

Disjunction and possibility*

Matthew Mandelkern
New York University, mandelkern@nyu.edu

June 18, 2024

Abstract

I argue that $\lceil p \text{ or } q \rceil$ can be interpreted as $(p \vee q) \wedge \diamond p \wedge \diamond q$, where \diamond is a possibility modal whose flavor can be epistemic, circumstantial, or deontic. I show that no extant theory can account for this generalization, and argue that the best way to do so is with a direct theory on which ‘or’ means $\lambda p.\lambda q.(p \vee q) \wedge \diamond p \wedge \diamond q$. I show that the resulting theory also yields an appealing account of both wide- and narrow-scope free choice inferences.

1 Introduction

I argue for a new theory of the meaning of natural language disjunction, on which ‘or’ means $\lambda p.\lambda q.(p \vee q) \wedge \diamond p \wedge \diamond q$, where \vee is classical disjunction, \wedge is classical conjunction, and \diamond is a possibility modal whose flavor is determined by context.

I motivate this theory by introducing a variety of patterns where $\lceil p \text{ or } q \rceil$ seems to be interpreted as $(p \vee q) \wedge \diamond p \wedge \diamond q$, where \diamond is either an epistemic possibility modal \diamond_e (roughly, ‘might’) or a circumstantial possibility modal \diamond_c (roughly, ‘could have’) or a deontic possibility modal \diamond_d (roughly, ‘may’). I show that no extant theories can account for these patterns. I argue that the best way to do so is the simplest one: maintaining that this is the meaning of ‘or’ in general. I compare my proposal to important precedents in [Zimmermann 2000](#); [Geurts 2005](#), arguing that my approach improves on those in a number of ways. Finally, I show how my theory yields an extremely simple account of both wide- and narrow-scope free choice inferences.

*Many thanks to Nathan Klindinst for extensive and insightful discussion, and for the tremendously helpful observations I note in the paper. Thanks also to Sam Carter, Cian Dorr, Patrick Elliott, Danny Fox, Melissa Fusco, Simon Goldstein, John Hawthorne, Ben Holguín, Salvador Mascarenhas, Jacopo Romoli, Richard Roth, Benjamin Spector, Stephen Yablo, Snow Zhang, and audiences at Heinrich-Heine Universität, the Institut Jean Nicod, PhLiP, UT Austin, the University of Edinburgh, UCL, and ZAS-Berlin for very helpful discussion and feedback.

2 Embedded epistemic possibility

I begin by describing a number of cases where an embedded disjunction with the form $\lceil p \text{ or } q \rceil$ is interpreted as $(p \vee q) \wedge \Diamond p \wedge \Diamond q$ (Throughout, I will use the symbols $\vee, \wedge,$ and \neg for classical disjunction, conjunction, and negation, and ‘or’ for natural language disjunction.)

A standard observation about disjunction is that it can give rise to epistemic possibility inferences. Suppose I assert (1):

- (1) John went to Paris or to Berlin.

It’s very natural to infer from my assertion of (1) that both disjuncts are epistemically possible for me: that is, for all I know, John went to Paris, and for all I know, John went to Berlin. This is easy enough to explain from the point of view of a classical theory of disjunction, together with Gricean reasoning (see Grice 1989; Stalnaker 1975). If I think that John went to Paris and am in a position to assert (1), then I should have just asserted ‘John went to Paris’, which is strictly stronger than (1) on a classical theory. Likewise if I think that John went to Berlin. Hence I neither think that John went to Paris nor that he went to Berlin. But, since I asserted (1), I must think that he went to Paris or Berlin. Hence I leave open that he went to Paris and that he went to Berlin.

While it is a commonplace that an *assertion* of a disjunction can give rise to epistemic possibility inferences, a novel observation, due to Nathan Klinedinst (p.c.), is that these epistemic possibility inferences can also be *embedded*. To see this, consider the following:

- (2) *A novel virus has been detected in Paris. Cian, Kit, and Hartry’s students have just returned from a trip to Europe, on which every student visited exactly one of Berlin or Paris (let us assume further that their classes are disjoint and non-empty). Unfortunately, some of the students’ records are mixed up. We know that Cian’s students all went to Berlin. But we don’t know which of Kit and Hartry’s students went where. A public health officer, apprised of these facts, announces:*
- a. Every student who went to either Berlin or Paris must be quarantined. But those who definitely went to Berlin can be released.

My claim is that the sequence in (2-a) has a coherent reading.

If ‘or’ simply means \vee , however, then the sequence would not be coherent. Any student S who definitely went to Berlin (i.e., any student in Cian’s class) also went to Berlin \vee Paris, so the two sentences in (2-a) would put contradictory constraints on S, namely, that she be quarantined and released.¹

¹I use NP disjunction in many examples, but am assuming that it reduces to clausal disjunction; in the examples I give, replacing the NP disjunction with clausal disjunction does not seem to make any difference

The coherent reading of (2-a) intuitively says that any student who definitely went to Berlin can be released, while any student who went to Paris \vee Berlin but might have gone to Paris must be quarantined; in other words:

- (3) Every student who went to either Berlin or Paris and might have gone to Berlin and might have gone to Paris must be quarantined. But those who definitely went to Berlin can be released.

There is of course no risk of incoherence from (3), since the restrictors are disjoint: no student who definitely went to Berlin is such that she also might have gone to Paris. Hence the coherent reading of (2-a) is plausibly one where $\lceil P \text{ or } B \rceil$ means $(P \vee B) \wedge \Diamond P \wedge \Diamond B$. In other words, the epistemic possibility inferences that an *assertion* like (1) gives rise to can apparently also be *embedded* in a construction like (2-a).

(2-a) also appears to have an incoherent reading. The coherent reading can naturally be brought out by emphasizing ‘or’ and by setting up a contrast with a subsequent clause (in this case, ‘definitely went to Berlin’); an incoherent reading can be more naturally brought out without these aids. (This is reminiscent of standard observations about embedding implicatures in non-upward-entailing environments, a point I’ll return to.) That both a coherent and classical reading exist is an important observation which any adequate theory will have to account for.

I should also note that not everyone seems able to access the non-classical readings I am targeting in this paper. I have presented these data to many audiences, and most people seem able to access most of the target readings, but not everyone can access all of them. This is an area where experimental testing would be helpful. Based on my judgments and those of many consultants, I am confident that these readings exist; an interesting further question, which I leave open, is why there is interpersonal variability in accessing them.

The reading of ‘or’ as $(p \vee q) \wedge \Diamond_e p \wedge \Diamond_e q$ can be brought out in other environments besides the restrictors of quantifiers. For instance, consider (4), where Philippa is a student who went to Europe:

- (4) If Philippa went to either Berlin or Paris, we must quarantine her. But if we’re sure she went to Berlin, we can let her go.

(4) seems like a corollary of (2-a). Intuitively, the antecedent of the first conditional is interpreted as ‘Philippa went to either Berlin or Paris and both are possible’.²

Suppose next that Cian, Kit, and Hartry each had exactly five students in their class.

to relevant judgments.

²Though the dialectic around a sequence like (4) is subtle, since on a variably strict theory of conditionals, no incoherence need arise from a classical interpretation of ‘or’; so the coherence of (4) does not form a strong argument for the existence of embedded epistemic possibility inferences.

Then, given the set-up, it seems that (5) has a prominent true reading:

- (5) Exactly ten of our students went to either Berlin or Paris, while another five definitely went to Berlin.

But note that every student went to Berlin \vee Paris, so it is false that exactly ten students went to Berlin \vee Paris; rather, exactly fifteen students went to Berlin \vee Paris. So the true reading of (5) cannot be one on which ‘or’ means \vee . Instead, it is intuitively a reading where ‘went to either Berlin or Paris’ means ‘went to Berlin \vee Paris and might have gone to either’, since there are exactly ten students like that (namely, the students in Kit and Hartry’s classes).

More cases like this are easy to generate, on a model with the present one; for instance, (6-a) appears to have a true reading in the case described here:

- (6) *Every first-year philosophy student must take exactly one of logic, semantics, or probability theory. Students may pre-register for more classes than they will ultimately take. Presently three students have registered for logic and for semantics, while another five have registered only for logic, and the last two just for probability.*
- a. Exactly three students will take either logic or semantics, another five will take logic, and the last two will take probability.

Since eight students will take logic \vee semantics, the true reading of (6-a) must be one where ‘or’ is interpreted non-classically; intuitively, as ‘three students are such that they will take logic or semantics and both are possible’. I leave it to readers to explore similar readings in the restrictors of quantifiers, antecedents of conditionals, and so on.

2.1 Deriving embedded epistemic possibility

We have now seen a variety of cases where $\lceil p \text{ or } q \rceil$ is interpreted $(p \vee q) \wedge \diamond_e p \wedge \diamond_e q$. It is, again, well-known that an *assertion* of $\lceil p \text{ or } q \rceil$ can give rise to the conjunction of epistemic possibility claims $\diamond_e p \wedge \diamond_e q$, and this is straightforward to derive on a broadly Gricean basis. But the standard Gricean reasoning does nothing to explain how epistemic possibility inferences could become part of the *embeddable content* of a disjunction.

There is, however, a well-developed tool for turning Gricean implicatures into part of the embeddable content of a sentence: namely, by positing a covert operator *EXH* at LF which means something like ‘only’ (Chierchia et al., 2012). So a natural thought at this point is to parse the embedded disjunction as $EXH(p \vee q)$ to try to derive our target readings. The problem is that there is no version of *EXH* that, when applied to $\lceil p \text{ or } q \rceil$, yields embedded epistemic possibility inferences.

We *can* use *EXH* to derive this reading if we throw into the mix a second covert operator,

usually written K , which means, in essence, ‘the speaker believes...’ (Meyer, 2013). If we can parse ‘went to Berlin or Paris’ as $EXH(K(p \vee q))$, then (given some further simple assumptions) we’ll generate $\neg K \neg p$ and $\neg K \neg q$ as part of the meaning of the sentence, which, given reasonable assumptions about the relation between K and \diamond_e , gets us what we want. The problem is that, while there have been arguments for the existence of K , it is implausible that K can be *embedded*: if K exists, then it must take highest scope over the overt material in an asserted sentence. The problem with embedding K is obvious: ‘John isn’t here’ doesn’t have a reading on which it *means* ‘I don’t know that John is here’; if it did, then you could assert ‘John isn’t here’, find out that he was there, and maintain that what you said was strictly speaking true, since you only said that you didn’t *know* that he was there. But if K could be embedded, then such a reading could be obtained by putting K under negation, i.e., as $\neg K J$. So, while having *matrix* K operators is theoretically viable, having embeddable K operators is a non-starter; and that’s what we would need to get *embedded* epistemic possibility inferences.

A different, more direct strategy for accounting for epistemic possibility inferences is to argue that in these cases we simply parse the embedded disjunction in the scope of an unpronounced possibility modal, so that the restrictor of (2-a), for instance, has the logical form $\diamond_e(p \vee q)$. To finish the story, we would need an account of embeddable free choice inferences on which this has a reading that entails $\diamond_e p \wedge \diamond_e q$ —for instance, by throwing in an EXH with a semantics that can derive free choice along the lines of (Bar-Lev and Fox, 2020). But this still won’t capture the intuitive reading of our target sentence, which entails not just $\diamond_e B \wedge \diamond_e P$ but also the classical disjunction $B \vee P$. To see this, compare (5) (repeated here) with (7):

- (5) Exactly ten students went to either Berlin or Paris, while another five definitely went to Berlin.
- (7) Exactly ten students might have gone to Berlin or Paris, while another five definitely went to Berlin.

The present account predicts that (5) and (7) are equivalent on the target readings, but this is wrong. To see this, elaborate the case by adding a sixteenth student who in fact went to Madrid, but might have gone to Berlin or Paris. So, to summarize: ten students went to either Berlin or Paris and might have gone to Berlin and might have gone to Paris; another five definitely went to Berlin; while a sixteenth went to Madrid but might have gone to Berlin and might have gone to Paris. In this case, (5) still has a true reading (the sixteenth student is neither such that she went to Berlin or Paris, nor such that she definitely went to Berlin); while (7) seems clearly false in this case, since eleven, not ten, students might have gone to Berlin or Paris.

This also brings out problems with a different analysis, due to Zimmermann 2000. On

that view, $\lceil p \text{ or } q \rceil$ simply means $\diamond_e p \wedge \diamond_e q$. That view wrongly predicts (5) to be equivalent to (8), which, like (7), only has a false reading in the emended scenario, while (5)’s target reading remains true:

- (8) Exactly ten students might have gone to Berlin and might have gone to Paris, while another five definitely went to Berlin.

Instead, the target reading is obtained by interpreting $\lceil p \text{ or } q \rceil$ as $(p \vee q) \wedge \diamond_e p \wedge \diamond_e q$. On my preferred theory, this is simply one contextual resolution of the meaning of ‘or’: that is, I maintain that $\lceil p \text{ or } q \rceil$ means $(p \vee q) \wedge \diamond p \wedge \diamond q$, with the flavor of \diamond resolved by context.

Let me return now to the observation above that ‘or’ also has a classical reading, one which renders the sequences above incoherent. Here is how my view can capture this observation. There is an interpretation of \diamond on which it is essentially trivial: the one we obtain by letting every world access *every possibility whatsoever*. Then, as long as p is possible *in some sense*, no matter how distant or etiolated, $\diamond p$ will be true. Let us write \diamond_{\top} for this resolution of the context sensitivity of \diamond . There is clearly more to say about this broadest possible notion of possibility, but I think it should be noted that even a contradiction can be possible in some linguistically important sense (say, epistemically possible for an agent who leaves it open that there are true contradictions). Indeed, we can spell out \diamond_{\top} linguistically, so that p is possible whenever p is a sentence of natural language. Then $\diamond_{\top} p$ is just equivalent to \top . If we resolve the context sensitivity of \diamond this way, then $\lceil p \text{ or } q \rceil$ will mean $(p \vee q) \wedge \diamond_{\top} p \wedge \diamond_{\top} q$, which is equivalent to $(p \vee q) \wedge \top \wedge \top$ and hence to the classical disjunction $p \vee q$. Thus a classical interpretation of disjunction falls out of my approach *as one resolution of the context sensitivity of disjunction*, along with an epistemic (and, as we will see, a circumstantial and deontic) interpretation.

3 Embedded circumstantial possibility

A natural question to entertain at this point is whether our target reading can be derived by a non-classical theory of disjunction on which each disjunct is raised as a possibility, as in the related approaches of alternative semantics (Alonso-Ovalle, 2006), inquisitive semantics (Ciardelli et al., 2018; Mascarenhas and Picat, 2019), state-based semantics (Hawke and Steinert-Threlkeld, 2020), or team semantics (Aloni, 2022). All these frameworks closely associate ‘or’ with epistemic possibility in ways that might be leveraged to account for our readings. I am not sure how the details of this would work. But instead of speculating about this, I will present new data which I think rule out a general response along any of these lines. These data show that $\lceil p \text{ or } q \rceil$ can be interpreted as $(p \vee q) \wedge \diamond_c p \wedge \diamond_c q$. I know of no account that can derive such a reading. In particular, the accounts just mentioned naturally

connect disjunction to *epistemic* possibility, but not to circumstantial possibility, and so I do not think any of them can be leveraged to account for the patterns I will bring out in this section.

Suppose ten guests are coming to Thanksgiving dinner. I have asked each guest to bring one pie. I know that seven guests are going to a store that carries apple and blackberry pies, and another three guests are going to a store that carries blackberry and cherry pies. I don't know which guest is going to which store. Given this set-up, I know (9), on at least one prominent reading:

- (9) Exactly seven guests will bring apple or blackberry, and another three guests will bring blackberry or cherry.

The reading of (9) which I know to be true, however, is not one on which 'or' means \vee , since for all I know, some of the three guests going to the blackberry/cherry store will get blackberry, in which case more than seven guests will bring apple \vee blackberry. Nor, crucially, is it a reading on which $\lceil p \text{ or } q \rceil$ is interpreted as $(p \vee q) \wedge \diamond_e p \wedge \diamond_e q$, because I don't know which guests are going where. So all *ten* guests are such that it's epistemically possible that they'll bring apple, epistemically possible they'll bring blackberry, and epistemically possible they'll bring cherry. So, I can't rule out that more than seven guests are such that they bring apple \vee blackberry and might bring apple and might bring blackberry.³

Instead, the target reading seems to be one where we interpret $\lceil A \text{ or } B \rceil$ as $(A \vee B) \wedge \diamond_c A \wedge \diamond_c B$, where \diamond_c is *circumstantial* possibility, so that (9) is interpreted:

- (10) Exactly seven guests will bring apple or blackberry and could bring apple and could bring blackberry, and another three guests will bring blackberry or cherry and could bring blackberry and could bring cherry.

(10) is something we know to be true, and gives an intuitive gloss of the target reading of (9).

A similar point can be made with disjunctions of disjunctions, though it requires a bit more theory to see the significance of the observation. The observation, due to Nathan Klinedinst (p.c.), is that asserting a disjunction of disjunctions like (11) can be felicitous:

- (11) Marie will bring apple or blackberry, or she'll bring blackberry or cherry.

If 'or' means \vee , then (11) is equivalent to (12):

- (12) Marie will bring apple or blackberry or cherry.

³Exhaustifying this reading also wouldn't help, since *none* of the guests are such that they will bring apple \vee blackberry, might bring apple, might bring blackberry, and *must* bring one of those.

But in general, sentences which are equivalent to one of their proper parts are infelicitous (Singh, 2007; Schlenker, 2009). Hence the standard explanation of the infelicity of a sentence like (13) is that it is equivalent to its part ‘Marie is in France’, and the standard explanation of the infelicity of a sentence like (14) is that it is equivalent to its part ‘Marie is in Paris or Berlin’:

(13) #Marie is in Paris or France.

(14) #Marie is in Paris or Berlin or Paris.

So the puzzle for a classical theory of disjunctions is why an assertion of (11) can be felicitous. Intuitively, what (11) says is that either Marie is going to a store with apple and blackberry, or a store with blackberry and cherry; in other words, a natural gloss on (11) is in (15):

(15) Either Marie will bring apple or blackberry and could bring either of those, or she’ll bring blackberry or cherry and could bring either of those.

And indeed, if $\lceil p \text{ or } q \rceil$ can be interpreted $(p \vee q) \wedge \diamond_c p \wedge \diamond_c q$, as I am arguing, then this is one interpretation of (11). Note, moreover, that, given this interpretation, (11) is not equivalent to any of its parts. On that approach (assuming that the inner ‘or’s are interpreted with a circumstantial flavor of modality and the main ‘or’ with an epistemic flavor), (11) has the meaning $((A \vee B) \wedge \diamond_c A \wedge \diamond_c B) \vee ((B \vee C) \wedge \diamond_c B \wedge \diamond_c C) \wedge \diamond_e ((A \vee B) \wedge \diamond_c A \wedge \diamond_c B) \wedge \diamond_e ((B \vee C) \wedge \diamond_c B \wedge \diamond_c C)$, which is not the meaning of any proper part of (11).

For a final kind of case we’ll turn to conditionals. Suppose that John is going to buy one pie at a store that has at least one of apple, blackberry, and cherry, but may not have all three. John prefers apple to blackberry to cherry. Given John’s preferences, (16) is assertable:

(16) If John buys apple or blackberry, he’ll buy apple. But, for all we know, John will buy blackberry, since they might not have apple at the store.

But, if ‘or’ means \vee , (16) should have a Moorean flavor. For in that case, the conditional ‘If John buys apple or blackberry, he’ll buy apple’ would entail (given that we know John will buy only one pie) that John won’t buy blackberry. For consider any world w where John buys blackberry. Given that he’s only going to buy one pie, in w he buys blackberry, hence apple \vee blackberry, but not apple; and hence the conditional has a true antecedent and false consequent. But everyone agrees that a conditional with a true antecedent and false consequent is itself false.⁴ So, if ‘or’ meant \vee , an assertion of (16) should be as incoherent as an assertion of (17):

⁴At least as long as the conditional does not itself embed further modals or conditionals, as our key conditionals do not.

- (17) # John will buy apple, not blackberry. But, for all we know, John will buy blackberry, since they might not have apple at the store.

Why is (16) felt to be coherent? A natural gloss on the conditional in (16) is (18):

- (18) If John buys apple or blackberry and they are both possible, then he'll buy apple.

And of course there is no puzzle about why (18) doesn't entail that John won't get blackberry, since the antecedent says nothing about what happens in case apple is not available. So if we can interpret 'A or B' in the antecedent of (16) as $(A \vee B) \wedge \diamond_c A \wedge \diamond_c B$, then our puzzle dissolves.

Note, again, that if we interpreted 'possible' in (18) *epistemically*, we would not derive the target reading, since we are sure that it is *epistemically* possible that John will buy apple and it is epistemically possible that John will buy blackberry. Hence adding these as conjuncts to the antecedent would not block the argument above that (18), together with the information in the context, entails that John won't buy blackberry. Instead, the sense of possibility here must once again be circumstantial.

In response to this case, it's natural to wonder whether intuitions here can be explained via the context-sensitivity of conditionals. Indeed, conditionals are context-sensitive; but that does not help explain intuitions in this case. If 'or' means \vee , then 'If John buys apple or blackberry, then he'll buy apple' entails that John won't buy blackberry *on any resolution of context sensitivity* in the conditional (given that John will buy only one pie). So, on any resolution of context sensitivity, (16) and (17) should be equally incoherent if 'or' means \vee .

We can build on the present paradigm to elicit *probability judgments* which further support the claim that 'or' can be interpreted non-classically. Suppose there are two urns on a table. Urn 1 contains seven orange marbles and three red marbles. Urn 2 contains seven red marbles and three blue marbles. An urn has been chosen at random with equal probability; we don't know which urn was chosen. In a moment, a marble will be chosen from that urn. In this situation, it seems that there is a reading of (19) on which it is reasonable to think it is probably true; and, correspondingly, there is a true reading of (20):

- (19) If an orange or red marble will be chosen, then an orange marble will be chosen.
(20) Probably, if an orange or red marble will be chosen, then an orange marble will be chosen.

If 'or' means \vee , however, then arguments similar to those above show that (19) is in fact unlikely (since it is false at every world where a red marble is chosen, which together comprise half of the state space); and, correspondingly, that (20) is false.⁵ Instead, (19) seems to have

⁵This case superficially resembles some of those discussed in Kaufmann 2004, but it is fundamentally

an interpretation on which it's equivalent to 'If a marble from Urn 1 is chosen, then it will be orange'. And this is the interpretation we obtain if the disjunction is interpreted as I have suggested: as 'the chosen marble is orange \vee red, and could be orange and could be red', in other words, the chosen marble is from Urn 1.

Obviously nothing essential turns on the details of the cases I have given. Here is another set of examples with a similar structure; the theoretical upshots are the same as in the cases we've seen, so I won't walk through them:

(21) *My ten students have to choose one of two summer programs, one physics-themed, with sites in France and Germany; and one philosophy-themed, with sites in France and Spain. Seven students have chosen the physics program, and the other three have chosen the philosophy program, but none have chosen their study sites. I don't know who has chosen which program. In addition, one of Cian's students, Marie, is also choosing between those programs. She is undecided about what program to do; but, ceteris paribus, she prefers to be in Germany over France, and France over Spain. Given this, we know:*

- a. Exactly seven of my students will go to France or Germany, and another three will go to France or Spain.
- b. All my students who are studying in France or Germany will study physics.
- c. Marie will go to France or Germany, or else to France or Spain.
- d. If Marie goes to France or Germany, she'll go to Germany; but she might still go to France, if she chooses the philosophy program.

Further examples are easy to generate on a model with these.

3.1 Accounting for embedded circumstantial possibility

The cases I have given in this section are ones where an embedded sentence with the form ' p or q ' is interpreted as $(p \vee q) \wedge \diamond_c p \wedge \diamond_c q$, where \diamond_c is, again, a circumstantial modal. On my view, ' p or q ' means $(p \vee q) \wedge \diamond p \wedge \diamond q$, where \diamond is a possibility modal whose flavor is determined by context. On this analysis, the target readings brought out in this section are obtained simply by resolving the context sensitivity of \diamond so that it is read circumstantially, as \diamond_c .

In the rest of this section, I'll compare this simple theory of these cases to the one theory I know of which can capture at least some of the data elicited in the last section. The theory I have in mind is from [Geurts 2005](#). Geurts's theory, as I understand it, is like Zimmermann's, except that the flavor of modality is contextually specified: that is, ' p or q ' is true just in case

different: adopting Kaufmann's theory, or anyone else's, of those cases will not help with this one, since the argument here depends only on the validity of Strong Centering, which is not at issue in that debate.

$\Diamond p \wedge \Diamond q$ is true, where \Diamond is a possibility modal whose flavor is determined by context. Hence, for instance, Geurts’s theory nicely predicts that ‘If John buys apple or blackberry, then he’ll buy apple’ does not entail that John will not buy blackberry, since the latter means the same thing as ‘If John can buy apple and John can buy blackberry, then he’ll buy apple’.

While my preferred theory directly builds on Geurts’s (and hence Zimmermann’s) theory, I think it improves on it in important ways: Geurts’s theory is extensionally inadequate in a number of ways. First, it shares Zimmermann’s problems with (8). This is reason enough, I think, to reject the theory, but I will make a few more points here, which will bring out the importance of the difference between my theory and Geurts’s: namely, the addition of the classical disjunction $p \vee q$ as part of the meaning of ‘or’, in addition to the modal conjuncts.

Recall (9):

- (9) Exactly seven guests will bring apple or blackberry, and another three guests will bring blackberry or cherry.

I claimed (9) has a prominent true reading when seven guests will go to a store that carries apple and blackberry, and another three to a store that carries blackberry and cherry; this true reading is, again, well-glossed by the reading our theory predicts of (9) in (10):

- (10) Exactly seven guests will bring apple \vee blackberry and could bring apple and could bring blackberry, and another three guests will bring blackberry \vee cherry and could bring blackberry and could bring cherry.

Geurts’s theory, by contrast, predicts that (9) is equivalent to (22):

- (22) Exactly seven guests could bring apple and could bring blackberry, and another three guests could bring blackberry and could bring cherry.

However, this is not, as far as I can tell, a reading of (9). To see this, suppose that we are sure that seven guests will go to a store that carries exactly apple and blackberry, and another three to a store that carries exactly blackberry and cherry, but we know that some of the guests will decline to buy a pie at all. Then (22) remains true. But in this case, (9) has no true reading, as far as I can tell. In other words, the classical disjunction still plays a role in the interpretation of ‘or’, even when it also has a modal meaning: $\lceil p \text{ or } q \rceil$ entails $\Diamond p \wedge \Diamond q$, but that does not exhaust its meaning; it also entails $p \vee q$.

This is closely related to a second problem with Geurts’s theory: while it predicts a non-classical reading of disjunction, it does not predict a classical reading. But there does seem to be a classical reading of disjunction, in addition to the non-classical readings that I have worked to bring out here. We have seen how to derive the classical reading on my theory by resolving the context sensitivity of \Diamond to the trivial reading \Diamond_{\top} . By contrast, no similar move

is available in Geurts’s theory: reading \diamond trivially would result in a reading of $\lceil p \text{ or } q \rceil$ where it is equivalent to $\diamond_{\top} p \wedge \diamond_{\top} q$, that is to \top , not to $p \vee q$.

Of course, a proponent of a Geurts-like theory could argue there is ambiguity in the meaning of ‘or’. But given that, as we have just seen, we anyways need the classical conjunct $p \vee q$ in order to have a truth-conditionally adequate theory of ‘or’, a better approach is to adopt my theory, on which the classical reading of ‘or’ can emerge as one resolution of disjunction’s context sensitivity.

This observation also lets us avoid what is widely taken to be a fatal objection to the Zimmermann/Geurts approach to disjunction: its predictions about the interpretation of disjunction under negation, and more generally in negative environments. On Geurts’s theory (and Zimmermann’s), a negated disjunction is equivalent to a classical disjunction:

$$\neg(p \text{ or } q) \equiv \neg(\diamond p \wedge \diamond q) \equiv \neg\diamond p \vee \neg\diamond q$$

Given the duality between \diamond and \square , the right-hand side is, in turn, equivalent to $\square\neg p \vee \square\neg q$. But this obviously misses something important about the interaction of disjunction with negation. ‘John is not in Berlin or Paris’ naturally licenses *both* the conclusion that John is not in Berlin and the conclusion that John is not in Paris; whereas on the present view, we can conclude only that *either* John can’t be in Berlin, or (in the classical sense) John can’t be in Paris.

By contrast, on my theory, we have the following equivalence:

$$\neg(p \text{ or } q) \equiv \neg((p \vee q) \wedge \diamond p \wedge \diamond q) \equiv (\neg p \wedge \neg q) \vee (\square\neg p \vee \square\neg q)$$

Now my suggestion is that, when embedded under negation (and, more generally, in downward entailing environments), we tend to resolve the context sensitivity of the modality in disjunction towards a trivial reading, so that the right-hand side ends up equivalent to $(\neg p \wedge \neg q) \vee (\square_{\top}\neg p \vee \square_{\top}\neg q)$ (where \square_{\top} is the dual of \diamond_{\top} , that is, the necessity modal interpreted with a universal accessibility relation). But, just as $\diamond_{\top} p$ is trivially true, likewise its dual $\square_{\top}\neg p$ is trivially false, since p is possible in this broadest sense. Hence, interpreted this way, $\neg(p \text{ or } q)$ comes out equivalent to $\neg p \wedge \neg q$, as desired.

So, on my theory, there is a resolution of the context-sensitivity of ‘or’ on which $\lceil p \text{ or } q \rceil$ ends up being equivalent to $p \vee q$. Write ‘or $_{\top}$ ’ for this resolution of context sensitivity. My suggestion is that we systematically tend to interpret ‘or’ as ‘or $_{\top}$ ’ when ‘or’ appears in negative environments. Why? Well, a trivial reading of disjunction’s modals *weakens* the meaning of the disjunction in a positive context and hence *strengthens* it in a negative context. Across a wide range of theoretical contexts, including the study of implicatures, homogeneity, reciprocity inferences, and bare plurals, it has been suggested that, given a range of available readings,

speakers tend to coordinate on the reading that results in the strongest overall meaning. Hence this is the standard explanation of the fact that ‘Susie had a cookie or an apple’ naturally communicates that she didn’t have both (‘or’ is strengthened with an exclusivity implicature, since this strengthens the overall meaning of the sentence), while ‘Everyone who had a cookie or an apple will get sick’ naturally communicates that everyone who had a cookie or an apple *or both* will get sick (‘or’ is not strengthened with exclusivity, since this would *weaken* the overall meaning of the sentence). Given the breadth of appeal to an interpretive heuristic like this, it seems reasonable to apply it to the case of disjunction as well.

Having said that, this heuristic is only a heuristic. Just as, for instance, exclusivity implicatures can be embedded in downward monotone or non-monotone environments (as in ‘Everyone who had a cookie or a banana can go, but everyone who had both needs to pay me’); so, likewise, non-trivial resolutions of the context sensitivity of ‘or’ can be brought out in downward monotone contexts with the right set-up, as we have seen (e.g., in the restrictor of ‘every’). And, as for non-monotone contexts like the scope of ‘exactly’, what reading we get seems to depend on context, again, just as for implicatures.

In sum: the crucial difference between my theory and Geurts’s is that, on my theory, a disjunction $\lceil p \text{ or } q \rceil$ still has a classical component to its meaning ($p \vee q$) along with the possibility conjuncts ($\diamond p \wedge \diamond q$). This is crucial for capturing the truth-conditions in our target cases, and also for accounting for the fact that disjunction has a classical reading—the default reading in negative contexts—which can be obtained on our approach by resolving context sensitivity so that the possibility conjuncts are trivialized.

4 Free choice

Zimmermann and Geurts’s theories were motivated by *free choice* inferences like these:

- (23) a. You may have an apple or a pear.
 b. So, you may have an apple.
- (24) a. The keys might be upstairs or might be downstairs.
 b. So, the keys might be upstairs.

Schematically, where \diamond is any possibility modal, *narrow scope free choice* is the inference from $\lceil \diamond(p \text{ or } q) \rceil$ to $\lceil \diamond p \text{ and } \diamond q \rceil$, and *wide scope free choice* is the inference from $\lceil \diamond p \text{ or } \diamond q \rceil$ to $\lceil \diamond p \text{ and } \diamond q \rceil$. These inferences are not valid in modal logic, given a classical meaning for ‘or’ and \diamond , nor is it easy to see how to derive them on a pragmatic basis (see [Kratzer and Shimoyama 2002](#) for the most prominent pragmatic approach, and [Fusco 2014](#) for a response).

Zimmermann and Geurts’s idea to capture free choice was this. Since on their theories $\lceil p \text{ or } q \rceil$ means $\diamond p \wedge \diamond q$, $\lceil \diamond(p \text{ or } q) \rceil$ means $\diamond(\diamond p \wedge \diamond q)$. Given the collapse principle $\diamond \diamond p \rightarrow \diamond p$

and $\diamond\diamond q \rightarrow \diamond q$, this entails $\diamond p \wedge \diamond q$ in classical modal logic. Things are similar for wide scope free choice.

This idea has had little uptake in the subsequent literature. My impression, however, is that this is not because of problems local to this account of free choice, but rather due to the disastrous predictions such accounts make about the interaction of disjunction with negation. As we have seen, by contrast, my theory, although it builds closely on Zimmermann and Geurts's, avoids this problem by adding in a classical conjunct to the meaning of 'or', so that $\lceil \text{Not}(p \text{ or } q) \rceil$ is equivalent to $\neg p \wedge \neg q$ provided we resolve the modal's context-sensitivity in a trivial way.

But my theory still accounts for free choice inferences, in just the same way as the Zimmermann/Geurts approach. A sentence with the form of (25-a) is predicted to have the meaning of (25-b):

- (25) a. $\diamond(p \text{ or } q)$
 b. $\diamond((p \vee q) \wedge \diamond_{or} p \wedge \diamond_{or} q)$

I am writing \diamond_{or} for the possibility modal generated by 'or', which does not necessarily have the same flavor as the overt modal \diamond . Given a classical meaning for \diamond , (25-b), in turn, entails $\diamond\diamond_{or} p \wedge \diamond\diamond_{or} q$, since we can distribute the outer \diamond over the conjunction. Likewise, on my theory, a sentence with the form (26-a) has the meaning of (26-b), which entails $\diamond_{or}\diamond p \wedge \diamond_{or}\diamond q$.

- (26) a. $\diamond p \text{ or } \diamond q$
 b. $(\diamond p \vee \diamond q) \wedge \diamond_{or}\diamond p \wedge \diamond_{or}\diamond q$

Now we make the following assumption: in non-downward-entailing environments, when 'or' appears in the vicinity of overt modals, we tend to interpret the modal generated by 'or' with the same flavor of modality as the overt modals. That gets us from the meanings generated above to $\diamond\diamond p \wedge \diamond\diamond q$. The final piece of the puzzle is to assume that, at least defeasibly, we can infer $\diamond p$ from $\diamond\diamond p$. (We need not assume this inference is universally valid, since free choice is not always valid.) Then we have a neat story about the derivation of both wide and narrow scope free choice inferences.

Furthermore, this story will be naturally blocked in negative environments, where, as many have observed, free choice does not naturally arise: the most prominent interpretation of (27) is one on which it is equivalent to 'You may not have an apple and you may not have a banana', not one on which it is equivalent to 'It's not the case that you may have an apple and you may have a banana':

- (27) You may not have an apple or a banana.

To capture this observation, we need only again rely on the assumption I argued for above:

as a default matter, ‘or’ is interpreted classically in negative environments, by resolving its context-sensitivity to the trivial reading \diamond_{\top} of \diamond_{or} .

To capture free choice readings for deontic modals (as in (23)), we must assume that the covert modal in ‘or’ can be interpreted deontically, in addition to epistemically and circumstantially. To argue clearly for the existence of such readings, we would need to find cases where deontic possibility and circumstantial possibility come clearly apart, which is somewhat difficult to do. However, such readings do seem possible. To vary our running case, suppose that ten guests are coming to dinner. All ten are going to the same store, which has abundant apple, blackberry, and cherry pies. So for each guest, it is circumstantially possible that he bring apple, blackberry, or cherry. But seven guests are part of a cult that abjures stonefruit, while another three guests are part of a cult that never eats apples. So seven guests are only deontically permitted to bring apple or blackberry, and another three guests are only deontically permitted to bring blackberry or cherry. In this case, it seems to me that (9), repeated here, has a true reading:

- (9) Exactly seven guests will bring apple or blackberry, and another three guests will bring blackberry or cherry.

Since each guest is such that it is both circumstantially and epistemically possible that they bring apple, blackberry, and cherry, it seems to me that the reading of (9) that we know must be one on which we interpret it as ‘Exactly seven guests will bring apple or blackberry and (deontically) may bring apple and may bring blackberry, while another three guests will bring blackberry or cherry and (deontically) may bring apple and may bring blackberry’.⁶

5 Epistemic possibility without ignorance

At the outset, I noted that Griceans have a natural story about how epistemic possibility inferences arise from *assertions* of disjunctions; I then argued, in §2, that that story does not extend to embeddings. But a different kind of objection to the Gricean derivation of epistemic possibility has recently been developed by Degano et al. (2023), to which my account also suggests a natural response. Recall that the standard Gricean derivation of possibility implicatures goes via *ignorance*: when S asserts ‘ p or q ’, we infer that both p and q are epistemically possible because we think that, if S knew either p or q , she would have asserted that instead (or their conjunction); hence S *is ignorant* of p and ignorant of q ; but S knows $p \vee q$; hence both p and q are epistemically possible for her.

But Degano et al. (2023) show that subjects still conclude q is epistemically possible

⁶This account of free choice does not extend to the superficially similar phenomenon of *simplification of disjunctive antecedents*. The latter may have a different explanation, however, of the sort recently explored in Klinedinst 2024.

when S asserts $\lceil p \text{ or } q \rceil$ *even when it is common ground that S knows p*. In a representative experiment, subjects were shown a depiction of four boxes, three open and one covered (the ‘mystery box’), and were told that the speaker can see what’s inside the three boxes but not the covered box. They were also told that the speaker knows that the mystery box always contains exactly the same contents as one of the open boxes. In the key stimulus, subjects were asked to assess ‘The mystery box contains a yellow ball or a blue ball’ along with a picture that depicts three uncovered boxes, each of which contains a yellow ball, and one of which also contains a blue ball. Hence the speaker is sure that the mystery box contains a yellow ball, and leaves it open that it also contains a blue ball. The key finding is that, in cases like this, subjects judged the assertion to be felicitous, *even though the speaker knows one conjunct to be true*. By contrast, if the set-up is the same but there is no blue ball, subjects find the same assertion to be infelicitous. This suggests that subjects still associate the assertion with the inference that blue is epistemically possible, even though they do *not* derive the inference that the speaker is ignorant about whether the box contains yellow. In a slogan: epistemic possibility inferences are drawn even in the absence of ignorance inferences.

This poses a serious challenge to the standard Gricean derivation of possibility implicatures, since that derivation goes via ignorance. So we need a way to derive epistemic possibility inferences that does not go by way of uncertainty inferences. My theory of disjunction gives us precisely that, since, on that theory, $\lceil p \text{ or } q \rceil$ simply entails $\Diamond p \wedge \Diamond q$.

6 Alternate approaches

I don’t know of any alternate extant theories that can capture the novel patterns I have brought out here. In concluding, I will briefly explore a variety of alternate approaches which have either occurred to me or been suggested by others, explaining why they either don’t work, or else seem interesting but less appealing than the route I have pursued.

6.1 Other non-classical treatments of disjunction

There have been many non-classical theories of disjunction proposed in the literature, in addition to the theories of Zimmermann and Geurts which I have discussed: dynamic (Beaver, 2001), state-based (Simons, 2005; Hawke and Steinert-Threlkeld, 2020), inquisitive (Groenendijk and Roelofsen, 2009), alternative (Hamblin, 1973; Alonso-Ovalle, 2006; Fine, 2017), and team (Aloni, 2022), among others. I won’t go into detail about any of these, however, because, as I noted above, while some provide interesting ingredients for generating *epistemic* possibility inferences, they cannot as far as I can tell generate circumstantial possibility inferences, and so will not be adequate to account for our data.

6.2 Embedded implicatures

It is very natural to think that the possibility conjuncts I have brought out are implicatures. After all, as we have seen, it is well known how to derive *unembedded* epistemic possibility implicatures; it has been frequently observed in the last few decades that implicatures of various kinds can be embedded; and, just as for embedded implicatures, the presence or absence of (non-trivial) possibility inferences arising from disjunction seems to be effected by monotonicity, focus, and the salience of contextual alternatives (‘went to Paris *or* Berlin’ vs. ‘*definitely* went to Berlin’, etc.).

This was indeed my first reaction to these cases. The problem is that no theory of embedded implicatures generates the target readings. We have already explored the unpromising options for this when it comes to epistemic possibility inferences. When we turn to the circumstantial reading, things look even worse: there just isn’t any way that I know of to supplement an extant theory about *EXH* so as to derive the circumstantial reading of ‘ p or q ’. We could explore a similar move as in the epistemic case, but it would be equally problematic. The idea would be to allow free embedding of a covert circumstantial necessity modal \Box_c which then gets exhaustified, so our target disjunction would have the parse $EXH(\Box_c(p \vee q))$. The problem with this approach is again overgeneration: just as ‘John isn’t in Paris’ can’t be interpreted as ‘I don’t know that John is in Paris’, likewise it can’t be interpreted as ‘John doesn’t have to be in Paris’. By contrast, as far as I can tell, similar worries about overgeneration do not apply to my account: ‘It’s not true that John is in Paris *or* Berlin, rather, he’s in Paris’ does have a coherent reading, just as my account predicts (one where the negation targets the possibility that John is in Berlin; it also has an incoherent reading, as my account also predicts).

A different approach would be to enrich *EXH* so that, whenever p is among the alternatives that *EXH* considers, so are $\Box p$ (and/or $\Diamond p$, where \Box and \Diamond are modals whose flavor is determined by context). I think this is a promising approach to derive the target readings, though I think that doing so requires an innocent inclusion (Bar-Lev and Fox, 2020) or recursive (Fox, 2006) approach to exhaustification. While this approach seems worth exploring in detail, I have two worries about it. First, the resulting theory of *EXH* will be so complex as to be somewhat unappealing, since this approach greatly increases the number of alternatives that must be considered. What’s more, as we will see shortly, my theory of ‘or’ also accounts for the free choice inferences that motivate the extra complexity of innocent inclusion or recursive exhaustification in the first place (indeed, as we will see, it does so more adequately than *EXH*-based theories). Second, there remain serious worries about overgeneration: would this approach still predict that ‘John isn’t in Paris’ has a reading where it means ‘John isn’t necessarily in Paris’, obtained by exhaustifying under negation? This depends on the details of implementation, but it seems at the very least that it would take considerable footwork to derive the target readings without overgeneration.

A different approach is to look for some other covert operator that can ‘embed implicatures’, in some broad sense, that serves our purposes. While something like that might indeed work, the only extant candidate I know of comes from an intriguing proposal in [Blumberg and Goldstein 2021](#) for a redundancy operator. However, while such an operator might help generate epistemic possibility inferences, I do not see how it could help derive circumstantial possibility inferences from disjunctions.

6.3 Embedded modals

A different approach would be to try to derive our target readings via covert possibility modals, so that, for instance, a sentence with the surface form ‘If Mark buys apple or blackberry, he’ll buy apple’ has the logical form $(\diamond A \vee \diamond B) > A$. Then, if we have some theory of free choice, we derive the reading. $(\diamond A \wedge \diamond B) > A$. Even granting some theory of free choice to get us to this point, the problem is that this derives readings like those we get on Geurts/Zimmerman’s theory, which are missing the classical component of disjunction (in this case, $A \vee B$). As we have seen, such readings are not in general available, and so this approach is empirically inadequate. More simply, there also remains a very substantial risk of overgeneration: if we can freely insert covert possibility modals into LFs, what prevents us from interpreting ‘John is in Paris’ as ‘John might be in Paris’?

6.4 Trivalence

Another option would be to adopt a trivalent approach like the dynamic theory developed in [Goldstein 2019](#) to account for free choice. Here is a version of Goldstein’s idea in a static framework:

$$\llbracket p \text{ or } q \rrbracket^{c,w} = \begin{cases} 1 & \llbracket p \vee q \rrbracket^{c,w} = \llbracket \diamond p \rrbracket^{c,w} = \llbracket \diamond q \rrbracket^{c,w} = 1 \\ 0 & \llbracket p \vee q \rrbracket^{c,w} = 0 \\ \# & \textit{otherwise} \end{cases}$$

Departing from Goldstein, we can stipulate that the flavor of \diamond is determined by the context. In other words, disjunction has the *truth* conditions that I have proposed, and the *falsity* conditions of classical negation.

This does a very nice job of capturing some of the cases I have brought out, and a nice job of capturing the default interaction of ‘or’ with negation. But it does not have enough flexibility to capture the full range of readings we have observed. On the one hand, it cannot obviously capture classical interpretations of ‘or’ in positive environments. For instance, suppose everyone went to a store with exactly one kind of pie, and everyone turned up with apple or with blackberry. Then there is a true reading of ‘Everyone brought apple or blackberry’,

where, intuitively, we interpret ‘or’ as \vee . We have seen how to derive this reading on my view: namely, by resolving the context-sensitivity in ‘or’ trivially. By contrast, it’s not obvious how to do so on the present theory. And, conversely, it cannot capture non-classical interpretation of ‘or’ in negative environments, as in ‘John isn’t in Berlin *or* Paris, he’s certainly in Berlin’.

There is room for maneuver here. For instance, an advocate of this view could follow my preferred approach in positing that \diamond_{\top} is one admissible resolution of the context-sensitivity of the modal \diamond , yielding a classical meaning for ‘or’ as one possibility in positive environments. However, once we have that possibility on the table, I don’t see what the payoff is of bringing trivalence into the picture. Alternately, we could add to our grammar a covert operator B which, given a sentence p with a trivalent meaning, yields a bivalent sentence Bp which is true whenever p is either true or undefined, and false otherwise, so that $\lceil B(p \text{ or } q) \rceil$ means $p \vee q$. But adding an operator like this is a cost. This would be, to my knowledge, a novel posit; by contrast, the well-known Bochvar floating- A operator takes a trivalent sentence p and returns a bivalent sentence Ap which is true whenever p is true and false whenever p is false or undefined. The Bochvar operator is what we would use to derive the non-classical meaning of ‘or’ in negative environments.

So adding trivalence to the picture does not seem to buy us anything in terms of coverage, and costs a fair amount in terms of complexity.

6.5 Error theory

A final approach would be to dismiss all the judgments I have elicited in this paper as errors of reasoning rather than reasonable targets of semantic theory. To see the appeal of this response, consider the conjunction fallacy, where subjects rank a conjunction $\lceil p \text{ and } q \rceil$ as being more likely than a single conjunct p (Tversky and Kahneman, 1983). This violates probability theory assuming that ‘and’ means \wedge . One response to this finding would be to propose a non-classical meaning for ‘and’, which yields a model where $\lceil p \text{ and } q \rceil$ is more likely than p . This is not a popular response, however; instead, it seems much more appealing to ascribe to subjects an error of some kind rather than vindicate their judgments with a new theory of ‘and’. In general, then, we cannot be overly quick to reach for semantic accounts of all patterns of speaker intuitions.

However, in this case, I don’t see a particularly natural way of spelling out an error theory in this domain. While I don’t have a decisive objection to such a theory, I want to emphasize a few things about this line of response.

First, error theories are generally best situated to account for cases where reflective judgments are systematically corrigible. In the case of the conjunction fallacy, most subjects will agree, on reflection, that there’s no way that $\lceil p \text{ and } q \rceil$ can be strictly more likely than p ; the puzzle that remains is to explain why they ever thought otherwise. By contrast, the judgments

I have elicited about disjunction seem more robust. If seven people are going to a store with apple and blackberry, and three people are going to a store with blackberry and cherry, it seems clearly assertable, after careful reflection, that exactly seven people will get apple or blackberry. There is clearly a reading of this on which it is not assertable, too, which can be brought out by emphasizing that some of the other three might get blackberry, and hence blackberry or apple. That is a reading which we also can account for (the reading of ‘or’ as ‘ or_\top ’). But the existence of this reading does not seem to crowd out the existence of the assertible one.

Second, an obvious point: an error theory needs to be a *theory*. We cannot just dismiss speaker judgments as erroneous and stop there; we need a predictive theory of what leads to those judgments. In particular, in this case an error theory would need to account for facts about *embeddings*. Unlike in the case of the conjunction fallacy, where subjects were asked to reason about unembedded conjunctions, the cases I have discussed are crucially ones where a disjunction is interpreted non-classically in *embedded* configurations. While an error theory of such configurations is of course still possible, it is not at obvious what it would look like.

Finally, an error theory needs to account for fine-grained *contrasts* in our data. Recall (16), repeated here:

- (16) If John buys apple or blackberry, he’ll buy apple. But, for all we know, John will buy blackberry, since they might not have apple at the store.

My claim was that (16) has a coherent reading. An error theory would say that this is because, even though the first sentence *in fact* entails that John won’t buy blackberry (given that he only buys one pie), speakers simply fail to draw this inference. Now compare (16) to the minimal variant in (28) which replaces ‘John buys apple or blackberry’ with ‘John doesn’t buy cherry’, which is *contextually equivalent* to ‘John buys apple or blackberry’, given that we know John will buy exactly one of apple, blackberry, and cherry, and assuming that ‘or’ is classical:

- (28) If John doesn’t buy cherry, he’ll buy apple. But, for all we know, John will buy blackberry, since they might not have apple at the store.

(28) seems strikingly less coherent than (16). The first sentence of (28) seems to clearly entail that John won’t buy blackberry: cherry and apple are the only options. Insofar as (28) has a coherent reading, the second sentence is felt to revise one of the commitments of the first.

The contrast between (16) and (28) is immediately explained on my theory: (16) has an interpretation where we interpret the antecedent as $(A \vee B) \wedge \diamond_c A \wedge \diamond_c B$, since that is one resolution of the context-sensitive meaning of ‘or’; on that interpretation, the first sentence of (16) does not entail $\neg B$. By contrast, (28) has no such interpretation, since negation remains

classical on my theory, and hence the first sentence of (28) entails $\neg B$, given that John will only buy one pie. In other words, the availability of the readings I have brought out seems to depend on exactly what words are used. That is what we would predict given a semantic theory of the contrast. Of course, an error theory might be able to predict this contrast too, but I don't see how such a story would go.

In sum: it is always worth considering whether the best theory of a given pattern of judgments is one that dismisses them as the result of a systematic error rather than as evidence of underlying knowledge of the meaning of connectives. In the present case, however, it is not at all obvious that such a theory can be developed in a systematic way that captures all of our observations.

7 Conclusion

I have argued that $\lceil p \text{ or } q \rceil$ means $(p \vee q) \wedge \diamond p \wedge \diamond q$, where \diamond is an existential modal whose flavor is determined by context, which can be interpreted epistemically, circumstantially, deontically, or trivially. Since the last interpretation yields a reading for $\lceil p \text{ or } q \rceil$ where it is contextually equivalent to $p \vee q$, this theory can account for the existence of classical interpretations of 'or' (as, for instance, in its prominent default interpretation in negative contexts). Other resolutions of the context sensitivity of 'or' account for the range of new patterns I have brought out here, where (embedded) disjunctions give rise to (embedded) possibility inferences of epistemic, circumstantial, and deontic flavors. Moreover, this approach accounts for wide and narrow scope free choice inferences, and for epistemic possibility inferences without uncertainty.

Other theories of the same data may be possible; indeed, I have flagged some possible alternatives along the way. But I do not know of any extant theory that can account for all the observations here. Strikingly, even though the data I have surveyed seem to be in the same vicinity as free choice inferences, existing theories of free choice do not account for them. By contrast, my theory accounts for both these data and for free choice. Its breadth of coverage, and the simplicity of its account of the myriad connections between disjunction and possibility, make it a plausible candidate for the meaning of 'or'.

OPEN QUESTIONS

References

- Aloni, M. (2022). Logic and conversation: The case of free choice. *Semantics & Pragmatics*, 15.
- Alonso-Ovalle, L. (2006). *Disjunction in alternative semantics*. PhD thesis, University of Massachusetts, Amherst.

- Bar-Lev, M. E. and Fox, D. (2020). Free choice, simplification, and Innocent Inclusion. *Natural Language Semantics*, 28:175–223.
- Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications: Stanford, CA.
- Blumberg, K. and Goldstein, S. (2021). A semantic theory of redundancy. Manuscript.
- Chierchia, G., Fox, D., and Spector, B. (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In Portner, P., Maienborn, C., and von Stechow, K., editors, *Handbook of Semantics*. Mouton.
- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2018). *Inquisitive Semantics*. Oxford University Press.
- Degano, M., Marty, P., Ramotowska, S., Aloni, M., Breheny, R., Romoli, J., and Sudo, Y. (2023). The ups and downs of ignorance. Manuscript.
- Fine, K. (2017). Truthmaker semantics. In Hale, B., Wright, C., and Miller, A., editors, *Blackwell Companion to the Philosophy of Language*. Blackwell.
- Fox, D. (2006). Free choice and the theory of scalar implicatures. unpublished manuscript.
- Fusco, M. (2014). Free choice permission and the counterfactuals of pragmatics. *Linguistics and Philosophy*, 37(4):275–290.
- Geurts, B. (2005). Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics*, 13(4):383–410.
- Goldstein, S. (2019). Free choice and homogeneity. *Semantics and Pragmatics*, 12(23):1–47.
- Grice, P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge.
- Groenendijk, J. and Roelofsen, F. (2009). Inquisitive semantics and pragmatics. In Larrazabal, J. M. and Zubeldia, L., editors, *Meaning, Content, and Argument: Proceedings of the ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric*. www.ilcl.uva.nl/inquisitive-semantics.
- Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language*, 10:41–53.
- Hawke, P. and Steinert-Threlkeld, S. (2020). Semantic expressivism for epistemic modals. *Linguistics and Philosophy*.
- Kaufmann, S. (2004). Conditioning against the grain. *Journal of Philosophical Logic*, 33:583–606.
- Klinedinst, N. (2024). Coordinated ifs and theories of conditionals. *Synthese*, 203(70).
- Kratzer, A. and Shimoyama, J. (2002). Indeterminate pronouns: The view from Japanese. 3rd Tokyo Conference on Psycholinguistics.

- Mascarenhas, S. and Picat, L. (2019). ‘Might’ as a generator of alternatives: the view from reasoning. In Blake, K., Davis, F., Lamp, K., and Rhyne, J., editors, *Proceedings of SALT*, volume 29, pages 549–561.
- Meyer, M.-C. (2013). *Ignorance and Grammar*. PhD thesis, Massachusetts Institute of Technology.
- Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics*, 2(3):1–78.
- Simons, M. (2005). Dividing things up: The semantics of *or* and the modal/*or* interaction. *Natural Language Semantics*, 13(3):271–316s.
- Singh, R. (2007). Assertability constraints and absurd assertions. Unpublished manuscript, MIT.
- Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, 5(3):269–86.
- Tversky, A. and Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4):293–315.
- Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8(4):255–290.