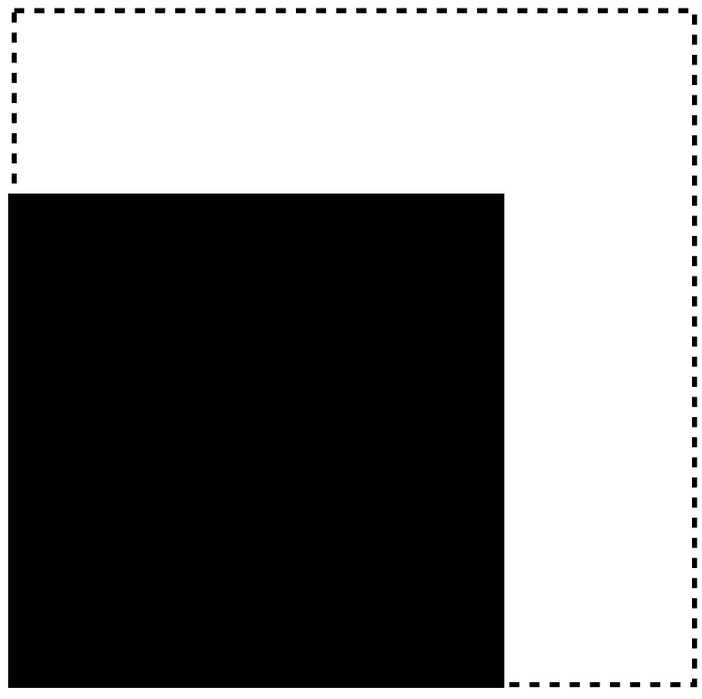

(1) Data

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

(B) Oil prices **have doubled** and demand for consumer goods **have plunged**.

(2) Different Readings

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

↘ doubling followed/caused plunge

(B) Oil prices **have doubled** and demand for consumer goods **have plunged**.

↘ doubling followed/caused plunge

(3) Theoretical Similarity

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

↘ doubling followed/caused plunge

Normally understood as a manner implicature that is signaled by the order of the two clauses.

(3) Theoretical Similarity

(B) Oil prices **have doubled** and demand for consumer goods **have plunged**.

↘↘ doubling followed/caused plunge

But the events are described in the same order in (B), and this doesn't signal a parallel reading.

(4) Anti-Grice Conclusion

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

↘ doubling followed/caused plunge

So, this shouldn't be understood as an implicature after all, lest our explanation overgenerate.

(5) Alternative Explanation

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

↘ doubling followed/caused plunge

"...a simple past tense verb like 'doubled' is interpreted temporally much like a pronoun. It is understood to describe what happened during a contextually specified interval, in much the same way as a pronoun is understood to describe a contextually specified individual."

Where is

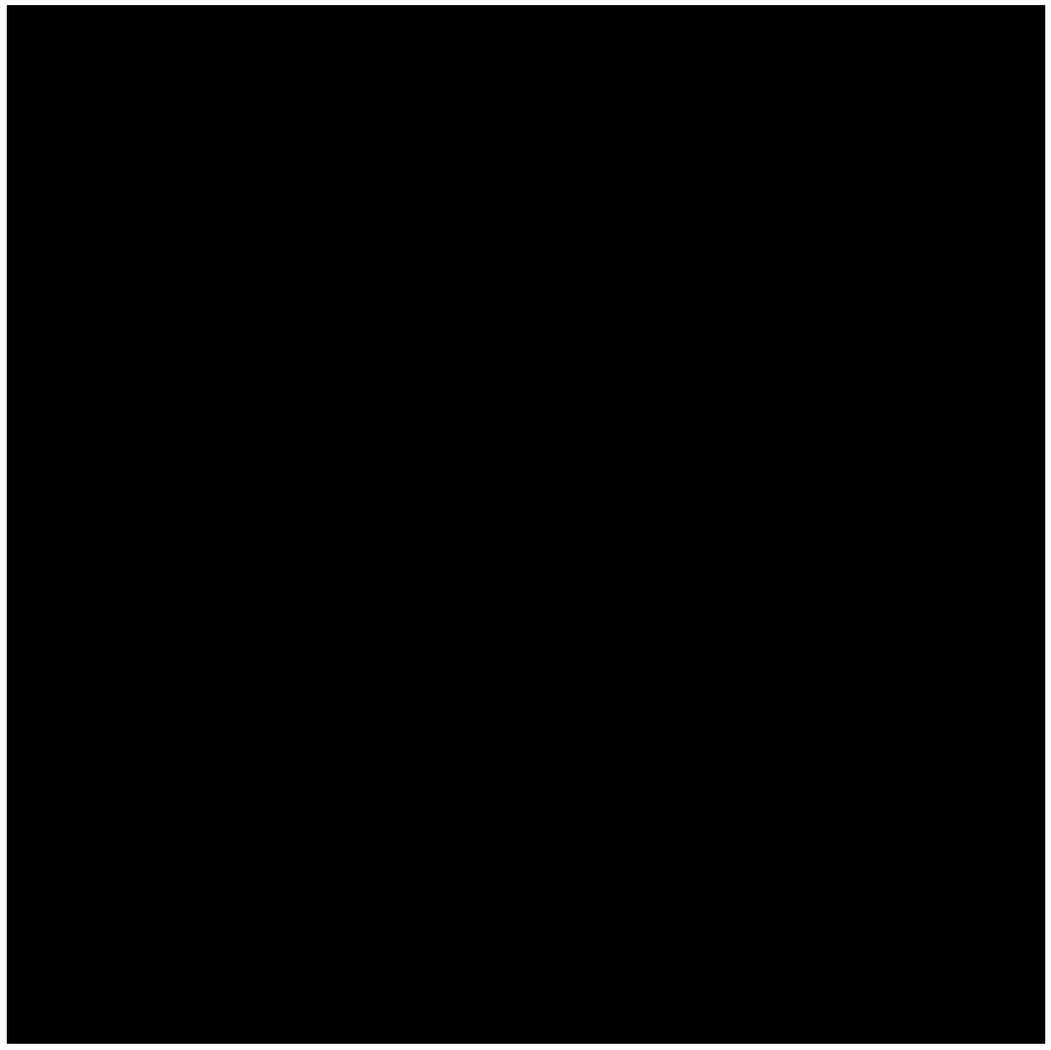

hiding here?

(5) Alternative Explanation

(A) Oil prices **doubled** and demand for consumer goods **plunged**.

↘ doubling followed/caused plunge

"...a simple past tense verb like 'doubled' is interpreted temporally much like a pronoun. It is understood to describe what happened during a **contextually specified** interval, in much the same way as a pronoun is understood to describe a contextually specified individual."

Ernie is eating lunch with his friend Tony, who has a casual interest in economics. Ernie takes out his phone and glances at his stock app, looks distressed, and blurts out: "Darn! All of my stock in Samsung, Nike, and Proctor & Gamble has been tanking!" Tony considers the question and responds: "[It's because of Iraq.]

(B) **Oil prices have doubled and demand for consumer goods have plunged.**

↘ **doubling followed/caused plunge**

(A) Can/could you pass the salt?

(B) Are you able to pass the salt?

Usually an indirect request.

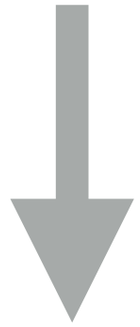

(A) Can/could you pass the salt?

(B) Are you able to pass the salt?

Not so much.

Even in these cases, however, we suspect that hearers must make an intuitive guess about a conventionalized indirect speech act, and so the listener's inference lacks the content of a Gricean calculation. On the ambiguity view, the question to ask is (106).

(106) What is a plausible convention that I could postulate to assign this utterance a likely intended interpretation?

—Lepore & Stone, p.105

[redacted]

in Convention Acquisition

(1) How do conventions arise?

(2) How do we become participants in pre-existing conventions?

Reasons why these questions are pressing:

1. Recent work by Chris Barker, Peter Ludlow, Josh Armstrong, etc., suggests that we're constantly renegotiating conventions and creating new ones, seamlessly and while we speak.

2. Ernie & Matt argue that we should think of the interpretation of particularized implicatures as the attempt to formulate a novel convention that applies to the case.

(i) We start by Xing in some unconventional, but fairly arbitrary way.

(ii) The fact that we've Xed in this way creates a precedent, which *somehow* causes us to keep Xing in that way.

(iii) This way of Xing becomes standard, and *some* mechanism keeps the standard from changing.

(iv) We now have a conventional way of Xing.


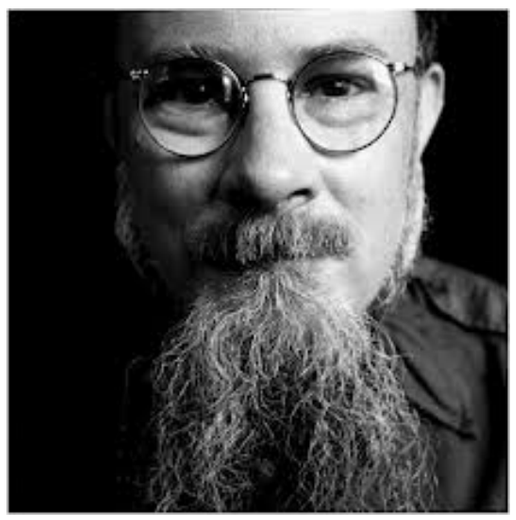


V. Hacquard, 'Bootstrapping Attitudes' (2014)

- **Develops a precise and empirically confirmed model of how children learn the semantics of attitude verbs**

- **Two inputs to the algorithm: syntax and speaker meanings.**

"...children can, and in fact, do hypothesize speaker meanings; sometimes to a fault" (20).

Conclusions

Even if Ernie & Matt's specific arguments in §2 are all sound, their big-picture philosophical conclusions have not been established.

Conclusions

Semantic, language-specific mechanisms are the oil in the interpretive engine, not the engine itself.

Conclusions

Even if we can't all learn to stop worrying and love the Black Box, we should at least all learn to respect it.