Disjunction and possibility*

Matthew Mandelkern
New York University, mandelkern@nyu.edu

December 24, 2024

Abstract

I explore cases where ⌜p or q⌝ is interpreted as $(p \vee q) \wedge \Diamond p \wedge \Diamond q$, where $\Diamond$ is a possibility modal whose flavor is either epistemic, circumstantial, or deontic. I argue that no extant theory can account for these interpretations. I propose that the best way to do so is with a direct theory on which ⌜p or q⌝ simply *means* $(p \vee q) \wedge \Diamond p \wedge \Diamond q$. In addition to accounting for these novel cases, the resulting theory also explains both wide- and narrow-scope free choice inferences, in a similar way to the theories of Zimmermann (2000); Geurts (2005); and also accounts for recent observations about the relation between disjunction and possibility from Degano et al. 2023; Feinmann 2023.

1 Introduction

I develop and motivate a theory of disjunction on which ⌜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. I motivate my theory by introducing a variety of cases where an embedded disjunction ⌜p or q⌝ is interpreted as $(p \vee q) \wedge \Diamond p \wedge \Diamond q$, where $\Diamond$ is either an epistemic possibility modal $\Diamond_e$ (§2) or a circumstantial possibility modal $\Diamond_c$ (§3) or a deontic possibility modal $\Diamond_d$ (§5). I argue that the best way to account for these patterns is the simplest one: maintaining that $(p \vee q) \wedge \Diamond p \wedge \Diamond q$ just is the meaning of ⌜p or q⌝ in general. I compare my proposal to alternate approaches (§2.3,§4), in particular important precedents in Zimmermann 2000; Geurts 2005, arguing that my approach improves on those in key ways (§4.1). Finally, I show how my theory yields an extremely simple account of both wide- and narrow-scope free choice inferences (§5), as well as of the derivation of possibility inferences form asserted disjunctions (§6).

*Many thanks to Nathan Klinedinst for extensive and insightful discussion, and for the tremendously helpful observations I note in the paper. Many thanks also to Justin Bledin, Sam Carter, Cian Dorr, Patrick Elliott, Danny Fox, Melissa Fusco, Simon Goldstein, John Hawthorne, Ben Holguín, Salvador Mascarenhas, Guillermo del Pinal, Jacopo Romoli, Richard Roth, Ginger Schultheis, 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 Epistemic possibility

Throughout, I use $\vee, \wedge$, and $\neg$ for classical disjunction, conjunction, and negation; $\Diamond$ for an unspecified possibility operator, and $\Diamond_e, \Diamond_c, \Diamond_d$ for epistemic, circumstantial, and deontic resolutions of $\Diamond$. I begin with cases where an embedded disjunction with the form $\ulcorner p$ or $q\urcorner$ is interpreted as $(p \vee q) \wedge \Diamond_e p \wedge \Diamond_e q$.

Before launching into cases, a caveat. Annoyingly, there seems to be some degree of interpersonal variation in accessing the readings I bring out here: many people have told me they get some but not all of the readings I report. However, every reading I report has been confirmed by many informants. More exploration of the empirical picture will surely be needed, but I am fairly sure all the judgments I report are possible.

2.1 Epistemic possibility from asserted disjunction

Start with a standard observation about disjunction: *asserting* $\ulcorner p$ or $q\urcorner$ can give rise to epistemic possibility inferences. Suppose I assert (1):

(1) John went to Paris or to Berlin.

It's then very natural to infer 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 1975; Stalnaker 1975). Here's the story, schematically, under the classical assumption that $\ulcorner p$ or $q\urcorner$ means $p \vee q$. I asserted $p \vee q$, and hence believe it. (Or know it—it doesn't matter for present purposes what norm of assertion we assume. Likewise, it doesn't matter, throughout this paper, if the possibilities in question are really epistemic or just doxastic.) If I believed p, then I would have asserted p, since it is logically stronger than $p \vee q$ and relevant to the question of $p \vee q$. So I don't believe p; hence there must be some $\overline{p}$-world compatible with my beliefs. Since all the worlds compatible with my beliefs are $p \vee q$-worlds, q is true there. Parallel reasoning shows that p is compatible with my beliefs, too.

Concretely, suppose I have asserted (1). If I believed that John went to Paris and (1) is relevant, then I should have just asserted 'John went to Paris', which is strictly stronger than (1) on a classical theory, and which is relevant if (1) is. Likewise, if I think that John went to Berlin, I should have said that. 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.

2.2 Embedded epistemic possibility

This much is commonplace (though not uncontested, a point we'll return to in §6). Our first new observation, due to Nathan Klinedinst is that *embedded* disjunctions can also give rise to embedded epistemic possibility inferences. This observation is significant because the Gricean story just sketched for deriving epistemic possibility from an *asserted* disjunction does not extend to embedded disjunctions.

To see this, consider the following case:

(2) *A novel virus has been detected in Paris. Cian's, Kit's, and Hartry's students have just returned from a trip to Europe, on which every student visited exactly one of Berlin or Paris. 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's and Hartry's students went where. The department's 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.

The sequence in (2-a) has a coherent reading (brought out most easily by accenting 'or'), intuitively equivalent to (3):

(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.

If 'or' meant $\vee$, however, then (2-a) would be incoherent: for any student S who definitely went to Berlin went to Berlin $\vee$ Paris, so the two sentences in (2-a) would put contradictory constraints on S, namely, that she must be quarantined and also can be released.[1]

Hence the disjunction 'went to Paris or Berlin' in (2-a) appears to have a reading where it means $(P \vee B) \wedge \Diamond_e P \wedge \Diamond_e B$. In other words, the epistemic possibility inferences that an *assertion* like (1) gives rise to can apparently also be *embedded* in the restrictor of a quantifier.

(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. The *incoherent* reading can be elicited by insisting that anyone who definitely went to Berlin went to Berlin or Paris. That both a coherent and classical reading exist is an important observation which any adequate theory will have to account for, and which I return to presently, in §2.4. In §4.1, I will say something about the distribution of readings, giving general reasons to think that non-classical interpretations of 'or' are more prominent in upward entailing contexts,

[1] I 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 to relevant judgments, but results in a clunkier sentence.

while the classical interpretation of 'or' is more prominent in downward entailing contexts, but that both readings can be accessed in both environments with the right kind of set-up. This generalization predicts that the coherent reading of (2-a) should be its less prominent reading, but that this reading should be accessible with the right contextual set-up and contrastive focus, which, indeed, seems to match observation.

The reading of $\ulcorner p$ or $q\urcorner$ as $(p \vee q) \wedge \Diamond_e p \wedge \Diamond_e q$ can be brought out in other environments. For instance, suppose Philippa is one of the students who went to Europe:

(4) If Philippa went to 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); insofar as it is, it's because the antecedent of the first conditional is interpreted as 'Philippa went to either Berlin or Paris and both are possible'.

Suppose next that Cian, Kit, and Hartry each have exactly five students in their classes, and their classes are disjoint. 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's 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:

(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 both logic and 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 two will take probability.

Since eight students, not three, will take logic $\vee$ semantics, the true reading of (6-a) must be one where 'will take either logic or semantics' is interpreted as 'will take logic or semantics and might take either'.

In these cases, $\ulcorner p$ or $q\urcorner$ is apparently interpreted as $(p\vee q)\wedge\Diamond_e p\wedge\Diamond_e q$. My preferred theory of these cases is simple: 'or' *just means* $\lambda p.\lambda q.(p \vee q) \wedge \Diamond p \wedge \Diamond q$, with the flavor of $\Diamond$ resolved by context. One way that the context-sensitivity of $\Diamond$ can be resolved is epistemically, in which case we get the readings I have brought out here.

2.3 Alternate approaches

I will continue to motivate and develop this approach throughout the paper. But first, I'll say a little about the significance of the embedded epistemic possibility inferences I have brought out, and why they are very difficult to explain on existing theories.

Return first to the Gricean theory. It is, again, well-known that an *assertion* of $\ulcorner p$ or $q\urcorner$ can give rise to the epistemic possibility claims $\Diamond_e$ and $\Diamond_e q$, and we have seen how this can be derived 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. For that reasoning crucially turned on the premise that the speaker *asserted* $\ulcorner p$ or $q\urcorner$ and hence *believes* $p \vee q$; when the speaker has not asserted $p \vee q$, but rather a more complex sentence which embeds (but does not entail) $p \vee q$, that reasoning will not go through.

There is 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' (see Chierchia et al. 2012 for an overview). 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 $\ulcorner p$ or $q\urcorner$, 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 in Paris' doesn't have a reading equivalent to 'I don't know that John is Paris'; if it did, then you could assert 'John isn't in Paris', 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 $\neg KP$. So, while having *matrix* K operators is theoretically viable, having embeddable K operators is a non-starter; but that's what we would need to derive embedded epistemic possibility inferences. (Perhaps we could find a way to rule out embedding K here, while allowing it to be embedded in the environments we've looked at. I'm skeptical this will be possible, but even if it is, such a strategy would not extend to the other cases we'll look at below.)

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 might need an account of embeddable free choice inferences on which $\Diamond_e(p \vee q)$ can entails $\Diamond_e p \wedge \Diamond_e q$—for instance, by adding into the LF an *EXH* operator with a semantics that can derive free choice (as in Fox 2006; Bar-Lev and Fox 2020). This approach obviously has exactly the same overgeneration problem as the embedded K approach: 'John isn't in Paris' can't mean 'John might not be in Paris'. Moreover, it is instructive to note that this approach would not capture the target reading of $\ulcorner p$ or $q \urcorner$, which is not $\Diamond_e p \vee \Diamond_e q$ or $\Diamond_e p \wedge \Diamond_e q$ but rather the *conjunction* of the latter with the classical disjunction, that is, $\Diamond_e p \wedge \Diamond_e q \wedge (p \vee q)$. 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 and might have gone to Paris, while another five definitely went to Berlin.

Suppose there is a sixteenth student who in fact went to Madrid, but we don't know that; she might have gone to Berlin or Paris. So ten students went to Berlin $\vee$ 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 did not go to Berlin or Paris, nor did she definitely go to Berlin); while (7) seems clearly false in this case, since eleven, not ten, students might have gone to Berlin or Paris. The classical disjunction $p \vee q$ is a *part* of the meaning of $\ulcorner p$ or $q \urcorner$, even if it is not the whole meaning.

This point is also problematic for a different analysis, due to Zimmermann 2000. On that view, $\ulcorner p$ or $q \urcorner$ simply means $\Diamond_e p \wedge \Diamond_e q$. This faces exactly the same objection: it predicts (5) and (7) to be equivalent, when they are not. I will have more to say about Zimmermann's view, and another in the same ballpark, in §4.1.

2.4 The classical reading of 'or'

For now, let me return to the important observation that 'or' also has a classical reading, one which renders the sequences above incoherent. Here is how my view can capture this observation. Consider the interpretation of $\Diamond$ on which every world accesses *every possible world whatsoever*. Then, as long as p is possible *in some sense*, no matter how distant or etiolated, $\Diamond p$ will be true. I'll write $\Diamond_\top$ for this *trivial* interpretation of $\Diamond$.

Note that even a contradiction can be possible in some linguistically relevant sense (say, epistemically possible for an agent who leaves it open that there are true contradictions). Indeed, we could spell out $\Diamond_\top$ linguistically, so that p is possible whenever p is a sentence of natural language. Then $\Diamond_\top p$ is equivalent to $\top$ for any p. There is clearly more to say about this interpretation of $\Diamond$, but for now I just want to note that if we resolve the context sensitivity of $\Diamond$ this way, then $\ulcorner p$ or $q\urcorner$ will mean $(p \vee q) \wedge \Diamond_\top p \wedge \Diamond_\top q$, which, given the present assumption, is equivalent to the classical disjunction $p \vee q$. Thus a classical interpretation of disjunction, which I'll write 'or$_\top$', falls out of my approach *as one resolution of the context sensitivity of disjunction*, along with epistemic, circumstantial and deontic interpretations. I'll have more to say about *when* such a reading arises in §4.1.

3 Embedded circumstantial possibility

A natural question to ask at this point is whether the non-classical readings of disjunction can be accounted for with one of a number of existing 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 the readings I have brought out so far. Instead of exploring these accounts in detail here, I will present new data which I think rule out a response along any of these lines. These data show that $\ulcorner p$ or $q\urcorner$ can be interpreted as $(p \vee q) \wedge \Diamond_c p \wedge \Diamond_c q$, where $\Diamond_c$ is, again, circumstantial possibility. I know of no account that can derive such a reading. In particular, the accounts just mentioned tie disjunction to *epistemic* possibility, but not to circumstantial possibility, and so I do not think any of them can account for the patterns I will bring out in this section.

3.1 'Exactly'

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 only apple and blackberry pies, and another three guests are going to a store that carries only blackberry and cherry pies. I don't know which guest is going to which store. Given this set-up, there is a prominent reading of (8) which I know to be true:

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

The reading of (8) 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 $\ulcorner p$ or $q\urcorner$ is interpreted as $(p \vee q) \wedge \Diamond_e p \wedge \Diamond_e q$. For 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.[2]

Instead, the target reading seems to be one where we interpret $\ulcorner A$ or $B\urcorner$ as $(A \vee B) \wedge \Diamond_c A \wedge \Diamond_c B$, so that (8) is interpreted:

(9) 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.

We know (9) to be true, and I think (9) gives an intuitive gloss of the target reading of (8). (There is, again, a reading of (8) which we don't know to be true, where 'or' is interpreted classically; the important point is the existence of the reading we know to be true.)

3.2 Disjoined disjunctions

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 (10) can be felicitous:

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

If 'or' means $\vee$, then (10) is equivalent to (11):

(11) Marie will bring apple or blackberry or cherry.

But in general, a sentence which is equivalent to a proper part of that sentence is infelicitous (see Singh 2007; if 'or' meant $\vee$, (10) would also be infelicitous on a local context theory of redundancy like that of Schlenker 2009). Hence a natural explanation of the infelicity of a sentence like like (12) is that it is equivalent to its part 'Marie is in Paris or Berlin':

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

So the puzzle for a classical theory of disjunctions is why an assertion of (10) can be felicitous. Intuitively, what (10) 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 (10) is in (13):

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

(13) 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 $\ulcorner p$ or $q\urcorner$ can be interpreted $(p \vee q) \wedge \Diamond_c p \wedge \Diamond_c q$, as I am arguing, then this is one interpretation of (10). Note, moreover, that, given this interpretation, (10) 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), (10) 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 (10).

3.3 Conditionals

The same reading can again be observed with disjunctions in conditional antecedents.[3] Suppose 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. Then it seems we know (14) (on at least one reading):

(14) 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.

(14) has a coherent reading, where the conditional is naturally glossed:

(15) If John buys apple or blackberry and can buy either one, then he'll buy apple.

But, if 'or' means $\vee$, (14) would be incoherent. For in that case, the conditional would entail that John won't buy blackberry. To see this, consider any world compatible with what we know where John buys blackberry and hence apple $\vee$ blackberry; since we know John will only buy one pie, he doesn't buy apple there; hence the conditional has a true antecedent and false consequent, and so is false.[4]

So, once again, we seem to have a case where $\ulcorner A$ or $B\urcorner$ is interpreted in the antecedent of the conditional in (14) as $(A \vee B) \wedge \Diamond_c A \wedge \Diamond_c B$. Note that, once again, the flavor of modality here needs to be circumstantial, not epistemic, 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 incoherence of (14).

Probability judgments provide further support for the existence of a circumstantial reading

[3] These constructions have received a fair amount of attention in the conditionals literature, but for completely different reasons.

[4] Pretty much everyone agrees that conditionals with true antecedents and false consequents are false as long as the conditional does not itself embed further modals or conditionals, which this conditional does not. Appealing to the context-sensitivity of conditionals does not help at all here, since if 'or' means $\vee$, then (14) will be incoherent *on any resolution of context sensitivity* of the conditional.

of 'or'. 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; we don't know which. In a moment, a marble will be drawn at random from that urn. In this situation, it seems that there is a reading of (16) on which it is reasonable to think it is probably true; and, correspondingly, there is a true reading of (17):

(16) If an orange or red marble will be chosen, then an orange marble will be chosen.

(17) 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 (16) only has a reading on which it is unlikely (since it is false at every world where a red marble is chosen, which comprise half of the state space); and, correspondingly, that (17) is false.[5] Instead, (16) seems (on the relevant reading) to mean:

(18) If an orange or red marble is chosen and both could be chosen (i.e., the marble is from Urn 1), then an orange marble will be chosen.

3.4 Another example

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.

(19) *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, all the following seem assertable:*

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 going to study in France or Germany will study physics.
c. Marie will go to France or Germany, or else to France or Spain.

[5] This case superficially resembles those discussed in Kaufmann 2004, but it is in fact very 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 principle that $p \wedge q$ entails $p > q$, which is not at issue in that debate.

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.

3.5 My analysis

We have seen a wide range of cases where an embedded sentence with the form $\ulcorner p$ or $q \urcorner$ is interpreted as $(p \vee q) \wedge \Diamond_c p \wedge \Diamond_c q$. My analysis of these cases is very simple. On my view, $\ulcorner p$ or $q \urcorner$ means $(p \vee q) \wedge \Diamond p \wedge \Diamond q$, where $\Diamond$ is a possibility modal whose flavor is determined by context. The target readings brought out in this section are obtained simply by resolving the context sensitivity of $\Diamond$ so that it is read circumstantially.

4 Alternate approaches

I don't know of any extant theories that can capture the novel patterns I have brought out here. In the rest of this section, I will briefly explore a variety of superficially promising approaches and explain why they do not work.

4.1 The modal list theory and negative environments

I'll start with the *modal list* theory of disjunction from Geurts 2005, which is the closest theory in the literature to the one I am advocating. Geurts's theory is like the theory of Zimmermann's briefly discussed in §2.3, except that the flavor of modality is lexically underspecified: on Geurts's theory, $\ulcorner p$ or $q \urcorner$ means $\Diamond p \wedge \Diamond q$, where $\Diamond$ is a possibility modal whose flavor is determined by context. My preferred theory builds on Geurts's (and hence Zimmermann's), but attributes to disjunction a strictly stronger meaning: on my theory, $\ulcorner p$ or $q \urcorner$ is the conjunction of $\Diamond p \wedge \Diamond q$ (Geurts's disjunction) *along with* the classical disjunction $p \vee q$. This turns out to be crucial for a number of reasons.

First, and most simply, $\ulcorner p$ or $q \urcorner$ clearly *does* entail $p \vee q$, in line with my theory, but pace Geurts. If I say that John is in Paris or Berlin, then you will judge me to have spoken falsely if John turns out to be in Madrid—no matter that it was possible, in whatever senses, that John was in Paris and that John was in Berlin. Zimmermann proposes that these judgments can be explained pragmatically: given the list of possibilities entailed by a disjunction, we infer that the list is exhaustive, and hence that the corresponding classical disjunction is true. But while this might suffice to account for some judgments, it suggests that the inference from $\ulcorner p$ or $q \urcorner$ to $p \vee q$ is defeasible, whereas in fact the inference seems as universally valid as any inference.

Next, Geurts' theory shares Zimmermann's problems with (5) and (7), which both theories predict to be equivalent, when they are plainly not. This is reason enough 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.

First, a structurally identical problem arises in the circumstantial cases: the readings we have brought out are ones where $\ulcorner p$ or $q\urcorner$ means $(p \vee q) \wedge \Diamond_c p \wedge \Diamond_c q$, *not* $\Diamond_c p \wedge \Diamond_c q$. For instance, Geurts' theory predicts (8) to be equivalent to (20):

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

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

But I don't think (20) is a reading of (8). 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 (20) remains true. But in this case, (8) has no true reading.

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. 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: resolving the context-sensitivity of $\Diamond$ to $\Diamond_\top$ would result in a reading of $\ulcorner p$ or $q\urcorner$ where it is equivalent to $\Diamond_\top p \wedge \Diamond_\top q$, which is equivalent to $\top$, not to $p \vee q$.

Of course, a proponent of the modal list theory could argue there is ambiguity in the meaning of 'or', which can either mean the conjunction of possibilities or the classical disjunction; indeed, Geurts himself suggests something like this to account for the interaction of 'or' with negation. But given that, as we have just seen, we need the classical conjunct $p \vee q$ in order to have a truth-conditionally adequate theory of 'or' *even when the possibility conjuncts are not idle*, 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, rather than through a distinct lexical entry.

Doing so also lets us avoid what is widely taken to be a fatal objection to the modal list approach: its predictions about the interpretation of disjunction under negation, and more generally in negative environments. On the modal list approach, a negated disjunction is equivalent to a classical *disjunction*:

$$\neg(p \textit{ or } q) \equiv \neg(\Diamond p \wedge \Diamond q) \equiv \neg\Diamond p \vee \neg\Diamond q$$

This obviously misses something important about the interaction of disjunction with negation. 'John is not in Berlin or Paris', on its most prominent reading, entails *both* that John is not in Berlin and that John is not in Paris; whereas on the modal list view, we can conclude only the disjunction that either John can't be in Berlin or he can't be in Paris.

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

$$\neg(p \ \mathit{or} \ q) \equiv \neg((p \vee q) \wedge \Diamond p \wedge \Diamond q) \equiv (\neg p \wedge \neg q) \vee (\Box \neg p \vee \Box \neg q)$$

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 the trivial reading $\mathrm{or}_\top$, so that the right-hand side ends up equivalent to $(\neg p \wedge \neg q) \vee (\Box_\top \neg p \vee \Box_\top \neg q)$ (where $\Box_\top$ 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, conversely its dual $\Box_\top \neg p$ is trivially false. Hence, interpreted this way, $\neg(p \ \mathit{or} \ q)$ comes out equivalent to $\neg p \wedge \neg q$, as desired.

Why would we systematically tend to interpret 'or' as '$\mathrm{or}_\top$' when 'or' appears in negative environments? This reading of disjunction *weakens* the interpretation of the disjunction in a positive context and hence *strengthens* it in a negative context. Across a wide range of theoretical contexts, including implicatures, homogeneity, reciprocity, and plurals, it has been suggested that, given a range of available readings, speakers tend to coordinate on the one 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. 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 set-up that brings out that intended reading. We have seen this in a number of cases, like the restrictor of 'every'. Here's another case that illustrates this nicely, building on Grice 1989:

(21) a. X or Y will be elected.
 b. That's not true, Y or Z will be elected.

A standard diagnosis of this case would be that the negation in the response is targeting the *assertability* of the first disjunction. But a simpler diagnosis, in light of everything we've seen here, is that the negation is straightforwardly targeting the asserted disjunction, in particular its possibility conjunct $\Diamond X$.

Finally, in 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, $\ulcorner p$ or $q\urcorner$ entails $p \vee q$, along with the possibility conjuncts $\Diamond p \wedge \Diamond q$. This is crucial for capturing the truth-conditions of disjunctions in both positive and negative contexts.

4.2 Other non-classical treatments of disjunction

There have been many other 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.

4.3 Embedded implicatures

A different path to explore is combining a *classical* theory of disjunction with some mechanism for embedding implicatures. Indeed, 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 affected 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 it derive the circumstantial reading of $\ulcorner p$ or $q\urcorner$. We could explore a similar move to the one briefly considered for 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, and might work if we adopt an innocent inclusion (Bar-Lev and Fox, 2020) or recursive (Fox, 2006) approach to exhaustification (since we need to generate $\Diamond p$ and $\Diamond q$ as alternatives to $\ulcorner p$ or $q\urcorner$ and then *add them* to the meaning). Still, I have a few worries. First, the resulting theory of *EXH* will be exceedingly complex, since this approach greatly increase 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. Finally, there remain 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 (which is standardly taken to be possible, if dispreferred)? 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. Conversely, my theory of disjunction can explain at least some of the data that motivate Blumberg and Goldstein, as in (22) (their (17)):

(22) [*Context: There are three detectives, and two possible suspects: Ann and Bill. One detective has already ruled out Ann, but the others haven’t ruled out either Ann or Bill.*]

   a. Exactly two detectives believe that Ann or Bill committed the crime.

My theory of ‘or’ straightforwardly predicts the true reading of (22), since exactly two detectives believe that Ann or Bill committed the crime and that both are possible.

4.4 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, translated into a static framework:

$$[\![p \text{ or } q]\!]^{c,w} = \begin{cases} 1 & [\![p \vee q]\!]^{c,w} = [\![\Diamond p]\!]^{c,w} = [\![\Diamond q]\!]^{c,w} = 1 \\ 0 & [\![p \vee q]\!]^{c,w} = 0 \\ \# & \textit{otherwise} \end{cases}$$

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 of 'or' trivially so that 'or' is interpreted as 'or$_\top$'. By contrast, it's not obvious how to do so on the present theory. And, conversely, it cannot capture non-classical interpretations 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 $\ulcorner B(p \text{ or } q)\urcorner$ means $p \vee q$. But this would be, to my knowledge, a new and otherwise unattested device; by contrast, the 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 gain us anything, and costs a fair amount in terms of complexity.

4.5 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 $\ulcorner p$ and $q\urcorner$ 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', according to which $\ulcorner p$ and $q\urcorner$ can indeed be 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 of the judgments I've elicited.

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 $\ulcorner p$ and $q\urcorner$ 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, to be sure, *also* a reading of this on which it is not assertable, 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 assertable one.

Second, an error theory needs to account for fine-grained *contrasts* in our data. Recall (14), repeated here:

(14) 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.

An error theory would say that the apparent coherence of (14) arises from a confusion: 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 (14) to the minimal variant in (23) 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:

(23) 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.

(23) seems strikingly less coherent than (14).

This contrast is immediately explained on my theory: (14) has an interpretation where the antecedent means $(A \vee B) \wedge \Diamond_c A \wedge \Diamond_c B$, and on that interpretation, the first sentence of (14) does not entail $\neg B$. By contrast, (23) has no such interpretation, and hence the first sentence of (23) entails $\neg B$, given that John will only buy one pie.

Perhaps an error theory could also explain the contrast here, but it is prima facie hard to see how: why would we find it difficult to see the entailment in the first case but easy in the second case, if the sentences have the same contextual meaning? This contrast suggests that the readings I have brought out arise *from the meaning of disjunction*.

In concluding, it is worth emphasizing that 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 possible, it is not obvious to me how that would go, or that the result would be simpler or more explanatory than the theory I've given. (Compare Williamson (2020)'s book-length defense of an error theory about conditionals; when he arrives at embeddings, he essentially throws up his hands. Accounting for systematic embedding data is exactly the place where a compositional semantic theory shows its strengths, and where, conversely, a globally-oriented error theory will struggle to make systematic predictions.)

It is always worth considering seriously 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 to systematically capture all of our observations.

5 Free choice

As I have noted, my theory builds closely on the modal list theory of disjunction, which is motivated by *free choice* inferences like these:

(24) a. You may have an apple or a pear.
b. So, you may have an apple.

(25) 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 $\ulcorner\Diamond(p \text{ or } q)\urcorner$ to $\ulcorner\Diamond p \text{ and } \Diamond q\urcorner$, and *wide scope free choice* is the inference from $\ulcorner\Diamond p \text{ or } \Diamond q\urcorner$ to $\ulcorner\Diamond p \text{ and } \Diamond q\urcorner$. 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).

The modal list approach was motivated by these inferences, and nicely accounts for them. On that theory, $\ulcorner\Diamond(p \text{ or } q)\urcorner$ 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. My impression is that this is due to the disastrous predictions this account makes about the interaction of disjunction with negation. As we have seen, by contrast, my theory, although it builds closely on the modal list approach, avoids its problems with negation by adding a classical conjunct to the meaning of 'or', so that $\ulcorner\neg(p \text{ or } q)\urcorner$ 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 modal list approach. (26-a), on my theory, is equivalent to (26-b):

(26) 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'. Given a classical meaning for $\Diamond$, (26-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, (27-a) has the meaning of (27-b), which entails $\Diamond_{or}\Diamond p \wedge \Diamond_{or}\Diamond q$.

(27) 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 (28) 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':

(28) 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 $\Diamond_{or}$. A similar move can account for the cancellability of free choice:

(29) You may have an apple or a banana, but I don't know which.

Here, a coherent reading can be obtained by overriding the default assumption that covert and overt modals are interpreted with the same flavor, so that the covert modals generated by 'or' in (29) are epistemic, and hence (29) has the coherent interpretation $(\Diamond_d A \vee \Diamond_d B) \wedge (\Diamond_e \Diamond_d A \wedge \Diamond_e \Diamond_d B) \wedge \neg K \Diamond_d A \wedge \neg K \Diamond_d B$. (A coherent reading could also be obtained by making the covert modals trivial.)

To capture free choice readings for deontic modals (as in (24)), 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 (8), repeated here, has a true reading:

(8) 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 (8) 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 Disjunction and epistemic possibility: Two more motivations

In concluding, I'll discuss two observations from the recent literature on disjunction and show that both provide support for my theory of disjunction.

The first is that disjunction gives rise to epistemic possibility inferences *even in the absence of ignorance inferences*. At the outset, I noted that Griceans have a standard story about how epistemic possibility inferences arise from *assertions* of disjunctions; my focus in §2 was on *embedded* disjunctions, where that story does not apply. But a striking objection to the Gricean derivation of epistemic possibility even for *asserted* disjunctions has recently been developed by Degano et al. (2023). Recall that the standard Gricean derivation of possibility implicatures goes via ignorance: when S asserts $\ulcorner p$ or $q \urcorner$, we infer that both p and q are epistemically possible for S 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 when S asserts $\ulcorner p$ or $q \urcorner$ *and it is common ground that S knows p*. In a representative experiment, subjects were shown a depiction of four boxes, three open and one concealed "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. Then subjects were asked to assess (30) along with a picture that depicts the three uncovered boxes each containing a yellow ball, with a blue ball in just one of them.

(30) The mystery box contains a yellow ball or a blue ball.

In this context, it is common ground that 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 disjunctive 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.

Schematically: when S asserts $\ulcorner p$ or $q \urcorner$, speakers conclude that q is epistemically possible for S even when it is common ground that S knows p, and hence even when the Gricean reasoning just reviewed is blocked. Asserted disjunctions license epistemic possibility inferences *even in the absence of ignorance*.

This poses a serious challenge to the standard Gricean derivation of possibility implicatures. By contrast, my theory of disjunction gives us a way to derive epistemic possibility inferences that does not go by way of uncertainty inferences.

The second observation from the recent literature comes from Feinmann (2023), who notes the following contrast:

(31) [*Context: It is common ground that Mary had either a cappuccino, an espresso, or a cup of tea, and that there is nothing else that she had.*]

A: Did Mary have a cappuccino, an espresso, or a cup of tea?

B: She had a cappuccino or an espresso.

A: ??OK, but was it a cappuccino that she had, or was it an espresso?

(32) [*Context: Same as in (31)*]

A: Did Mary have a cappuccino, an espresso, or a cup of tea?

B: She had a coffee.

A: ✓OK, but was it a cappuccino that she had, or was it an espresso?

As Feinmann argues, this contrast is surprising from the point of view of most theories of implicature, on which 'She had a cappuccino or an espresso' and 'She had a coffee' give rise to identical implicatures. By contrast, our theory has a ready explanation of this: if the disjunctive expression but not the indefinite lexicalizes possibility inferences, then A's response in (31) will be difficult to make sense of, since the speaker has just communicated that both disjuncts are possible for her.

7 Conclusion

I have argued that $\ulcorner p$ or $q\urcorner$ 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 contextual resolution of $\ulcorner p$ or $q\urcorner$ where it is 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, for epistemic possibility inferences without uncertainty, and for contrasts between disjunction and indefinites.

Other theories of the same data may be possible. 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 compelling candidate for the meaning of 'or'.

My goal has been to lay out and motivate this new theory in a relatively brief way. I leave open many interesting and important question. I will briefly survey some of them in concluding.

One is to what extent the embedded possibility phenomena I have described can be found in other kinds of constructions, like indefinite noun phrases; and, insofar as they can, to what degree my account can be extended. Thus, for instance, it has been claimed that indefinite noun phrases give rise to free choice inferences, so that 'You may have a stone fruit' naturally licenses the inference that you may have a plum and that you may have a peach, if those are salient stone fruits. It seems fairly straightforward to integrate the present approach into different theories of indefinites, but Feinmann's observation discussed above may suggest this is not the right path to pursue.

Another obvious question is how conjunction fits into my account: should we assume that 'and' means $\wedge$, thus giving up de Morgan's laws; or should we instead maintain de Morgan and treat $p \wedge q$ as equivalent to $\ulcorner\neg(\neg p \text{ or } \neg q)\urcorner$? Differences between these accounts will be hard to detect given the assumption that we tend to interpret 'or' as 'or$_\top$' in negative contexts, but it seems plausible to me that 'and' means $\wedge$, because I cannot see any evidence that $\ulcorner p$ and $q\urcorner$ is ever consistent with $\neg p$, whereas $\ulcorner\neg(\neg p \text{ or } \neg q)\urcorner$ is. Concretely, I think there is a coherent reading of (33-a) (modulo its complexity), but not of (33-b):

(33) a. It's not the case that the murderer isn't the butler *or* isn't the gardener; in fact, it isn't the butler.
b. The murderer is the butler and isn't the butler.

Another set of issues concerns how the possibility inferences generated by disjunction interact with other kinds of content, like presupposition, alternatives, and implicature (of the kind raised in Marty and Romoli 2022). I think cases like this speak in favor of the possibility of interpreting the context-sensitivity of disjunction differently in different dimensions of content: so, for instance, in 'Susie is not aware that John may have an apple or a pear', 'or's modals are interpreted deontically in the presupposed dimension and trivially in the main dimension, yielding the desired reading. Chatain and Schlenker 2024).

Another question concerns the phenomenon of *simplification of disjunctive antecedents*, which is superficially similar to free choice, but which is not explained by my theory. The latter may have a different explanation, however, of the sort recently explored in Klinedinst 2024.

Another question concerns the relative availability of readings that I have brought out, and striking interpersonal variation in this respect. This is an area where detailed experimental work would be helpful.

A final question concerns potential connections between the phenomena I have studied here and a variety of other phenomena studied in the philosophical literature, like issues about confirmation and updating discussed by Yablo (2022); Krämer (2022).

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.

Chatain, K. and Schlenker, P. (2024). Local pragmatics redux: Presupposition accommodation without covert operators. Manuscript, ENS.

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 Heusinger, 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.

Feinmann, D. (2023). Ignorance is a problem. *Snippets.*

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, H. P. (1975). *Logic and Conversation.* Harvard.

Grice, P. (1989). Indicative conditionals. In *Studies in the way of words*, chapter 4. Harvard University Press.

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.illc.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).

Krämer, S. (2022). Mighty belief revision. *Journal of Philosophical Logic*, 51:1175–1213.

Kratzer, A. and Shimoyama, J. (2002). Indeterminate pronouns: The view from Japanese. 3rd Tokyo Conference on Psycholinguistics.

Marty, P. and Romoli, J. (2022). Presupposed free choice and the theory of scalar implicatures. *Linguistics and Philosophy*, 45:91–152.

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.

Williamson, T. (2020). *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford.

Yablo, S. (2022). Tangled up in grue. Handout for a talk in Geneva.

Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8(4):255–290.