

A Satisfying Account of Conjunction Fallacies*

Ash Asudeh
Department of Linguistics
University of Rochester

December 4, 2020
Workshop on Natural Language Meaning

1 Overview

-

1.1 Outline

- Section 2 lays out the more general project developed in Asudeh and Giorgolo (2020).
- Section 3 lays out the background on conjunction fallacies and issues of compositionality raised by previous approaches.
- Section 4 lays out the intuitions behind our model.
- Section 5 discusses how monads can model uncertainty as probability.
- Section 6 discusses how a previous monadic solution that we have proposed (Giorgolo and Asudeh 2014) needs to appeal to general Gricean reasoning to fully explain conjunction fallacies.
- Section 7 introduces tropical semirings, explains how this predicts the pattern of experimental results without appeal to Gricean pragmatics, and shows that the result captures the duality of conjunction and disjunction.
- Section 8 concludes.

2 Our project

- The Enriched Meanings project (Asudeh and Giorgolo 2020) is about using a concept from category theory and functional programming — monads — to model some murkier aspects of natural language meaning.¹
- Monads have been successfully used in the semantics of programming languages to characterize certain classes of computation (Moggi 1989, Wadler 1992, 1994, 1995).
- Our original inspiration came from Shan (2001), who sketched how monads could potentially be used to offer new solutions to certain problems in natural language semantics and pragmatics.
- Before I answer the question of what we mean by enriched meanings it is useful to build up some common ground about semantic theory.

We adhere to the fundamental assumption of semant1(C)-288(of)-28-46(up)-2hntal

There has often been a tendency in formal semantics to generalize lexical meanings to the worst case.

- For example, based on the necessity of treating some nominals as generalized quantifiers, Montague (1973) famously treated all nominals as generalized quantifiers.
 - Similarly, in order to capture the semantics of intensional verbs, Montague treated simpler, extensional verbs equivalently, but had their lexical entries disregard the intensional parameters. In other words, because some verbs need intensional parameters, all verbs have them.
 - Yet another example is the use of assignment functions to interpret variables: Every interpretation is relativized to an assignment function, whether what is being interpreted contains variables or not. Again, because assignment functions are needed in some cases, they are there in all cases.
- If we want linguistic semantics to be a linguistically and philosophically sound science, we should avoid generalizing to the worst case.

Serious steps were taken in this direction through the development of type shifting rules (Partee and Rooth 1983)

- Thus, enriched meanings provide a principled way to capture a mixture of simpler and more complex meanings and also provide empirical benefits in shedding new light on tricky phenomena at the semantics–pragmatics boundary.
- Conjunction fallacies, today's topic, are an instance of such phenomena.

3 Background

- Tversky and Kahneman(1983) noticed that, in a task asking for ratings of the relative likelihoods of different propositions being true, the majority of experimental participants consistently rated the likelihood of the conjunction of two propositions being true as higher than the likelihood that one of the conjoined propositions is true.
- One of their examples is the well-known `Linda paradox' or `Linda effect'.

Experimental participants were given statements like (1).

- (1) Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

As part of the experimental task, they were asked to rank the probability that various statements were true of Linda; the resulting ranking for the relevant cases is given in (2)

Representativeness tends to covary with frequency but not necessarily.

- A crucial point of Tversky and Kahneman's analysis is that this heuristic operates on the conjunction as a whole, or as they put it:

[T]he judged probability (or representativeness) of a conjunction cannot be computed as a function (e.g., product, sum, minimum, weighted average) of the scale values of its constituents. (Tversky and Kahneman 1983: p. 305)

This is essentially a claim that the conjunction is non-compositional, which is rather problematic for any attempt to integrate their observations with linguistic theories of meaning composition.

- Aspects of

4 The intuitions behind our model

- Our model starts from the assumption that people employ multiple strategies when evaluating the likelihood of conjoined uncertain events, an assumption shared with Yates and Carlson (1986).

However, Yates and Carlson(1986) propose a less principled “signed summation model” that is not based on probability theory, unlike ours.

They also assume that unrelated computational processes underpin the different strategies.

- We show that it is possible to assume a single uniform process that computes the likelihood of the conjunction of two events from their two relative likelihoods.

This uniform process may use different but related representations of uncertainty, expressed as alternative algebraic structures, yielding different results.

We can explain the results reported in the literature in terms of an algebraic structure known as a semiring.

This mathematical object is at the heart of monads

- This yields more than one possibility. How is the choice between them made?
- We assume that the observation made in the literature (Hertwig and Gigerenzer 1999) that conjunction fallacies arise only under specific conditions and can be cancelled under other conditions is explained in terms of cognitive/computational economy

In fact, the same computational structure, the monad, can be used together with a number of different underlying semirings, one of them being the probability semiring.

We predict that, in general, if people are presented with a task where there are “no stakes” they will reason based on a representation corresponding to a semiring defined over a relatively simple set with generally simple operations.

Using this strategy will normally lead people to make overconfident estimates, which may result in conjunction fallacies.

If, on the other hand, people are forced to evaluate the consequences of their judgements, such as in the context of a gaming scenario, or if explicitly primed to think in frequentist terms, then they can switch to a more complex representation, with properties that better

In the standard Gricean approach the core meaning is further elaborated by conversational implicatures— additional meanings that are not explicitly stated by the speaker but inferable by the speaker based on conversational goals (Grice's Cooperative Principle and conversational maxims) and real world knowledge.

Crucially, conversational implicatures are determined by context and can be cancelled by explicitly stating that they do not hold.

– For instance, it is a consequence of the standard Gricean view that a numeral like two is

5 Monads and uncertainty

- As a preface to what follows, we need to familiarize ourselves with the related notion of semirings.
- A semiring is:




Some set A ,

With two distinguished elements 0 and 1 , and

Two binary operations $+$ and \cdot ,

Such that the following axioms are satisfied, for all x, y and $z \in A$:

$(x + y) + z = x + (y + z)$
 $(x \cdot y) \cdot z = x \cdot (y \cdot z)$
 $x + (y \cdot z) = (x + y) \cdot z$
 $x \cdot (y + z) = (x \cdot y) + (x \cdot z)$
 $x + 0 = x$
 $x \cdot 1 = x$
 $0 + x = x$
 $1 \cdot x = x$
 $x + 1 = 1 + x$
 $x \cdot 0 = 0 \cdot x = 0$
 $0 + 1 = 1 + 0 = 1$






Figure 2: $t(T)$






Figure 4: Example of a function from $f a; b; c$ to probability distributions over truth values.




Figure 5: Pure joint probability for $f a; b; c$ and $f T; F$.

- But bind actually generates a distribution probability over the final type of its second argument, so the result we get is instead the one shown in Figure 6.
- What bind returns for T in Figure 6 is the sum of a-T, b-T and c-T in Figure 5; similarly for F and a-F, b-F, and c-F. You can check this by eye.
- Thanks to these proofs we can substitute any semiring for the probability semiring and still use our monadic calculus to derive the meaning of complex expressions.
- Giorgolo and Asudeh (2014) offered a semiring that we called the one semiring as a candidate solution for the problem of conjunction fallacies.
- We used the results of Yates and Carlson results, which are summarized in Table 1.




Figure 6: The result of applying bind to its two arguments

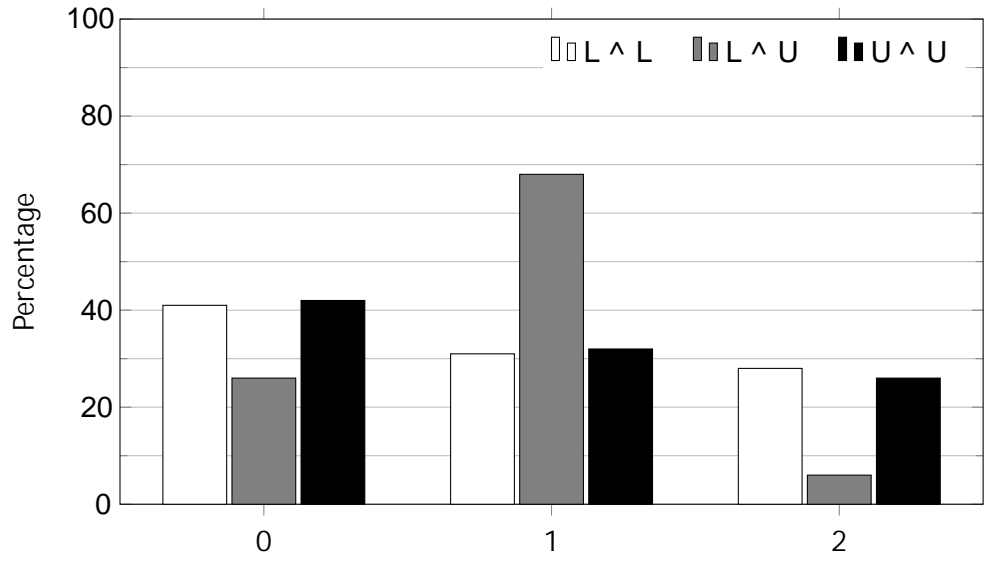
Likelihood of event A	Likelihood of event B	Observed rating	Intuition
U(nlikely)	U	P(A and B)	

+	I	U	P	L	C		I	U	P	L	C
I	I	U	P	L	C	I	I	I	I	I	I
U	U	U	P	L	C	U	I	U	U	U	U
P	P	P	P	L	C	P	I	U	P	P	P
L	L	L	L	L	C	L	I	U	P	L	L
C	C	C	C	C	C	C	I	U	P	L	C

Table 2: The one semiring.

- The one semiring is shown in Table

Thus, the statement that Linda is a bank teller creates an unnecessary implicature that would cause participants to reconsider the context in an aberrant way, perhaps calculating that the experimenters for some bizarre reason believe that bank tellers are generally interested in “issues of discrimination and social justice”, but this likely clashes with the real world knowledge



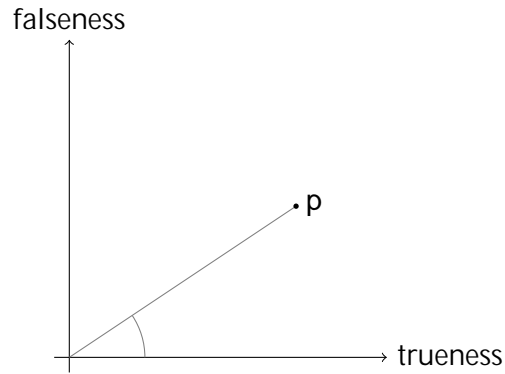


Figure 8: The likelihood of the proposition p represented as a point in the trueness/falseness space.

- The **max-plus** semiring is specified in a similar way.

The semiring is defined on the set $\mathbb{R} \cup \{-\infty\}$, the set of real numbers together with negative infinity. The two operations are defined as follows:

$$(18) \quad x + y = \max(x, y)$$

$$(19) \quad x \cdot y = x + y$$

where \max is a function that returns the greatest of two numbers. The unit of $+$, 0 , is $-\infty$, and the unit of \cdot , 1 , is 0 .

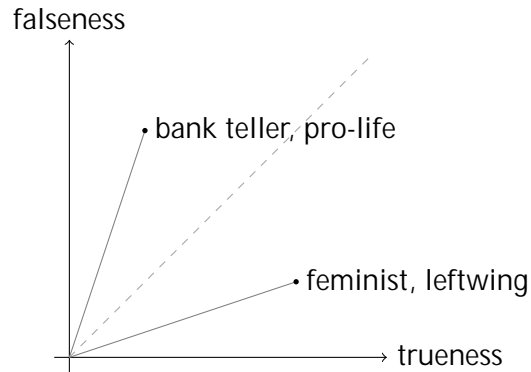


Figure 9: Examples of likely and unlikely propositions in the trueness/falseness space.

- We noted above that, given the equivalence of the **min-plus** and **max-plus** tropical semirings, there is no reason to assume that people use only one and not the other.

We therefore assume that people have both tropical semirings at their disposal and may use either one or even both.

- In other words, we assume that people have cognitive equivalents of both semirings.

In the Linda scenerio, we assume as part of the experimental conditions, that the proposition Linda is a bank teller is unlikely and the proposition Linda is a feminist is likely.

- This is the same assumption that Tversky and Kahneman (1983) make.

We therefore assign each some value in the monadic space — on the false side of 45 in the case of the unlikely proposition and on the true side in the case of the likely proposition — and keep the assumption constant across both tropical semirings.

In the case of the conjunction of the unlikely and likely propositions — which leads to the observed Linda effects/conjunction fallacies in the psychological literature — the two semirings yield identical ordering results, as shown in Figure 10a and Figure 10b.

- For both semirings, the unlikely proposition that Linda is a **bank teller** comes out as more false than the conjunction of this unlikely proposition with the likely proposition that Linda is a **feminist** and this conjunction in turn comes out as more false than the likely proposition on its own.
- Thus, both semirings capture the Linda effect.

- So how does this explain the results reported in Figure 7?

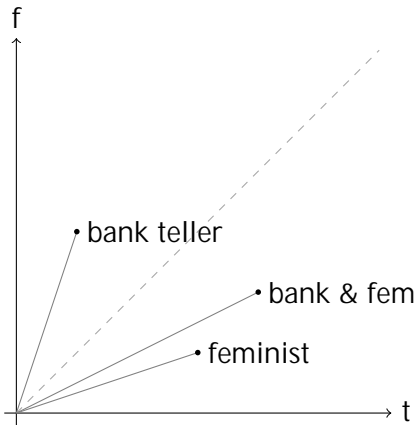
Notice that Figure 7 (and a similar graph in Yates and Carlson 1986) is reporting total number of responses and thus summing across all experimental participants.

The mass of responses clearly peaks at one conjunction error (nearly 70%) in the case of the conjunction of an unlikely and a likely proposition, with a smaller peak at zero conjunction errors (about 25%).


- We assume that the small mass of responses at two conjunction errors is noise.

There are two possible explanations for this, both of which converge on the same result.

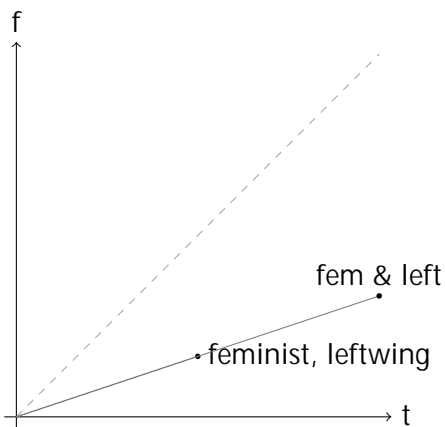
1. If an individual in any particular response exercises the choice to use the **min-plus** or **max-plus** semiring, they will commit a single conjunction error (i.e., a Linda effect).




(a) Conjunction of an unlikely and a likely proposition, given the min-plus tropical semiring.



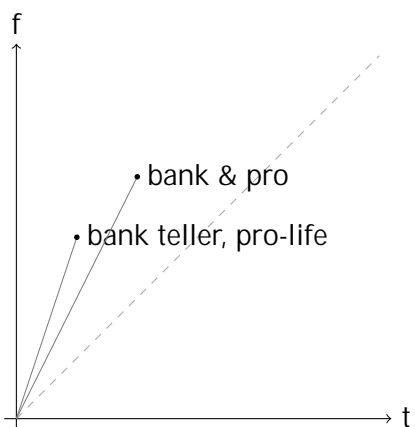
(b) Conjunction of an unlikely and a likely proposition, given the max-plus tropical semiring.




(c) Conjunction of two likely propositions, given the min-plus tropical semiring.



(d) Conjunction of two likely propositions, given the max-plus tropical semiring.



(e) Conjunction of two unlikely propositions, given the min-plus tropical semiring.



(f) Conjunction of two unlikely propositions, given the max-plus tropical semiring.

Figure 10: Conjunction of propositions, given the min-plus (left column) or max-plus (right column) tropical semirings.

1. The first possible explanation is that experimental participants have a preference for the tropical semiring that makes the greatest distinction: This would mean preferring the **max-plus** semiring in the case of a conjunction of two likely propositions and preferring the **min-plus** semiring in the case of a conjunction of two unlikely propositions.

This would still correctly predict the (near) tie between Likely/Likely and Unlikely/Unlikely at zero conjunction errors in Figure 7.

And it also allows for the almost uniform distribution between Likely/Likely and Unlikely/Unlikely in the other cases (one and two conjunction errors in Figure 7).

2. The second possible explanation is that experimental participants try both semirings, but in the face of the inconsistent results that this delivers (i.e., the **min-plus** and **max-plus** semirings give inconsistent answers for the Likely/Likely case and also for the Unlikely/Unlikely case), at least some participants back off to proper probabilistic reasoning and use the probability semiring, such that more responses occur in the zero conjunction error part of the response mass for both Likely/Likely and Unlikely/Unlikely.

Both of these explanations would thus predict the distribution of data, but using different mechanisms. Therefore there are distinct empirical predictions at play here which could somehow be tested experimentally.

- The tropical model also behaves well logically, as illustrated in Figure 11.

The key observation is that — despite the fact that conjunction does not behave absolutely logically (likewise disjunction), which is after all necessary to model the Linda effect — conjunction and disjunction still form duals.

We use the likely proposition that Linda is a feminist and the unlikely proposition that Linda is a bank teller to demonstrate this.

The dashed line in Figure 11 once again represents total uncertainty as to the truth/falsity of a proposition.

The **feminist** proposition is assigned a vector in the true space, so below the dashed line, whereas the **bank teller** proposition is assigned the symmetrical vector in the false space, so above the dashed line.

The **min-plus** semiring treats the conjunction of **bank teller** and **feminist** as more true than false; this vector is labelled “ $\text{bank} \wedge_{\min} \text{fem}$ ” in Figure 11.

– We also observe that the **min-plus** semiring dashed line,

This is, in our view, a serious defect which we have sought to remedy by offering a compositional treatment instead.

In particular, we have developed a model in which the behaviour of natural language and is closely connected to the basic logical meaning of conjunction as a propositional connective, which upholds a key aspect of Grice's program (Grice 1989) and takes seriously his Modified Occam's Razor: 'Senses are not to be multiplied beyond necessity' (Grice 1978: 118).

References

- Ash Asudeh and Gianluca Giorgolo. *Enriched Meanings: Natural Language Semantics with Category Theory*. Oxford University Press, Oxford, 2020.
- Chris Barker. Continuations and the nature of quantification. *Natural Language Semantics* 10(3): 211–242, 2002.
- Chris Barker. Quantificational binding does not require c-command. *Linguistic Inquiry*, 43(4):614–633, 2012.
- Chris Barker. Negative polarity as scope marking. *Linguistics and Philosophy* 41(5):483–510, 2018.
- Chris Barker. NL as the logic of scope and movement. *Journal of Logic, Language and Information* 28(2):217–237, 2019.
- Chris Barker and Chung-chieh Shan. Types as graphs: Continuations in Type Logical Grammar. *Journal of Logic, Language and Information* 15(4):331–370, 2006.
- Chris Barker and Chung-chieh Shan. Donkey anaphora is in-scope binding. *Semantics and Pragmatics*, 1(1):1–46, June 2008.
- Chris Barker and Chung-chieh Shan. *Continuations and Natural Language* Oxford University Press, Oxford, 2014.
- Chris Barker, Raffaella Bernardi, and Chung-chieh Shan. Principles of interdimensional meaning interaction. In Man Li and David Lutz, editors, *Proceedings of SALT 20* pages 109–127, 2010.
- Johan van Benthem. *Language in Action: Categories, Lambdas, and Dynamic Logic*. North-Holland, Amsterdam, 1991. Reprinted in 1995 by the MIT Press, Cambridge, MA.
- Dylan Bumford. Incremental quantification and the dynamics of pair-list phenomena. *Semantics and Pragmatics* 8(9):1–70, 2015.
- Dylan Bumford. Split-scope definites: Relative superlatives and Haddock descriptions. *Linguistics and Philosophy* 40(6):549–593, 2017a.
- Dylan Bumford. *Split-Scope Effects in Definite Descriptions* PhD thesis, New York University, 2017b.
- Dylan Bumford and Chris Barker. Association with distributivity and the problem of multiple antecedents for singular Different. *Linguistics and Philosophy* 36(5):355–369, 2013.
- Simon Charlow. *On the Semantics of Exceptional Scope* PhD thesis, New York University, 2014.
- Simon Charlow. Post-suppositions and semantic theory. Ms., Rutgers University. To appear in *Journal of Semantics* 2018.

- Simon Charlow. The scope of alternatives: Indefiniteness and islands. *Linguistics and Philosophy* pages 1–46, 2019a. First online: August 26, 2019. Early access version and pagination.
- Simon Charlow. Where is the destructive update problem? *Semantics and Pragmatics* 12(10):1–17, 2019b. Early access version and pagination.
- Reuben Cohn-Gordon. Monad transformers for natural language: Combining monads to model effect interaction. Ms., Stanford University, nd.
- Jan van Eijck and Christina Unger. *Computational Semantics with Functional Programming* Cambridge University Press, Cambridge, 2010.
- Gianluca Giorgolo and Ash Asudeh. One semiring to rule them all. In *Proceedings of the 36th Annual Cognitive Science Society Conference*, pages 116–121, 2014.
- Gianluca Giorgolo and Christina Unger. Coreference without discourse referents: A non-representational DRT-like discourse semantics. In B. Plank, T. Kim Sang, and T. Van de Cruys, editors, *Computational Linguistics in the Netherlands 2009*, number 14 in LOT Occasional Series, pages 69–81. LOT, 2009.
- Michele Giry. A Categorical Approach to Probability Theory. In *Categorical Aspects of Topology and Analysis, Proc. int. Conf., Ottawa 1981, Lect. Notes Math. 915, 68-85 (1982)*1982.
- H. Paul Grice. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics: Speech Acts* volume 3, pages 41–58. Academic Press, San Diego, CA, 1975.
- H. Paul Grice. Further notes on logic and conversation. In Peter Cole, editor, *Syntax and Semantics: Pragmatics* volume 9, pages 113–127. Academic Press, San Diego, CA, 1978.
- H. Paul Grice. *Studies in the Way of Words* Harvard University Press, Cambridge, MA, 1989.
- James A. Hampton. Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14(1):12–32, 1988.
- Ralph Hertwig and Gerd Gigerenzer. The “conjunction fallacy” revisited: How intelligible is the conjunction fallacy? *Journal of Experimental Psychology: Applied*, 10(1):1–16, 2004.

Yong Lu. The conjunction and disjunction fallacies: Explanations of the Linda problem by the equate-to-differentiate model. *Integrative Psychological and Behavioral Science*, 50(3):507–531, 2016.

Jirka Maršík. *Effects and Handlers in Natural Language*. PhD thesis, Université de Lorraine, 2016.

Barbara Mellers, Ralph Hertwig, and Daniel Kahneman. Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, 12:269–275, 2001.

Eugenio Moggi. Computational lambda-calculus and monads. In *LICS*, pages 14–23. IEEE Computer Society Press, 1989.

Mehryar Mohri. Weighted automata algorithms. In Manfred Droste, Werner Kuich, and Heiko Vogler, editors, *Handbook of W Wpages3r596(14–]TJ 13216(154.226]TJSpringer-)55er)-36erlagges3r59Berlin*

Herbert A. Simon. Rational choice and the structure of the environment. *Psychological Review* 63 (2):129–138, 1956.

Imre Simon. Recognizable sets with multiplicities in the tropical semiring. In Michal P. Chytil,