Independent Alternatives: Ross's Puzzle and Free Choice

Forthcoming in *Philosophical Studies*.

Richard Booth

Abstract

Orthodox semantics for natural language modals give rise to two puzzles for their interactions with disjunction: Ross's puzzle and the puzzle of free choice permission. It is widely assumed that each puzzle can be explained in terms of the licensing of 'Diversity' inferences: from the truth of a possibility or necessity modal with an embedded disjunction, hearers infer that each disjunct is compatible with the relevant set of worlds. I argue that Diversity inferences are too weak to explain the full range of data. Instead, I argue, modals with embedded disjunctions license 'Independence' inferences: from the truth of a modal with an embedded disjunction, hearers infer that each disjunct is an independent alternative among the relevant set of worlds. I then develop a bilateral inquisitive semantics for modals that predicts the validity of these Independence inferences. My account vindicates common intuitions about both Ross's puzzle and the puzzle of free choice permission, and explains the full range of data.

1 Introduction

Orthodox semantic theories of natural language modals make the seemingly correct prediction that we can make true, underspecific statements about what is necessary or possible.[1] For example, if it is necessary that I mail the letter *overnight*, orthodox theories predict that my utterance of 'I ought to mail the letter' is true, even though this sentence is not fully specific about what I have to do. When combined with the standard theory of disjunction, however, this apparently correct prediction gives rise to two well-known puzzles.[2] First, there is *Ross's Puzzle* (Ross, 1941), which is the puzzle of reconciling the fact that orthodox

[1]By 'orthodox', I mean any theory that makes necessity and possibility modals upward monotonic, e.g. the standard Kripke semantics that uses accessibility relations, and the context-sensitive semantics of Kratzer (2012b). 'Underspecific' has many different uses in the philosophy of language; the one I employ here has precedent in Zimmermann (2006) and Fara (2013).

[2]By 'standard theory of disjunction', I mean especially theories like that of Partee and Rooth (1983), where disjunction is treated as the Boolean dual of conjunction.

semantic theories predict the following inference to be valid with the appearance that it is *invalid*:[3]

(1) a. Alicia ought to mail the letter. $\Box M$
b. So, Alicia ought to either mail the letter or burn it. $\Box(M \vee B)$

I will call an inference that exhibits the pattern in (1) a *Ross inference*.

Second, there is the *Puzzle of Free Choice Permission* (von Wright, 1968; Kamp, 1973). Here, the challenge is to reconcile the apparent *validity* of the following inference with the fact that orthodox theories predict it to be invalid:

(2) a. Alicia may either mail the letter or burn it. $\Diamond(M \vee B)$
b. So, Alicia may burn the letter. $\Diamond B$

I will call an inference that exhibits the pattern in (2) a *free choice inference*.

There have been two kinds of solution to the puzzles. Pragmatic accounts defend the semantic predictions of the orthodox theories and give a pragmatic explanation of the apparent (in)validity in each case. Semantic accounts, by contrast, reject orthodoxy in favor of revisionary semantics for modals and/or disjunction in order to vindicate the appearance of (in)validity.

Both kinds of solution must explain *why* speakers' intuitions about the (in)validity of these arguments are in opposition to the predictions of the orthodox semantics. Since a speaker's judgment of the (in)validity of an argument depends on what meanings the speaker assigns to the premises and the conclusion, providing such an explanation involves taking a stand on what speakers intuitively take these sentences to mean. In fact, both semantic and pragmatic solutions to the puzzles have typically agreed on this point: over and above their orthodox truth conditions, sentences of the form $\Diamond(p \vee q)$ and $\Box(p \vee q)$ convey that p and q are each compatible with the set of relevant worlds. In this paper, I argue that this analysis cannot explain the full range of data. Instead, I argue that claims of the form $\Diamond(p \vee q)$ and $\Box(p \vee q)$ convey something stronger: that p and q are each *independent* alternatives among the set of relevant worlds. I then develop an original revisionary semantics that validates these inferences, and thereby vindicates our unorthodox intuitions about both puzzles.

The plan is as follows. In the next section (§2), I explain why orthodox semantic theories make the counter-intuitive predictions that they do in our two

[3]In keeping with much of the recent literature, I use declarative 'ought' claims to illustrate Ross's puzzle. Ross's original example was presented using imperatives: where 'Post the letter!' entails 'Post the letter or burn it!' Although I do not discuss imperatives in this paper, my account can be straightforwardly extended to them.

puzzles, and in §3, I survey a range of different modals for which the two puzzles arise. In §4, I argue that the puzzles need to be explained in terms of what I call *Independence inferences*. In §5, I show that these inferences are validated when a simple topological relationship — a minimal covering relation — holds between the truth sets of the disjuncts and the set of relevant worlds over which a modal quantifies. In §6, I use this relation to outline an inquisitive semantics for modals that validates the Independence inferences, and I show how it vindicates our intuitions in the two puzzles. Then, in §7, I extend the model to a 'bilateral' version in order to generate better predictions for the interactions between modals and negation. Finally, in §8, I show that the bilateral theory generates some truth value gaps, and argue that these well-placed gaps shed new light on the original puzzles.

2 Orthodox Predictions

Orthodox semantic theories treat necessity modals like 'ought' and possibility modals like 'may' as universal and existential quantifiers, respectively, over sets of worlds.[4] Where $R(w)$ is the set of worlds relevant at the world w,[5] the orthodox theory of 'ought' ($\Box$) and 'may' ($\Diamond$) roughly amounts to the following, with ϕ^w giving the truth value of ϕ at a world w (for simplicity, I ignore some syntactic complexity, the dependence of R on context, and semantic interactions

[4]See Kripke (1963). In Kratzer (1977, 1991), the dominant orthodox theory of natural language modals, there are several worthwhile complications of the basic quantificational idea, but they do not matter for our purposes. See Portner (2009) for a textbook presentation of various semantic theories of natural language modals.

[5] In the Kratzerian dialect (Kratzer, 1977, 1991), where modals are relative to a *modal base* f and an *ordering source* g, we define $R(w)$ as follows:

$$R(w) = \mathsf{max}_{g(w)}(\bigwedge f(w))$$

where $\mathsf{max}_{g(w)}$ is a function that takes a proposition and returns the subset of worlds that are maximal with respect to the order determined by g.

In the Kripkean dialect, where modals are sensitive to an *accessibility relation* R between worlds in a set W, we define $R(w)$ as follows:

$$R(w) = \{v \in \mathsf{W} \mid (w, v) \in \mathsf{R}\}$$

with tense):

$$\Box\phi^{w} = \text{TRUE} \text{ iff for every } w' \in R(w) : \phi^{w'} = \text{TRUE}$$
$$\Diamond\phi^{w} = \text{TRUE} \text{ iff for some } w' \in R(w) : \phi^{w'} = \text{TRUE}$$

That is, $\Box(\phi)$ is true iff ϕ is true at every relevant world, while $\Diamond(\phi)$ is true iff ϕ is true at at least one relevant world. As a result, orthodox semantic theories make modals like 'ought' and 'may' *upward monotonic* operators:

> For any sentential operator Δ: Δ is upward monotonic iff whenever ϕ entails ψ, $\Delta(\phi)$ entails $\Delta(\psi)$.

The upward monotonicity property means that in order to speak truly about what is necessary or possible, one need not be completely specific. For example, consider 'ought'. By the orthodox semantics above, 'Alicia ought to mail the letter overnight' is true at w iff every world in $R(w)$ is one where 'Alicia mails the letter' is true. Since 'Alicia mails the letter overnight' entails 'Alicia mails the letter', whenever that latter condition holds, every world in $R(w)$ will also be one where 'Alicia mails the letter' is true. Thus, the less specific sentence 'Alicia ought to mail the letter' is true whenever the more specific sentence 'Alicia ought to mail the letter overnight' is.

By the same token, however, orthodox theories make the counter-intuitive predictions they do in our two puzzles. Both are a direct result of the fact that a disjunction $\phi \vee \psi$ is asymmetrically entailed by its disjuncts, ϕ and ψ. This means that in general, $\Box\phi$ asymmetrically entails $\Box(\phi \vee \psi)$, and $\Diamond\phi$ asymmetrically entails $\Diamond(\phi \vee \psi)$:

$$\Box\phi \vDash \Box(\phi \vee \psi)$$
$$\Diamond\phi \vDash \Diamond(\phi \vee \psi)$$

The Ross inference in (1) is just an instance of this schema. Since 'Alicia mails the letter' entails 'Alicia either mails the letter or burns it', orthodox theories predict that 'Alicia ought to mail the letter' entails 'Alicia ought to either mail the letter or burn it' ($\Box M \vDash \Box(M \vee B)$).

The invalidity of the free choice inference has the same source. Suppose there is only one relevant world, w. Further, suppose M is true and B is false at w. Then $\Diamond M$ is true, since w is a relevant world where M is true. Since $\Diamond$ is an upward monotonic operator, it follows that $\Diamond(M \vee B)$ is automatically true. However, since every relevant world (i.e., w) makes B false, we also have that $\Diamond B$ is false.

This presents a counterexample to the validity of the argument in (2): the premise $\Diamond(M \vee B)$ is true while the conclusion, $\Diamond B$ is false. So the free choice inference is invalid.

3 Range

Before moving on to my analysis of the two puzzles, I want to briefly note their generality. While they have sometimes been treated as arising only for a special class of *deontic* modals like 'ought' and 'may',[6] recent research has made clear that they arise for a much wider class of modals,[7] and are not tied to any particular 'flavors' of modality.[8] To see this, let us schematize the inferences as follows, where Δ is a modal operator:[9]

$$\begin{array}{rl} \text{(Ross schema)} & \Delta(p) \text{ therefore } \Delta(p \vee q) \\ \text{(Free choice schema)} & \Delta(p \vee q) \text{ therefore } \Delta p \text{ and } \Delta q \end{array}$$

Consider the following instances of the Ross schema with non-deontic modals:

(3) a. An object in motion $\left\{\begin{array}{c}\text{must}\\ \text{will}\\ \text{has to}\\ \text{is likely to}\\ \text{might}\\ \text{can}\end{array}\right.$ stay in motion.

[6] The puzzles were discovered in early work on deontic logic and the logic of imperatives (von Wright, 1968; Kamp, 1973; Ross, 1941). Some recent research has continued this focus on deontic flavors of modality, including Barker (2010), Cariani (2013), Fusco (2015), and Starr (2016). Of course, as some of these authors mention, there are likely ways to extend these accounts, tailored to the deontic case, to other flavors of modality.

[7] See Zimmermann (2006), Yablo (2014), Abreu Zavaleta (2019) for some examples of invalid, non-deontic Ross inferences, and Zimmermann (2000); Nickel (2010); Romoli and Santorio (2017); Willer (2021) for some discussion of valid, non-deontic free choice inferences.

[8] It is standardly assumed that modal auxiliaries like 'must' are context sensitive, and in various contexts can express different *flavors* of necessity, such as metaphysical, epistemic, deontic, and so on.

[9] Throughout, by 'modal' or 'modal operator', I mean a sentential operator that is standardly analyzed as shifting the world of evaluation for its complement proposition. 'Modal' as I am using it thus includes not just modal auxiliaries like 'must' and 'may', but also attitude verbs like 'want' and 'believe'.

b. # So, an object in motion {must / will / has to / is likely to / might / can} either stay in motion or come to absolute rest.

The subject matter in these examples encourages physical, metaphysical, or epistemic interpretations of the modals, and yet the inference appears just as invalid as the one in (1). Similarly, the pattern appears invalid for many attitude verbs, which are usually given a necessity modal semantics. For some examples:

(4) a. Alicia {seeks / intends / hopes / wants / expects} to mail the letter.

b. # So, Alicia {seeks / intends / hopes / wants / expects} to either mail the letter or burn it.

(5) a. Alicia {said / claims / believes / thinks} that Bulmaro mailed the letter.

b. # So, Alicia {said / claims / believes / thinks} that Bulmaro either mailed the letter or burned it.

Each seems to be a complete non sequitur. See Zimmermann (2006), Yablo (2014), Abreu Zavaleta (2019) for further discussions of Ross inferences with some of these verbs.

Similarly, as Zimmermann (2000); Nickel (2010); Romoli and Santorio (2017); Willer (2021) argue, the free choice pattern appears to be valid for non-deontic possibility modals. For some examples:

(6) a. Given their skill, Alicia or Bulmaro $\left\{\begin{array}{l}\text{might}\\ \text{could}\end{array}\right.$ have won the game.

b. So, Alicia $\left\{\begin{array}{l}\text{might}\\ \text{could}\end{array}\right.$ have won and Bulmaro $\left\{\begin{array}{l}\text{might}\\ \text{could}\end{array}\right.$ have won.

(7) a. Hydrangeas can grow either pink or blue flowers.

b. So, hydrangeas can grow pink flowers and they can grow blue flowers.

Each inference seems as well-supported as the original in (2), despite the fact that the modals take on a non-deontic interpretation.

This data suggests that the intuitive status of the two inferences has nothing to do with a deontic interpretation of the modals involved.[10] An adequate account of the puzzles, therefore, should be general enough to accommodate modals of any flavor.

4 Independence

As mentioned in the introduction, a speaker's judgment of the (in)validity of an argument depends on what meanings the speaker takes the premises and conclusion to have. In judging the Ross inference to be invalid, a speaker senses that $\Box(M \vee B)$, for example, intuitively means something that the truth of $\Box M$ fails to guarantee. And in judging a free choice inference to be valid, a speaker senses that $\Diamond(M \vee B)$ intuitively means something that should guarantee the truth of both $\Diamond M$ and $\Diamond B$. Thus, a crucial step toward providing a solution to the two puzzles is to identify what intuitive meanings speakers take sentences of the form $\Delta(p \vee q)$ to have (I will use Δ as an variable ranging over possibility and necessity modals that give rise to either of the puzzles), such that orthodox theories fail to correctly predict that these sentences have this intuitive meaning.

Both pragmatic and semantic accounts of the puzzles have typically agreed on this issue.[11] They say that over and above its orthodox truth conditions, a modal claim of the form $\Delta(p \vee q)$ conveys that each disjunct is *compatible* with the relevant set of worlds. That is, they say that the information that $\Delta(p \vee q)$

[10] This is not to claim that there are no modals for which the Ross inference is valid — perhaps there are. My claim is a weaker one: given these examples, the solution we propose to the puzzles should be flexible enough to handle any flavor of modality.

[11] See, for example, von Fintel (2012) for a pragmatic version of this thesis and Simons (2005) for a semantic one. As I mention below, Menéndez Benito (2005, 2010) has dissented in the analogous case of free choice 'any' under possibility modals, and building on this work, Aloni and Ciardelli (2013) dissent in the analogous case of imperatives.

intuitively conveys, and which orthodox theories fail to predict, is underwritten by what I will call *Diversity* inferences (where $\Rightarrow$ is ambiguous between semantic entailment, implicature, or some other robust form of licensing):

Diversity Inferences:

$$\begin{aligned}\Box(p \vee q) &\Rightarrow \Diamond p\\ &\Rightarrow \Diamond q\\ \Diamond(p \vee q) &\Rightarrow \Diamond p\\ &\Rightarrow \Diamond q\end{aligned}$$

For modals without lexicalized duals, like 'want' and 'intend', we can give a meta-linguistic characterization of the Diversity inferences as follows:

Meta-Linguistic Diversity Inferences:

Where Δ is a necessity or possibility modal, $R(w)$ is the set of relevant worlds associated with the modal Δ at w, and p is the set of worlds in which p is true:

$$\begin{aligned}\ulcorner\Delta(p \vee q)\urcorner \text{ is true at } w &\Rightarrow R(w) \cap p \neq \varnothing\\ &\Rightarrow R(w) \cap q \neq \varnothing\end{aligned}$$

The Diversity inferences, not supported by the orthodox semantics, promise to explain our intuitive judgments in both puzzles. For free choice, this is trivial: the inference just is a special case of Diversity. For the Ross inference, the thought is usually that because of the Diversity inferences, the conclusion $\Box(M \vee B)$ conveys $\Diamond B$, and since the premise $\Box M$ does not guarantee this, the argument is judged to be invalid.[12] I will call the theory that the Diversity inferences capture the unorthodox content conveyed by modals with disjunctive complements, and that these inferences explain the gap between the orthodox semantics and our intuitions about the two puzzles, the 'Diversity analysis'.

Unfortunately, the Diversity analysis cannot explain all of the relevant data. In particular, there are two pieces of data it cannot explain: what I will call *extended Ross's puzzle*; and what I will call *independence conditional inferences*. First, let me introduce the former. The Diversity analysis says that the argument in (1) is (or merely appears) invalid because the conclusion entails (or implicates) the

[12] See, for example, Wedgwood (2006); von Fintel (2012).

permissibility of the added disjunct — a fact not guaranteed by the truth of the premise. If that analysis were correct, then adding a premise that explicitly guarantees the permissibility of the added disjunct should change whether or not the argument is judged to be invalid. However, this prediction is not borne out. Even in cases where both disjuncts are guaranteed to be permissible by the premises, the inference pattern seems invalid (see Sayre-McCord (1986) and Fusco (2015) for this point). For example:

Extended Ross's Puzzle.

(8) a. Alicia ought to mail the letter. $\Box M$
b. Alicia may use the phone. $\Diamond P$
c. # So, Alicia ought to either mail the letter or use the phone. $\Box(M \vee P)$

The conclusion in (8c) does not seem to follow. But the premises guarantee that the Diversity inferences are supported.[13] So the defectiveness of the inference in (8) cannot be a result of the failure of the Diversity inferences to be licensed by the premises. Notice that like the original Ross inference, the extended pattern is also defective for non-deontic modals. For example, consider the case of epistemic 'must' and 'might':

(9) a. Alicia must have mailed the letter. $\Box M$
b. Alicia might have used the phone. $\Diamond P$
c. # So, Alicia must have either mailed the letter or used the phone. $\Box(M \vee P)$

The second type of data that the Diversity analysis cannot explain arises from the interaction between modals and conditionals.[14] As is well known, conditionals of the form 'If ϕ, then $\Delta(\psi)$' (where Δ is a modal) often give rise to 'restriction' readings, where the semantic function of the antecedent clause (ϕ) seems to be to restrict the domain of worlds relevant for the modal Δ in the consequent to the subset of relevant worlds where the antecedent is true ($R(w) \cap \phi$).[15] When

[13] I am making the standard assumption that modals, like other quantifiers, presuppose that their domains are non-empty. Thus, $\Box M$ entails that there is a relevant world where M is true. This means that whenever $\Box M$ is true, $\Diamond M$ is also true.

[14] See also Fusco (2015) for this data in the case of 'ought', and an alternative account of it.

[15] See Kratzer (2012a,b) for the classic theory of this phenomenon in the case of non-attitude modals, and Blumberg and Holguín (2019) for recent work on restriction effects in the case of attitude verbs.

modals capable of such restriction have disjunctive complements, their acceptance appears to license the inference of what I call *independence conditionals*:

Independence Conditionals.

(10) a. Alicia ought to either write an essay or give a presentation. $\Box(E \vee P)$
b. So, if she doesn't write an essay, she ought give a presentation. $\neg E \rightarrow \Box P$
c. So, if she doesn't give a presentation, she ought to write an essay. $\neg P \rightarrow \Box E$

(11) a. Alicia may either write an essay or give a presentation. $\Diamond(E \vee P)$
b. So, if she doesn't write an essay, she may give a presentation. $\neg E \rightarrow \Diamond P$
c. So, if she doesn't give a presentation, she may write an essay.$\neg P \rightarrow \Diamond E$

Given the standard assumption that a necessity modal like 'ought' presupposes that its domain is non-empty, the restriction readings of these conditionals have the following truth conditions:

(10b) True iff there are some relevant worlds where Alicia doesn't write an essay, and in all of these, she gives a presentation ($R(w) \cap \neg E \neq \varnothing$ and $R(w) \cap \neg E \subseteq P$).

(10c) True iff there are some relevant worlds where Alicia doesn't give a presentation, and in all of these, she writes an essay ($R(w) \cap \neg E \neq \varnothing$ and $R(w) \cap \neg E \subseteq P$).

(11b) True iff in some relevant world where Alicia does not write an essay, she gives a presentation ($(R(w) \cap \neg E) \cap P \neq \varnothing$).

(11c) True iff in some relevant world where Alicia does not give a presentation, she writes an essay ($(R(w) \cap \neg P) \cap E \neq \varnothing$).

Neither orthodox semantic theories nor the Diversity analysis predicts that (10b-10c) should follow from (10a) or that (11b-11c) should follow from (11a). To see this, suppose that (10a) and (11a) are true. Orthodox semantic theories allow that (10a) and (11a) can be true if there are no relevant B-worlds, or if there are no relevant M-worlds. In the former case, the conditionals (10b) and (11b) would be false, while in the latter, (10c) and (11c) would be false. Thus, orthodox semantic theories do not validate these conditional inferences. But the Diversity inferences cannot explain their plausibility either. The truth conditions of (10b-10c) and (11b-11c) require not just that there is some relevant E-world and some relevant P-world. Rather, (10b) and (11b) each require more specifically that there is a relevant P-and-not-E-world, and (10c) and (11c) each require that there is a relevant E-and-not-P-world.[16] In other words, these conditionals require that E and P are *independently* relevant alternatives, and, I claim, the acceptance of (10a) and (11a) appears to guarantee this.

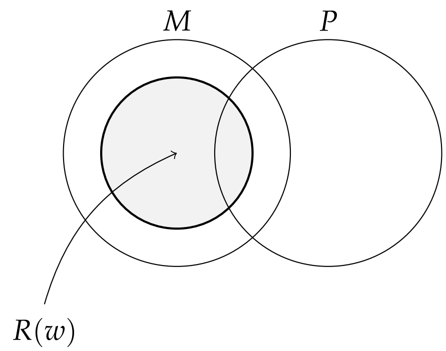

Figure 1: Diversity without Independence: $R(w) \cap M \neq \varnothing$ and $R(w) \cap P \neq \varnothing$

In fact, this requirement of *independent* relevance also offers a plausible explanation of what goes wrong in the extended Ross's puzzle above. The conclusion (8c) does not merely convey that using the phone is allowed. It conveys that Alicia may fulfill her obligation(s) by using the phone *independently* of whether she mails the letter. The premises of (8) do not guarantee this. If (9) is true at w, then there is some world v in $R(w)$ where Alicia uses the phone. But given (8a), it follows that Alicia *also* mails the letter at v. This situation is illustrated in Figure 1, where the relevant set of worlds, $R(w)$, makes the premises (8a) and (9) true. The problem, according to me, is that while using the phone is permissible, it is not an independent option on a par with mailing the letter. In fact, the premises of (8) entail that using the phone is precisely *not* an independent way for Alicia to fulfill her obligation(s), for if they are true, Alicia may use the phone only if she *also* mails the letter. For using the phone to constitute an independent way for Alicia's obligation(s) to be fulfilled, as I am claiming the conclusion in (8c) requires, there must also be a relevant world where she uses the phone without mailing the letter. That is just to say the conclusion (8c) requires both $\Diamond(M \wedge \neg P)$ and $\Diamond(P \wedge \neg M)$ to be true. In other words, the set of relevant worlds has to intersect the relative complements of the disjuncts, as illustrated in Figure 2.

[16] Again, in the case of (10b) and (10c), this assumes that natural language necessity modals like 'ought' presuppose that their domains are non-empty.

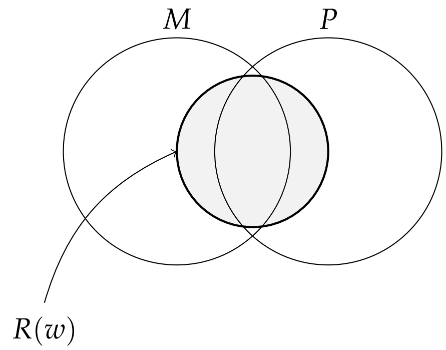

Figure 2: Independence: $R(w) \cap M \wedge \neg P \neq \emptyset$; $R(w) \cap P \wedge \neg M \neq \emptyset$

The Independence requirement I am proposing has some precedent in Menéndez Benito (2010), which argues that a similar condition governs the interaction between possibility modals and free choice 'any' (and Spanish 'cualquiera').[17] Suppose we are playing a card game and before you is the whole suit of diamonds on one side, and the whole suit of hearts on the other. If the rules of the game require that you now take one of these suits, then saying 'You may take any diamond card' is deeply misleading, as Menéndez Benito argues. For while it is true that for each diamond card, there is a permissible possibility in which you take it, the sentence seems to convey the stronger permission to take each diamond card *independently* of taking the others.[18] We may adapt the example to show something similar for the interaction between possibility modals and *disjunctive* complements. Suppose now that you have only one option: to take the pair consisting of the ace and ten of diamonds. Orthodox semantic treatments of 'may' would say that the following sentence is true:

(12) # You may either take the ace or the ten of diamonds.

But (12) seems to misdescribe things. Orthodox semantic theories, by themselves, do not have the resources to explain why, since taking the pair is a relevant possibility, and doing so makes the embedded disjunction true. But the Diversity

[17] I will not discuss free choice 'any' in this paper, but I think it will be clear enough how my account could be extended to that case in order to capture the data Menéndez Benito puts forward.

See also Aloni and Ciardelli (2013), which applies Menéndez Benito's insight to the case of imperatives.

[18] The term used by Menéndez Benito (2005, 2010)) is not *independence* but *exclusivity*. I have opted for 'independence' here since I do not want to suggest any erroneous connections to 'exclusive' disjunction.

analysis cannot explain why (12) is misleading either, since taking the pair involves taking the ace and taking the ten, making each disjunct true. By contrast, an extension of our analysis of (8) does offer an explanation for why (12) is misleading: the disjuncts do not pick out *independent* alternatives. Since the only relevant option is to take the whole pair, you cannot take the ace of diamonds *without* taking the ten of diamonds, or vice versa. On my view, then, (12) is infelicitous because the alternatives it describes are not independently possible.[19]

In sum, I argue that both possibility and necessity modals license what I will call *Independence* inferences. Accepting $\Delta(p \vee q)$ disposes us to accept that p and q are each independent alternatives in the relevant domain of possibilities: i.e., that both (i) p-without-q, and (ii) q-without-p are compatible with the relevant set of worlds. For necessity and possibility modals with lexicalized duals, we can schematize these inferences as follows:

Independence Inferences:

$$\begin{aligned} \Box(p \vee q) &\Rightarrow \Diamond(p \wedge \neg q) \\ &\Rightarrow \Diamond(q \wedge \neg p) \\ \Diamond(p \vee q) &\Rightarrow \Diamond(p \wedge \neg q) \\ &\Rightarrow \Diamond(q \wedge \neg p) \end{aligned}$$

For modals without lexicalized duals like 'want', or 'intend', we can give a meta-linguistic characterization of Independence as follows:

Meta-Linguistic Independence Inferences:

Where Δ is a necessity or possibility modal, $R(w)$ is the set of relevant worlds associated with the modal Δ at w, and p is the set of worlds in which p is true:

$$\begin{aligned} \ulcorner\Delta(p \vee q)\urcorner \text{ is true at } w &\Rightarrow R(w) \cap (p \backslash q) \neq \varnothing \\ &\Rightarrow R(w) \cap (q \backslash p) \neq \varnothing \end{aligned}$$

Note that, since $\Diamond(p \wedge \neg q)$ entails $\Diamond p$, the Independence inferences are strictly stronger than the Diversity inferences.

Of course, the Independence inferences only make sense if the embedded disjunction $p \vee q$ satisfies 'Hurford's constraint' against redundant disjuncts —

[19] In fact, on the theory I go on to develop in this paper, (12) will be neither true nor false in the described situation. I discuss the status of (12) further in §8.

that is, if neither p nor q entails the other (see Hurford (1974) and Gazdar (1979)). I will assume, following recent work on the topic, that felicitous disjunctions which appear to flout Hurford's constraint at the level of surface grammar are interpreted at the level of logical form via 'exhaustification' operators that enforce conformity to the constraint.[20] This means that for the interpreted structures we are interested in, we may safely assume Hurford's constraint holds and neither *p-and-not-q* nor *q-and-not-p* are contradictions.

In the rest of this paper, I will provide a positive account of the interactions between modals and disjunctions that predicts the Independence inferences. Before moving on to that project, I want to note that it begins with an important theoretical choice: whether to explain the Independence inferences pragmatically or semantically. Presenting a conclusive argument in favor of one or the other strategy would not be possible here. Instead, I will limit myself to noting some problems that pragmatic accounts face. I think these problems suggest that we should opt for a semantic account. First, if hearers make the Independence inferences pragmatically (as Wedgwood (2006) and von Fintel (2012) argue for the Diversity inferences), then it is unclear why Ross inferences should seem so persistently *invalid*. Pragmatic inferences in general are usually drawn when they seem plausible. But a pragmatic account of the Independence inferences hypothesizes that when encountering the argument in (1), a hearer draws the Independence inferences precisely when they are *implausible*, given the premise. Indeed, on a pragmatic account, the hearer is supposed to draw the Independence inferences, sense their implausibility, and rather than withdraw them, mistake a semantically valid argument for an invalid one.[21]

[20]For example, take:

(13) Alicia must have either some or all of the ice cream.

The surface grammar of the embedded disjunction, 'Alicia had either some or all of the ice cream,' flouts Hurford's constraint, since having *all* of the ice cream entails having *some* of it. This means that the modal claim (13) could not possibly license both of the Independence inferences, since one of them would be:

(14) Alicia may have all but not some of the ice cream.

(14) cannot be true, since the embedded conjunction, 'Alicia has all but not some of the ice cream,' is a contradiction. The assumption I will make is that in recognizing this, a hearer assigns to (13) a logical form roughly equivalent to, 'Alicia must have either *merely* some, or all of the ice cream" (where the exhaustification operator *merely* has the semantic function of denying that Alicia has all of the ice cream. See Simons (2001), Katzir and Singh (2013), Meyer (2013, 2014), and Ciardelli and Roelofsen (2017) for recent discussion.

[21]An anonymous reviewer suggests that the independence inferences might be computed by

Furthermore, recent empirical research on pragmatic approaches to the free choice inferences suggests they may have some trouble explaining data concerning processing times (Chemla and Bott, 2014) and the interaction between these inferences, on the one hand, and non-monotonic contexts or presuppositions, on the other (see Romoli and Santorio (2019); Gotzner et al. (2020)). A semantic account of the Independence inferences has the potential to do better on all of these points. For these reasons, the solution to our puzzles that I outline below will derive the validity of the Independence inferences as part of the semantics of sentences of the form $\Delta(p \vee q)$.

5 Independence and Minimal Covers

Now that we have seen that the Independence inferences are needed to explain the puzzles, what consequence does this have for the semantics of modals with embedded disjunctions? First, let us consider necessity modals. Validating the Independence inferences means that the truth conditions of $\Box(p \vee q)$ will involve two requirements. First, there is the one that orthodox accounts predict: that $\Box(p \vee q)$ is true only if $p \vee q$ is true in all of the relevant worlds. Second, p and q must pick out independent alternatives among the relevant set of worlds — that is, the relative complements of the truth sets of the disjuncts (the set of p-without-q worlds, the set of q-without-p worlds) must each have a non-empty intersection with the set of relevant worlds. Together, these two requirements mean that in order for $\Box(p \vee q)$ to be true, the truth sets of the disjuncts must form a *minimal cover* of the set of relevant worlds (as illustrated by Figure 2 above).[22] A minimal

the combination of the Diversity inferences together and an *exclusive* reading of the embedded disjunction. On this hypothesis, the Independence inferences would only be licensed in cases where it is reasonable to suppose that there is no relevant world that makes the *conjunction* of the disjuncts true. While I cannot provide a conclusive argument against this hypothesis here, I think that the apparent invalidity of examples like (8) and (9) provide evidence against it. In those cases, the premises can be true together, and when they are both true, there is a relevant world that the conjunction of the disjuncts true — where Alicia mails the letter and uses the phone. In such a case, it should be natural for speakers to opt for an *inclusive* reading of the disjunction embedded in the conclusion, and recognize that it follows on the orthodox semantics (even supplemented with the Diversity inferences). But both inferences seem just as invalid as the original (1). This suggests to me that the Independence inferences do not arise only on an exclusive reading of the disjunction involved, and thus that the Independence inferences are not computed in the way this hypothesis claims.

[22]Simons (2005) also uses a *covering* relation as a helpful way of summarizing the modal/disjunction interaction. For Simons, modals with disjunctive complements truth conditionally require that the disjuncts form a *supercover* of the relevant set of worlds. C is a supercover

cover C of a set S is a collection of sets such that (i) their union contains S (they *cover S*); and (ii) no proper subset of C is such that its union contains S (they do so *minimally*):

$$\begin{aligned} C \text{ is a } \textbf{cover} \text{ of } S \text{ iff } & \; S \subseteq \bigcup C \\ C \text{ is a } \textbf{minimal cover} \text{ of } S \text{ iff } & \; C \text{ is a cover of } S \text{ and} \\ & \; \text{there is no } C' \subset C : C' \text{ is a cover of } S \end{aligned}$$

To break this down: the first requirement associated with the truth of $\Box(p \vee q)$ — that $p \vee q$ be true throughout the relevant set of worlds — is equivalent to the requirement that the truth sets of the disjuncts ($\{p, q\}$) form a *cover* of the relevant set of worlds. The second requirement — that among the relevant worlds, there are both p-without-q-worlds and q-without-p-worlds — is equivalent to the requirement that the cover be a *minimal* one: neither the truth set of p nor that of q suffices on its own to cover the set of relevant worlds.

We can also use the notion of a minimal cover to explain what validating the Independence inferences will mean for the truth conditions for sentences of the form $\Diamond(p \vee q)$. First, there is the part that orthodox semantics predict: since $\Diamond(p \vee q)$ is a possibility claim, its truth requires that there is at least one relevant world where $p \vee q$ is true. Equivalently, the truth sets of the disjuncts must form a cover of a *non-empty* subset of relevant worlds. Second, to support the Independence inferences, the relative complements of the contents of the disjuncts (the set of p-without-q worlds, the set of q-without-p worlds) must each have a non-empty intersection with the set of relevant worlds $R(w)$. Together, these requirements mean that the truth sets of the disjuncts must form a minimal cover of some non-empty subset of relevant worlds. This is illustrated in Figure 3, where $\{p, q\}$ forms a minimal cover of the shaded subset of relevant worlds, R'.

Using the notion of a minimal cover, our target truth conditions for sentences

of S iff it is a cover of S and every member of C has a non-empty intersection with S. Note that every minimal cover is a super cover but not vice versa; and that a supercover semantics validates the Diversity, but not Independence, inferences. See also Nygren (2019), which systematically explores the logic of a supercover semantics.

At the end of her paper (§6), Simons briefly considers various pragmatic 'add-on' requirements that she thinks may govern the felicity of disjunctions in certain contexts. One of the three requirements she outlines resembles the minimal covering relation I define here. However, she does not explore this pragmatic constraint in much detail, and clearly does not think, as I argue here, that it is part of the literal, truth conditional semantics of modals with embedded disjunctions.

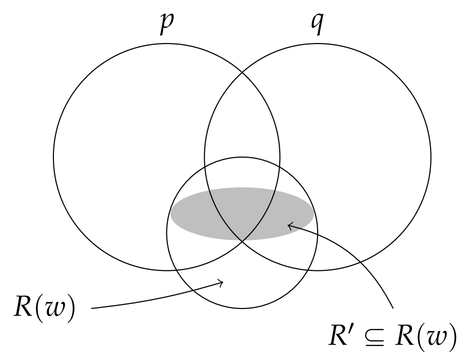

Figure 3: Independence for possibility modals: the shaded region R' contains both p-without-q worlds, and q-without-p worlds.

of the form $\Box(p \vee q)$ and $\Diamond(p \vee q)$ are as follows:

$$\Box(p \vee q) \text{ is true at } w \text{ iff } \{p, q\} \text{ is a minimal cover of } R(w)$$
$$\Diamond(p \vee q) \text{ is true at } w \text{ iff there is a non-empty } R' \subseteq R(w) \text{ such that } \{p, q\} \text{ is a minimal cover of } R'$$

Since these truth conditions make reference to the semantic values of the disjuncts of the complement clause, generating these truth conditions compositionally requires that we adopt a framework in which we can recover, from the semantic value of a disjunction, the semantic values of its disjuncts. The traditional possible worlds framework, on which the semantic value of a sentence is simply its truth set, makes this impossible. For example, given a disjunction that denotes $\{w_1, w_2, w_3\}$, there is no way to tell whether it was composed from disjuncts denoting $\{w_1, w_2\}$ and $\{w_3\}$, or from disjuncts denoting $\{w_1\}$ and $\{w_2, w_3\}$. The traditional framework thus ensures that propositional operators like modals and attitude verbs are blind to the disjunctive structure of their arguments.

Two related frameworks that allow a modal to see the disjunctive structure of its argument are *alternative semantics* (Kratzer and Shimoyama (2002), Alonso-Ovalle (2006), Aloni (2007)) and *inquisitive semantics* (Ciardelli and Roelofsen (2011) Aloni and Ciardelli (2013), Ciardelli et al. (2018), Ciardelli et al. (2017)). For reasons I leave to a footnote,[23] I will use a version of inquisitive semantics.

[23]One of the key differences between the two frameworks arises in cases when the set of worlds where 'q' is true is a subset of the worlds where 'p' is true ($q \subseteq p$). On the traditional possible worlds analysis of propositions and disjunction, in this case the proposition denoted by '$p \vee q$' is

I will present the model in two stages.[24] First, in §6, I will outline a semantics for modals in a fragment of propositional inquisitive semantics, where propositions consist of alternatives relevant for their *truth*. This will allow me to highlight the basics of the interaction between modals and disjunctions, and show how my semantics accounts for the original two puzzles as well as the data from §4. Then, I will note in §7 that this implementation of the theory faces two problems resulting from its interaction with negation: (i) it no longer supports impossibility and unnecessity distribution inferences over disjunctions (examples (19) and (20) below); and (ii), it gives up the duality between possibility and necessity modals. In response to these issues, I extend the account to a *bilateral* system, where propositions consist of *two* kinds of alternatives: those relevant for their truth and those relevant for their falsity.[25]

6 Minimal Covering Semantics

I will use a simple formal language to model the semantics I am arguing for, consisting of atomic formulae $p, q, ...$, the standard connectives $\neg, \wedge, \vee$, a possibility and a necessity modal $\Diamond, \Box$, and two special connectives: $\rightarrow$, a restrictor conditional; and !, the issue-cancelling operator of inquisitive semantics.

In standard possible worlds semantics, the semantic value of a sentence, a *classical proposition*, is just a set of worlds (or its characteristic function). Inquisitive semantics adds some complexity: an *inquisitive proposition* is a set of sets of worlds. The semantic value of an atomic formula is the set of all sets that only include worlds where it is true. In other words, it denotes its truth set, plus all of the subsets of that set. Let me illustrate with an example of two atomic sentences, p and q. Suppose p is true at w_1 and w_2, while q is true at w_2 and w_3.

identical to the one denoted by just 'p'. In inquisitive semantics, the same is true, the propositions denoted by 'p' and '$p \vee q$' come out the same. But standard versions of alternative semantics distinguish between these propositions (Roelofsen (2013), Ciardelli et al. (2017)). This means that alternative semantics, but not inquisitive semantics, gives up the traditional explanation of Hurford's constraint in terms of redundancy (Ciardelli and Roelofsen, 2017).

[24] A formal summary of the semantic framework and results is contained in the appendix at the end of this paper.

[25] The basic behavior of the connectives in the bilateral approach is formally similar to the 'radical inquisitive semantics' of Groenendijk and Roelofsen (2010) and Aher (2012), the dual update semantics in Willer (2018), the bilateral 'state-based' semantics of Aloni (2018), and the bilateral truthmaker semantics of Yablo (2014) and Fine (2017a,b). While these other theorists share a similar semantic framework, and some also share an interest in our two puzzles, none of these accounts offers a theory that supports the *Independence* inferences we are interested in in this paper.

Then, where $[\phi]$ is the inquisitive proposition denoted by ϕ, we have:

$$[p] = \{\{w_1, w_2\}, \{w_1\}, \{w_2\}, \varnothing\}$$
$$[q] = \{\{w_2, w_3\}, \{w_2\}, \{w_3\}, \varnothing\}$$

In inquisitive semantics, conjunctions and disjunctions of propositions are treated just as in traditional possible world semantics: they denote the operations of set intersection and union. The proposition denoted by $p \wedge q$, then, is just the set containing everything in common between the proposition denoted by p and the one denoted by q. The proposition denoted by $p \vee q$ is just the set containing everything that is included either in the one denoted by p or the one denoted by q. As applied to our simple example, we have:

$$[p \wedge q] = [p] \cap [q] = \{\{w_2\}, \varnothing\}$$
$$[p \vee q] = [p] \cup [q] = \{\{w_1, w_2\}, \{w_2, w_3\}, \{w_1\}, \{w_2\}, \{w_3\}, \varnothing\}$$

Two pieces of information made available by inquisitive propositions are important for our purposes. The first is the *informative content* of a sentence ϕ, written $\mathsf{info}(\phi)$. This is just the set of worlds where ϕ is true, corresponding directly to the traditional possible worlds proposition associated with ϕ. We can recover the informative content of any inquisitive proposition simply by taking the set of all of the worlds that are members of any element of the proposition (equivalently, the union of all sets of worlds in the proposition, $\mathsf{info}(\phi) = \bigcup[\phi]$). So, to continue our example, we have:

$$\mathsf{info}(p) = \bigcup[p] = \{w_1, w_2\}$$
$$\mathsf{info}(q) = \bigcup[q] = \{w_2, w_3\}$$
$$\mathsf{info}(p \wedge q) = \bigcup[p \wedge q] = \{w_2\}$$
$$\mathsf{info}(p \vee q) = \bigcup[p \vee q] = \{w_1, w_2, w_3\}$$

The second piece of information that the inquisitive semantics framework makes available is what sets it apart from the traditional possible worlds framework. This is the notion of the *alternatives* offered by a proposition ϕ, written $\mathsf{alt}(\phi)$. These correspond to the *largest* sets included in the inquisitive proposition denoted by a sentence, and will give us what we are after: the ability to recover from the semantic value of a disjunction the semantic values of each disjunct.[26] In our example, for instance, the alternatives offered by $p \vee q$ are the largest sets in $[p \vee q]$, i.e. $\{w_1, w_2\}$ and $\{w_2, w_3\}$. By no accident, these sets are identical to the informative content, or truth sets, of the two disjuncts, p and q:

$$\mathsf{alt}(p \vee q) = \{\{w_1, w_2\}, \{w_2, w_3\}\} = \{\mathsf{info}(p),\ \mathsf{info}(q)\}$$

Compare the case of $[p \wedge q]$, which has only one largest member, the singleton containing w_2, where both p and q are true. In this basic propositional fragment of inquisitive semantics, *only* disjunctions offer multiple alternatives.[27] Every sentence without a disjunction offers only a single alternative, namely, its informative content (i.e., for non-disjunctive ϕ, $\mathsf{alt}(\phi) = \{\mathsf{info}(\phi)\}$).

Using the notion of the alternatives offered by a proposition, then, we can give a semantics for modals that is sensitive to the disjunctive structure of their complement clauses. We will say that a necessity modal claim of the form $\Box\phi$ is true at w just in case the alternatives offered by ϕ form a minimal cover of the relevant set of worlds, $R(w)$. Similarly, a possibility modal claim of the form $\Diamond\phi$ will be true at w just in case the alternatives offered by ϕ form a minimal cover of a non-empty subset R' of the relevant set of worlds, $R(w)$:

$$\begin{aligned} &\Box\phi \text{ is true at } w \text{ iff } \mathsf{alt}(\phi) \text{ forms a minimal cover of } R(w) \\ &\Diamond\phi \text{ is true at } w \text{ iff there is a non-empty } R' \subseteq R(w) \\ &\qquad\qquad\qquad \text{such that } \mathsf{alt}(\phi) \text{ forms a minimal cover of } R' \end{aligned}$$

These truth conditions, seemingly tailored to the case when ϕ is a disjunction, reduce to the orthodox semantics for modals when ϕ contains no disjunction. Let me explain why. As mentioned, when ϕ contains no disjunction, it offers a single alternative, corresponding to its truth conditions, i.e. $\mathsf{alt}(\phi) = \{\mathsf{info}([\phi])\}$. Now, for a singleton set like $\{\mathsf{info}([\phi])\}$ to form a *minimal* cover of a (non-empty) set S is just for it to cover S *simpliciter*, i.e. for $S \subseteq \mathsf{info}([\phi])$. Thus, for a necessity modal $\Box$, when ϕ contains no disjunction, it follows that $\Box\phi$ is true at w iff ϕ is

[26]'Largest sets' here means the sets in the proposition such that there is no proper superset also in the proposition. Officially:

$$\mathsf{alt}(\phi) = \{s \in [\phi] \mid \text{for every } t \in [\phi],\ \text{if } s \subseteq t \text{ then } s = t\ \}$$

[27]In richer versions of inquisitive semantics, other expressions like existential quantifiers and interrogative operators also introduce multiple alternatives.

true at every relevant world, i.e. iff $R(w) \subseteq \mathsf{info}(\phi)$. For possibility modals, when ϕ contains no disjunction, $\Diamond\phi$ is true at a world w iff there is a non-empty subset R' of relevant worlds such that R' is a subset of $\mathsf{info}(\phi)$. In other words, $\Diamond\phi$ is true iff there is at least one relevant world where ϕ is true.

In order to discuss the payoffs of my semantics, we must first talk about the semantics of negation ($\neg$) and the restrictor conditional ($\rightarrow$). In basic inquisitive semantics, the negation $\neg\phi$ of a proposition ϕ denotes the set of all sets of worlds that have no overlap with $\mathsf{info}(\phi)$. To illustrate by way of our running example, suppose there is just one more world, w_4, where both atoms p, q are false. Then $[\neg p]$ is the set of all sets of worlds that do not overlap with $\mathsf{info}(p) = \{w_1, w_2\}$. This means we have:

$$[\neg p] = \{\{w_3, w_4\}, \{w_3\}, \{w_4\}, \varnothing\}$$

I will use the conditional $\rightarrow$ solely to express conditional restriction effects. So we will say that $\phi \rightarrow \psi$ is true at w iff ψ is true at w when the set of relevant worlds $R(w)$ is altered so as to only include worlds where ϕ is true: $R(w) \cap \mathsf{info}(\phi)$.[28]

Now let us turn to the results, for which we will suppose that for atomic p, q, $p \vee q$ satisfies Hurford's constraint, i.e. neither p nor q entails the other.

Independence Inferences. Suppose that either $\Diamond(p \vee q)$ or $\Box(p \vee q)$ is true at w. Either way, it follows that $\{\mathsf{info}(p), \mathsf{info}(q)\}$ forms a minimal cover of some non-empty subset of the relevant worlds, R' ($R' = R(w)$ in the case of $\Box(p \vee q)$). Then there are at least two relevant worlds in R', call them v, u, such that one of them, say v, is included in $\mathsf{info}(p) \backslash \mathsf{info}(q)$, while the other, u, is included in $\mathsf{info}(q) \backslash \mathsf{info}(p)$. Now, $\Diamond(p \wedge \neg q)$ is true iff the alternatives offered by $p \wedge \neg q$ form a minimal cover of a non-empty subset of relevant worlds. Since $p \wedge \neg q$ just offers a single alternative, corresponding to the worlds where it is true, and v is a $\neg q$-world, it follows that $\{v\}$ is a subset of relevant worlds that is minimally covered by the alternatives offered by $p \wedge \neg q$, i.e. $\mathsf{info}(p \wedge \neg q)$. So $\Diamond(p \wedge \neg q)$ is true. A similar argument shows that $\{u\}$ makes the other Independence inference, $\Diamond(q \wedge \neg p)$, true.[29]

Independence Conditionals. I will spell out just the necessity modal case, but the extension to possibility modals is straightforward. Assume that $\Box(p \vee q)$ is true. Then the alternatives offered by $p \vee q$, i.e. $\{\mathsf{info}(p), \mathsf{info}(q)\}$, form a minimal cover

[28] See the dynamic model update conditional in van Ditmarsch et al. (2008), or the appendix of this paper, for precise versions of such a conditional.

[29] In our running example, the truth of $\Diamond(p \vee q)$ or $\Box(p \vee q)$ requires that w_1, w_3 be among the relevant alternatives.

of the relevant set of worlds. This means again that every relevant world is either a p-world or a q-world, and further, that there are at least two relevant worlds, $v \in \mathsf{info}(p) \backslash \mathsf{info}(q)$ and $u \in \mathsf{info}(q) \backslash \mathsf{info}(p)$. On our semantics for the restrictor conditional, $\neg p \rightarrow \Box q$ is true iff $\Box q$ is true when we ignore all relevant worlds except the not-p worlds. Doing leaves us only with q-and-not-p-worlds like u. Since all such worlds are in the truth set of q, $\{\mathsf{info}(q)\}$ forms a minimal cover of them. Thus, $\Box q$ is true on the restriction to the $\neg p$ worlds, so $\neg p \rightarrow \Box q$ is true on our original assumption. A parallel argument shows $\neg q \rightarrow \Box p$ is true.

Free Choice. The free choice inference is predicted to be valid. Suppose $\Diamond(p \vee q)$ is true. Then the set of alternatives offered by $p \vee q$, namely $\{\mathsf{info}(p), \mathsf{info}(q)\}$, forms a minimal cover of some subset of the relevant worlds. Since the cover is *minimal*, it follows that there is at least one relevant p-and-not-q-world, call it v, and one q-and-not-p-world, call it u. Thus, there is a non-empty subset of relevant worlds, namely $\{v\}$, that is minimally covered by $\{\mathsf{info}(p)\}$. So $\Diamond p$ is true. Likewise, there is a non-empty subset of relevant worlds, namely $\{u\}$, that is minimally covered by $\{\mathsf{info}(q)\}$. So $\Diamond q$ is true. Thus, the free choice inference is validated: $\Diamond(p \vee q)$ entails $\Diamond p$ and $\Diamond q$.

Ross Inference. The Ross inference, by contrast, is invalid. Suppose p stands for 'Alicia mails the letter', and further, that $\Box p$, 'Alicia ought to mail the letter', is true. Then, the alternatives offered by p, namely $\{\mathsf{info}(p)\}$, cover the relevant set of worlds, $R(w)$, where Alicia fulfills her obligation. Now, consider the truth conditions of $\Box(p \vee q)$, where q is an independent disjunct like, 'Alicia burns the letter'. $\Box(p \vee q)$ is true iff the alternatives offered by $p \vee q$, namely $\{\mathsf{info}(p), \mathsf{info}(q)\}$, form a minimal cover of $R(w)$. But, they cannot. Since $\Box p$ is true, $\{\mathsf{info}(p)\}$ is a strictly smaller cover of $R(w)$. Thus, $\Box(p \vee q)$ is not true, and the argument is invalid.

Extended Ross's Puzzle. Suppose $\Box p$ and $\Diamond q$, are true. Because $\Box p$ is true, $\{\mathsf{info}(p)\}$ covers the relevant worlds. And since $\{\mathsf{info}(p)\}$ is a minimal cover of the relevant worlds, $\{\mathsf{info}(p), \mathsf{info}(q)\}$ is not. Thus, the conclusion of the extended Ross's puzzle, $\Box(p \vee q)$, is not true. It does not matter whether q is compatible with the relevant alternatives; the inference is invalid because q does not pick out an *independent* alternative relative to p.

Flexibility. Before turning to some difficulties faced by the present minimal covering semantics in the next section, I want to note that standard inquisitive semantics posits the linguistic resources to allow for some flexibility when it comes to drawing the Independence inferences. Indeed, standard inquisitive semantics posits the existence of an 'issue-cancelling' operator (denoted by '!'), which has the effect of eliminating the distinction between alternatives in an inquisitive proposition: any previously distinguished alternatives are collapsed into a single, undifferentiated one. This operator has been posited to distinguish between the semantic values of two questions one can ask using a sentence like 'Does Alicia speak Hindi or Tamil?':

(15) a. Does Alicia speak Hindi-or-Tamil$^{\uparrow}$?
b. Does Alicia speak Hindi$^{\uparrow}$ or Tamil$^{\downarrow}$?

(where $\uparrow/\downarrow$ indicate rising/falling intonation. The intended reading of the first question, (15a), is polar; it can be resolved with *Yes* (Alicia speaks at least one of Hindi and Tamil) or *No* (Alicia speaks neither language). By contrast, the second question, (15b), is not polar, and the conventional ways to resolve it directly are to either say that Alicia speaks Hindi, or to say that she speaks Tamil. In inquisitive semantics, the difference in the conventional resolutions of these questions is explained by a difference in the *alternatives* they offer. Most relevant for our purposes, (15b) treats Hindi and Tamil as distinct alternatives, but (15a) treats Hindi-or-Tamil as a single alternative. In order to generate *different* alternatives for sentences with the same surface structure, inquisitive semantics hypothesizes that the polar question (15a) contains an operator (represented with '!') that erases the distinction between the Hindi and Tamil alternatives. For our purposes, the semantics of ! is important insofar as it gives rise to the following identities. For any sentence ϕ:[30]

$$\mathsf{alt}(!\phi) = \{\mathsf{info}(\phi)\}$$
$$\mathsf{info}(!\phi) = \mathsf{info}(\phi)$$

So $!\phi$ is true iff ϕ is, but its proposition always offers a single alternative corresponding to its truth conditions. Above, I explained that for a non-disjunctive formula ϕ, my minimal covering semantics assigns orthodox truth conditions to $\Box\phi$ and $\Diamond\phi$. Since the ! operator makes even a disjunction offer just a single alternative, the same reasoning extends to arbitrary formulae of the form $!\phi$. Even when ϕ contains disjunction, my semantics assigns orthodox truth conditions to $\Box!\phi$ and $\Diamond!\phi$.

With this operator, we can explain how in certain special contexts, the Independence inferences may not be licensed. For example, in an epistemology class,

[30] The official semantics for the operator is:

$$[!\phi] = \wp(\mathsf{info}(\phi))$$

one might hear the following:

(16) a. Given the evidence, Smith must own a Ford.
b. So, given the evidence, Smith must either own a Ford or live in Barcelona.

On my view, the premise of (16) rules out that *Smith lives in Barcelona* is an independent alternative relative to his owning a Ford. This would normally ensure that the conclusion is not true. In order to make the discourse coherent, then, interpreters may insert the issue-cancelling operator (!) just above the disjunction. When the inference in (16) is heard as acceptable, this is because it is not actually of the form $\Box p \Rightarrow \Box(p \vee q)$. Rather, it has the form $\Box p \Rightarrow \Box!(p \vee q)$. In support of this hypothesis, notice that just as with the question (15a), a monotonous intonation seems to encourage acceptability:

(17) a. Given the evidence, Smith must own a Ford.
b. So, given the evidence, Smith must either own-a-Ford-or-live-in-Barcelona.

In contrast, an alternating intonation, as with (15b), makes it harder to accept:

(18) a. Given the evidence, it must be that Smith owns a Ford.
b. # So, Smith must either own a Ford$^{\uparrow}$ or live in Barcelona$^{\downarrow}$.

Note that this form of flexibility is very different from a pragmatic account of the Independence inferences. On a semantic account like mine, the unorthodox content captured by the Independence inferences is part of the default, literal meaning of modals with disjunctive complements, and *not* drawing these inferences requires special interpretive work.

7 A Bilateral Version

The minimal covering semantics improves upon the orthodox semantics when it comes to predicting the meanings of unembedded modal claims. But, like some other semantic accounts of the puzzles, it is thereby worse than the orthodox semantics at predicting the meanings of modal claims embedded under downward entailing operators like negation.[31] By assigning stronger-than-orthodox truth conditions to a bare disjunctive modal claim ($\Delta(p \vee q)$), the theory assigns weaker-than-orthodox falsity conditions to it, and thus assigns weaker truth conditions to its negation $\neg\Delta(p \vee q)$. In particular, it gives up the validity of the

[31] For example, as Alonso-Ovalle (2006) shows, the semantics of Simons (2005) suffers these problems with negation. See Aloni (2018) and Willer (2018) for alternative bilateral solutions to these problems, and Aloni (2007) for a unilateral response to these issues based on ambiguity.

following two patterns of inference, both of which are validated by the orthodox semantics:

Impossibility Distribution.

(19) a. Alicia may not mail the letter or burn it. $\neg\Diamond(M \vee B)$
b. So, Alicia may not mail the letter. $\neg\Diamond M$
c. So, Alicia may not burn the letter. $\neg\Diamond B$

Unnecessity Distribution.

(20) a. Alicia doesn't have to either mail the letter or burn it. $\neg\Box(M \vee B)$
b. So, Alicia doesn't have to mail the letter. $\neg\Box M$
c. So, Alicia doesn't have to burn the letter. $\neg\Box B$

A second problem is that the basic minimal covering semantics gives up predicting the *duality* of possibility and necessity modals. Duality (the truth conditional equivalence of $\neg\Box\phi$ and $\Diamond\neg\phi$ on the one hand, and $\neg\Diamond\phi$ and $\Box\neg\phi$ on the other) is one of the fundamental virtues of the orthodox analysis, which derives the duality of necessity and possibility modals from the duality of the universal and existential quantifiers. To illustrate where duality fails on the minimal covering semantics of the previous section, consider the following pair:

Duality.

(21) a. Alicia may not mail the letter or burn it $\neg\Diamond(M \vee B)$
b. Alicia must not mail the letter or burn it $\Box\neg(M \vee B)$

Intuitively, these two sentences are equivalent. Orthodox semantic theories predict as much. The basic minimal covering semantics I have sketched, however, strengthens the truth conditions of $\Diamond(M \vee B)$, weakening the truth conditions of $\neg\Diamond(M \vee B)$. It says that (21a) is true iff the alternatives offered by $M \vee B$ (Alicia's mailing the letter, Alicia's burning the letter) do not form a minimal cover of the relevant set of worlds. This would hold, for example, in case in all of the relevant worlds, Alicia burns the letter (i.e. if $\Box B$ were true). Meanwhile, it assigns orthodox truth conditions to (21b): it is true iff in every relevant world, M and B are both false (so, e.g., it entails $\Box\neg M$ and $\Box\neg B$). Thus, while the truth of (21b) prohibits Alicia from burning the letter, the basic minimal covering semantics allows (21a) to be true even if Alicia *must* burn the letter. Clearly, the predicted truth conditions of (21a) are far too weak.

Fortunately, a natural extension of the semantics solves these problems. The key is to define the falsity conditions of our sentences independently of their truth conditions, and to allow for truth value gaps.[32] I will use the same language as before, but now assign *bilateral inquisitive propositions* to sentences. These bilateral inquisitive propositions will be modeled as *pairs* of regular inquisitive propositions that have only one member in common: the empty set. To return to our simple example of four worlds, we now assign to p a *positive* part of its bilateral proposition (written $[p]^+$), corresponding to the unilateral inquisitive proposition from before:

$$[p]^+ = \{\{w_1, w_2\}, \{w_1\}, \{w_2\}, \varnothing\}$$

and we define the negative part (written $[p]^-$), as the set of worlds where p is false, plus all of its subsets (corresponding to what was previously the unilateral value of $\neg p$):

$$[p]^- = \{\{w_3, w_4\}, \{w_3\}, \{w_4\}, \varnothing\}$$

Extending these principles to our other atomic sentence, q, we have:

$$[q]^+ = \{\{w_2, w_3\}, \{w_2\}, \{w_3\}, \varnothing\}$$
$$[q]^- = \{\{w_1, w_4\}, \{w_1\}, \{w_4\}, \varnothing\}$$

In order to make use of these negative parts, I will treat negation as an involution — it swaps the negative part of a proposition for its positive part, and vice versa:

$$[\neg\phi]^+ = [\phi]^-$$
$$[\neg\phi]^- = [\phi]^+$$

I will treat the positive contribution of conjunction and disjunction the same as before, namely as set intersection and union, respectively:

$$[\phi \wedge \psi]^+ = [\phi]^+ \cap [\psi]^+$$
$$[\phi \vee \psi]^+ = [\phi]^+ \cup [\psi]^+$$

As for the negative contributions of these connectives, I will, for simplicity, draw on the De Morgan equivalences. The negative part of a conjunction will be the

[32]See Groenendijk and Roelofsen (2010), Aloni (2018) and Willer (2018) for other versions of bilateral inquisitive semantics with similar motivations.

union of the negative parts of its conjuncts, and the negative part of a disjunction will be the intersection of the negative parts of its disjuncts:[33]

$$[\phi \wedge \psi]^- = [\phi]^- \cup [\psi]^-$$
$$[\phi \vee \psi]^- = [\phi]^- \cap [\psi]^-$$

I will also extend definitions of the informative content of a sentence, and the alternatives offered by a sentence, to the bilateral model. Where we used to have just the positive alternatives offered by a sentence, we now have its positive ($\mathsf{alt}([\phi]^+)$) and negative ($\mathsf{alt}([\phi]^-)$) alternatives, defined in the same way as before. And where we used to have just the positive informative content of a sentence, we will now have its positive ($\mathsf{info}([\phi]^+)$) and negative ($\mathsf{info}([\phi]^-)$) informative contents, which correspond to the sets of worlds where it is true and false, respectively.

Figure 4 illustrates some simple examples of propositions. In each subfigure, the four circles are worlds corresponding to the classical valuations of the atomic sentences p, q (with $\overline{p}$ indicating p is false), and correspond to our informal model, starting with w_1 in the upper right, and w_2, w_3, w_4 moving counterclockwise around the square. Solid lines represent positive alternatives, while dotted lines represent negative alternatives.

Now, let us return to our modals. We define the *truth* conditions of the modal propositions as we did previously, though now with the detail that they depend

[33]It is easy to see that given the semantics of negation, conjunction, and disjunction I outline here, the semantics predicts that each of the classical De Morgan equivalences holds. Some of these equivalences are controversial, especially when embedded under modals or conditionals. The theory I give here thus inherits some of this controversy. For example, it validates *Dual Free Choice*: when p and q are atoms that obey Hurford's constraint, $\neg\Box(p \wedge q) \vDash \Diamond\neg p \wedge \Diamond\neg q$. For given duality between $\Box$ and $\Diamond$, $[\neg\Box(p \wedge q)] = [\Diamond\neg(p \wedge q)]$. Then, given the De Morgan equivalences, $[\Diamond\neg(p \wedge q)] = [\Diamond(\neg p \vee \neg q)]$. Finally, given the validity of free choice, clearly $\Diamond(\neg p \vee \neg q)$ entails $\Diamond\neg p$ and $\Diamond\neg q$. The obvious culprit appears to be the De Morgan equivalence between $\neg(p \wedge q)$ and $\neg p \vee \neg q$, which are not equivalent on the unilateral model of the previous section or in standard inquisitive semantics. One way to modify the present system in order to invalidate dual free choice would be to change the rule for the negative part of conjunction, so that it is not equivalent to a disjunction of negations. Since this issue independent of the original puzzles and the data adduced in §4, I do not want to take a stand on it here. For simplicity and completeness, I have opted to validate all of the De Morgan equivalences. For further discussion of the De Morgan equivalences in modal and conditional contexts, see Fox (2007); Chemla (2009); Ciardelli et al. (2018); Romoli and Santorio (2019); Marty et al. (ms).

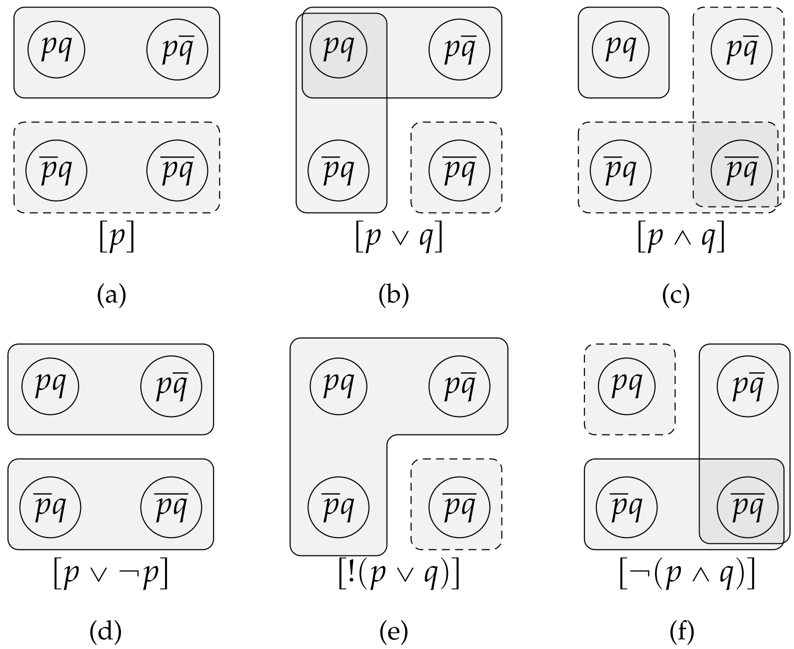

Figure 4: Examples of bilateral propositions.

on the *positive* alternatives offered by the proposition:

$\Box\phi$ is true at w iff $\mathsf{alt}([\phi]^+)$ forms a minimal cover of $R(w)$
$\Diamond\phi$ is true at w iff there is a non-empty $R' \subseteq R(w)$
such that $\mathsf{alt}([\phi]^+)$ forms a minimal cover of R'

Ensuring duality means making $\neg\Box\phi$ equivalent to $\Diamond\neg\phi$, and $\neg\Diamond\phi$ equivalent to $\Box\neg\phi$. Since the negative alternatives of a proposition are identical to the positive alternatives of its negation (i.e. $\mathsf{alt}([\phi]^-) = \mathsf{alt}([\neg\phi]^+)$), duality leads to the following *falsity* conditions:

$\Box\phi$ is false at w iff there is a non-empty $R' \subseteq R(w)$
such that $\mathsf{alt}([\phi]^-)$ forms a minimal cover of R'
$\Diamond\phi$ is false at w iff $\mathsf{alt}([\phi]^-)$ forms a minimal cover of $R(w)$

Duality. Now let me unpack this semantics by showing how it recovers the truth conditional equivalence of (21a) and (21b). Recall those sentences:

(21)	(21a)	Alicia may not mail the letter or burn it	$\neg\Diamond(M \vee B)$
	(21b)	Alicia must not mail the letter or burn it	$\Box\neg(M \vee B)$

I will show that these sentences are true and false in the same circumstances. First, truth. By our semantics for negation, $\neg\Diamond(M \vee B)$ is true just in case $\Diamond(M \vee B)$ is false. $\Diamond(M \vee B)$ is false iff the negative alternatives offered by $M \vee B$ form a minimal cover of the set of relevant worlds. Since the negative alternatives offered by $M \vee B$ are equal to the positive alternatives offered by $\neg(M \vee B)$, it follows that $\Diamond(M \vee B)$ is false exactly whenever $\Box\neg(M \vee B)$ is true. Thus, $\neg\Diamond(M \vee B)$ is true iff $\Box\neg(M \vee B)$ is true.

Now for falsity. $\neg\Diamond(M \vee B)$ is false iff $\Diamond(M \vee B)$ is true, i.e. when the positive alternatives of $M \vee B$ form a minimal cover of a non-empty subset of relevant worlds. $\Box\neg(M \vee B)$, on the other hand, is false iff the negative alternatives of $\neg(M \vee B)$ form a minimal cover of a non-empty subset of relevant worlds. Since the negative alternatives of $\neg(M \vee B)$ are identical to the positive alternatives of $M \vee B$, the two sentences are false in exactly the same worlds.

By ensuring duality, we once again predict the validity of impossibility and unnecessity distribution over disjunction (19)-(20). Consider impossibility distribution, repeated here:

Impossibility Distribution.

(19a)	Alicia may not mail the letter or burn it.	$\neg\Diamond(M \vee B)$
(19b)	So, Alicia may not mail the letter.	$\neg\Diamond M$
(19c)	So, Alicia may not burn the letter.	$\neg\Diamond B$

By Duality, the premise $\neg\Diamond(M \vee B)$ is equivalent to $\Box\neg(M \vee B)$, which by the De Morgan laws is equivalent to $\Box(\neg M \wedge \neg B)$, which is true just in case $\mathsf{alt}([\neg M \wedge \neg B]^+) = \{\mathsf{info}([M]^-) \cap \mathsf{info}([B]^-)\}$) is a minimal cover of $R(w)$. If that latter condition holds, then $\{\mathsf{info}([M]^-)\} = \mathsf{alt}([M]^-)$ is also a minimal cover of $R(w)$, meaning that $\Diamond(M)$ is false at w, so $\neg\Diamond M$ is true. Similarly for $\neg\Diamond B$.

A similar argument works for unnecessity distribution, repeated here:

Unnecessity Distribution.

(20a)	Alicia doesn't have to either mail the letter or burn it.	$\neg\Box(M \vee B)$
(20b)	So, Alicia doesn't have to mail the letter.	$\neg\Box M$
(20c)	So, Alicia doesn't have to burn the letter.	$\neg\Box B$

When (20a) is true, $\Box(M \vee B)$ is false. It follows that the negative alternatives of $M \vee B$, i.e. the singleton containing all worlds where both M and B are false, cover a non-empty set of relevant worlds. If that's the case, then the larger set of worlds where just M is false *also* covers that set. The same goes for the larger set of worlds where just B is false. But then, this means that $\Diamond\neg M$ and $\Diamond\neg B$ are true. By Duality, these are equivalent to $\neg\Box M$ and $\neg\Box B$.

Moving to this bilateral system thus allows us to regain the orthodox predictions about the truth conditions of disjunctive modal claims embedded under negation and other downward entailing operators.

8 Truth Value Gaps

Adopting the bilateral framework I have outlined here requires accepting some truth value gaps, but I think that they are well placed.[34] First, if we assume atomic formulae are either true or false in every world, then it is only sentences containing modals that generate truth value gaps. Further, there are no truth value gluts. So, the bilateral system ensures that the non-modal fragment behaves classically.

In semantics, truth value gaps are most often taken to indicate presupposition failure. But other kinds of gaps have been postulated. In particular, predicates of plurals seem to generate gaps as a result of 'homogeneity' effects.[35] Consider the following pair:

(22) a. Alicia and Bulmaro saw the movie. They liked it.
b. Alicia and Bulmaro saw the movie. They did not like it.

The second sentence of (22a) is true iff *both* Alicia and Bulmaro liked the movie; while its negation in (22b) is true iff *neither* Alicia nor Bulmaro liked it. There is a gap between these truth conditions: if Alicia but not Bulmaro liked the movie, both (22a) and (22b) seem to misdescribe the situation. On some theories of plural predication, this is because the second sentence of each example is neither true nor false in that situation. The truth value gap arises because the predicate 'liked it' expects its plural arguments to be homogeneous with respect to it: either *all* or *none* of the individuals in a collection satisfy it.

[34]The bilateral, dynamic inquisitive semantics of Willer (2018) also postulates some truth value gaps for modal sentences.

[35]See for example, Schwarzschild (1993) and Križ (2015, 2016). It is controversial whether homogeneity effects give rise to gaps, and if they do, whether these gaps should be thought of as presuppositions or not.

For possibility modals, there appears to be a similar kind of truth value gap: when $\Diamond(p \vee q)$ is true, p and q are both possible (by free choice). When $\Diamond(p \vee q)$ is false (i.e. when $\neg\Diamond(p \vee q)$ is true), p and q are both *impossible* (by impossibility distribution). If that is correct, then $\Diamond(p \vee q)$ is neither true nor false when only one of p, q is possible. And this is exactly what my account predicts. In this respect, my account concurs with Goldstein (2019)'s theory of possibility modals, which writes homogeneity directly into their semantics. In contrast to Goldstein's homogeneity account, my theory posits a second source of truth value gap for sentences like $\Diamond(p \vee q)$. On my account, $\Diamond(p \vee q)$ is also neither true nor false in case p and q are both possible but not *independently* so: that is, when $\Diamond p$ and $\Diamond q$ are true, but either $\Diamond(p \wedge \neg q)$ or $\Diamond(q \wedge \neg p)$ is not. This second source of truth value gap helps explain the anomalous status of (12) in the described scenario. To recall, you are playing a card game, and your only option is to pick up the pair consisting of the ace and the ten of diamonds. Someone says:

(12) You may either take the ace of diamonds or the ten of diamonds.

(12) seems to misdescribe things — it appears to say that there are at least two (independent) options. But its negation seems equally off the mark:

(12′) You may not take the ace of diamonds or the ten of diamonds.

My bilateral semantics can explain why both (12) and (12′) seem to misdescribe things: neither of them are true.

For necessity modals, truth value gaps arise precisely in the cases at issue in Ross's puzzle. Suppose 'Alicia ought to mail the letter' ($\Box M$) is true at w. The truth of the Ross inference conclusion, 'Alicia ought to either mail the letter or burn it' ($\Box(M \vee B)$) requires that M and B are independent alternatives in the relevant set of worlds. If the Ross premise $\Box M$ is true, then B is not an independent alternative, so $\Box(M \vee B)$ is not true. But neither is it false. By duality and one of the De Morgan equivalences, it is false just in case $\Diamond(\neg M \wedge \neg B)$ is true. The truth of that latter sentence obviously requires that there are relevant not-M-worlds. But since $\Box M$ is true, there are none. My account thus predicts that the premise does not merely *fail* to ensure the truth of the conclusion; it ensures the conclusion is *not* true.[36]

If my account is correct in this prediction, then it offers a diplomatic resolution of the disagreement over Ross's puzzle between orthodox and revisionary semantic accounts. Proponents of revisionary semantic accounts have insisted

[36] This may be why Ross inferences are so strongly repugnant: the premise not only fails to ensure the truth of the conclusion: it ensures the conclusion is *neither true nor false*.

that $\Box M$ does not entail $\Box(M \vee B)$. Defenders of the orthodox semantics sense that it must, since $\neg\Box(M \vee B)$ entails $\neg\Box M$. If these sentences are always true or false, then at most one of these claims can be correct. But my account rejects this assumption. In particular, when $\Box M$ is true, $\Box(M \vee B)$ is neither true nor false, so $\Box M$ does not entail $\Box(M \vee B)$, even though $\neg\Box(M \vee B)$ entails $\neg\Box M$. Thus, we may agree with both parties on these points.

9 Conclusion

I have argued that both Ross's puzzle and the puzzle of free choice permission should be explained in terms of the licensing of *Independence inferences*, which are stronger than the standard Diversity analysis predicts. I have also given a bilateral inquisitive semantics for the interaction between modals and disjunctions that generates the validity of these inferences, and showed how it can explain the full range of data adduced in this paper. Of course, there are many related issues that remain to be investigated: for example, whether similar puzzles for the logic of conditionals with embedded disjunctions ought to be explained in terms of (suitably translated) Independence inferences; and whether there is some deeper computational or philosophical explanation for the semantic interaction between modals and disjunctions I defend here. These questions will have to be left for future research.

Appendix: Bilateral Minimal Covering Semantics

Language.
Given a countable set of atomic sentence letters, At, and $p \in \mathsf{At}$, wffs are defined by the following grammar:

$$p \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid\, !\phi \mid \phi \to \psi \mid \Diamond\phi \mid \Box\phi$$

Models.
A model $\mathcal{M}$ is a triple $\mathcal{M} = (W_{\mathcal{M}}, R_{\mathcal{M}}, V_{\mathcal{M}})$ where $W_{\mathcal{M}}$ is a set of worlds, $R_{\mathcal{M}}$ is a function from worlds to sets of worlds ($R_{\mathcal{M}} : W \mapsto \wp(W)$), and $V_{\mathcal{M}}$ is a function from atomic sentences to truth sets ($V_{\mathcal{M}} : \mathsf{At} \mapsto \wp(W)$).

Bilateral Propositions.
A bilateral proposition P in a model $\mathcal{M}$ is a pair (P^+, P^-) of downward-closed (relative to the subset relation) sets of sets of worlds, such that their intersection

is the singleton containing the empty set. In other words, where $P^\circ \in \{P^+, P^-\}$, and $s, t \subseteq W_\mathcal{M}$:

$$\begin{aligned} P^\circ &\subseteq \wp(W_\mathcal{M}) \\ \text{if } s \in P^\circ \text{ and } t \subseteq s &\Rightarrow t \in P^\circ \\ P^+ \cap P^- &= \{\varnothing\} \end{aligned}$$

Let $\mathcal{P}_\mathcal{M}$ be the set of all bilateral propositions in $\mathcal{M}$.

Information and Alternatives
For $P \in \mathcal{P}_\mathcal{M}$:

$$\begin{aligned} \mathsf{info}(P^+) &= \bigcup P^+ \quad \text{(the set of worlds where } P \text{ is true)} \\ \mathsf{info}(P^-) &= \bigcup P^- \quad \text{(the set of worlds where } P \text{ is false)} \\ \mathsf{alt}(P^+) &= \{s \in P^+ \mid \neg\exists t \in P^+ : \ t \supset s\} \quad \text{(the positive alternatives offered by } P\text{)} \\ \mathsf{alt}(P^-) &= \{s \in P^- \mid \neg\exists t \in P^- : \ t \supset s\} \quad \text{(the negative alternatives offered by } P\text{)} \end{aligned}$$

Minimal Cover

$$\begin{aligned} C \text{ is a cover of } S \text{ iff } \ & S \subseteq \bigcup C \\ C \text{ is a minimal cover of } S \text{ iff } \ & C \text{ is a cover of } S \text{ and} \\ & \text{there is no } C' \subset C : C' \text{ is a cover of } S \end{aligned}$$

Model Update.
The ϕ-update of an accessibility function $R_\mathcal{M}$, written $R_\mathcal{M} \upharpoonright \phi$, is defined as follows:

$$R_\mathcal{M} \upharpoonright \phi = \lambda w \in W_\mathcal{M}.\ \mathsf{info}^+([\phi]_\mathcal{M}) \cap R_\mathcal{M}(w)$$

So $R_\mathcal{M} \upharpoonright \phi$ is the function that takes w to the subset of $R_\mathcal{M}(w)$ where ϕ is true in $_\mathcal{M}$. We use this notion to define a model update. The ϕ-update of a model $\mathcal{M}$, written $(\mathcal{M} \upharpoonright \phi)$ is defined:

$$\mathcal{M} \upharpoonright \phi = (W_\mathcal{M}, R_\mathcal{M} \upharpoonright \phi, V_\mathcal{M})$$

Semantics.
The proposition denoted by a wff ϕ in a model $\mathcal{M}$ is denoted $[\phi]_\mathcal{M} = ([\phi]^+_\mathcal{M}, [\phi]^-_\mathcal{M})$ (I drop the model subscript for readability except when important). For atomic

$p \in \mathsf{At}$:

$$[p] = (\wp(V(p)), \wp(W \backslash V(p)))$$

Non-modal complex ϕ:

$$\begin{aligned}
[\neg\phi] &= ([\phi]^-, [\phi]^+) \\
[\phi \wedge \psi] &= ([\phi]^+ \cap [\psi]^+, [\phi]^- \cup [\psi]^-) \\
[\phi \vee \psi] &= ([\phi]^+ \cup [\psi]^+, [\phi]^- \cap [\psi]^-) \\
[!\phi] &= (\wp(\mathsf{info}([\phi]]^+)), \wp(\mathsf{info}([\psi]^-))) \\
[\phi \to \psi]_{\mathcal{M}} &= ([\psi]^+_{\mathcal{M}\restriction\phi}, [\psi]^-_{\mathcal{M}\restriction\phi})
\end{aligned}$$

For modal ϕ:

$$\begin{aligned}
[\Box\phi]^+ &= \wp(\{w \in W \mid \mathsf{alt}([\phi]^+) \text{ is a minimal cover of } R(w)\}) \\
[\Box\phi]^- &= \wp(\{w \in W \mid \text{ there is a non-empty } R' \subseteq R(w) \text{ such that} \\
&\qquad \mathsf{alt}([\phi]^-) \text{ is a minimal cover of } R'\} \\
[\Diamond\phi]^+ &= \wp(\{w \in W \mid \text{ there is a non-empty } R' \subseteq R(w) \text{ such that} \\
&\qquad \mathsf{alt}([\phi]^+) \text{ is a minimal cover of } R'\}) \\
[\Diamond\phi]^- &= \wp(\{w \in W \mid \mathsf{alt}([\phi]^-) \text{ is a minimal cover of } R(w)\})
\end{aligned}$$

Truth and Falsity
A sentence ϕ is true at a point M, w iff $w \in \mathsf{info}([\phi]^+_{\mathcal{M}})$.
A sentence ϕ is false at a point M, w iff $w \in \mathsf{info}([\phi]^-_{\mathcal{M}})$.

Entailment
A sentence ϕ entails a sentence ψ in a model $\mathcal{M}$ (written $\phi \vDash_{\mathcal{M}} \psi$), when the following condition holds:

$$\phi \vDash_{\mathcal{M}} \psi \text{ iff } \mathsf{info}^+([\phi]_{\mathcal{M}}) \subseteq \mathsf{info}^+([\psi]_{\mathcal{M}})$$

Summary of Results.
As mentioned above, since I am using this simple language to model natural language disjunction, I am interested in how modals interact with disjunctions that obey *Hurford's constraint*. For this reason, I focus on the set of all models such that for $p, q \in \mathsf{At}$, neither entails the other; we will call these *admissible*.

Definition (Admissible model)**.** A model $\mathcal{M}$ is *admissible* iff for $p, q \in \mathsf{At}$:

$$V_{\mathcal{M}}(p) \nsubseteq V_{\mathcal{M}}(q)$$
$$V_{\mathcal{M}}(q) \nsubseteq V_{\mathcal{M}}(p)$$

Let $\mathfrak{M}$ be the set of all admissible models, and we will say that $\phi \vDash_{\mathfrak{M}} \psi$ iff for every $\mathcal{M} \in \mathfrak{M}$, $\phi \vDash_{\mathcal{M}} \psi$.

Fact (Independence inferences)**.**

$$\Box(p \vee q) \vDash_{\mathfrak{M}} \Diamond(p \wedge \neg q)$$
$$\Box(p \vee q) \vDash_{\mathfrak{M}} \Diamond(q \wedge \neg p)$$
$$\Diamond(p \vee q) \vDash_{\mathfrak{M}} \Diamond(p \wedge \neg q)$$
$$\Diamond(p \vee q) \vDash_{\mathfrak{M}} \Diamond(q \wedge \neg p)$$

Fact (Independence Conditionals)**.**

$$\Box(p \vee q) \vDash_{\mathfrak{M}} \neg p \rightarrow \Box q$$
$$\Box(p \vee q) \vDash_{\mathfrak{M}} \neg q \rightarrow \Box p$$
$$\Diamond(p \vee q) \vDash_{\mathfrak{M}} \neg p \rightarrow \Diamond q$$
$$\Diamond(p \vee q) \vDash_{\mathfrak{M}} \neg q \rightarrow \Diamond p$$

Fact (Ross Inference)**.**

$$\Box p \nvDash_{\mathfrak{M}} \Box(p \vee q)$$

Fact (Free Choice)**.**

$$\Diamond(p \vee q) \vDash_{\mathfrak{M}} \Diamond p$$

Fact (Modal Duality)**.** *For any ϕ in any model $\mathcal{M}$:*

$$[\Box\phi]_{\mathcal{M}} = [\neg\Diamond\neg\phi]_{\mathcal{M}}$$
$$[\Diamond\phi]_{\mathcal{M}} = [\neg\Box\neg\phi]_{\mathcal{M}}$$

Fact (Impossibility Distribution Over $\vee$)**.**

$$\neg\Diamond(p \vee q) \vDash_{\mathfrak{M}} \neg\Diamond p$$
$$\neg\Diamond(p \vee q) \vDash_{\mathfrak{M}} \neg\Diamond q$$

Fact (Unnecessity Distribution Over $\vee$)**.**

$$\neg\Box(p \vee q) \vDash_{\mathfrak{M}} \neg\Box p$$
$$\neg\Box(p \vee q) \vDash_{\mathfrak{M}} \neg\Box q$$

I am very grateful for discussion, comments, and advice on previous drafts of this paper to Jessica Collins, Melissa Fusco, Ben Holguín, Ezra Keshet, Tamar Lando, Karen Lewis, Janum Sethi, Achille Varzi, the participants of the 2021 Eastern APA session for this paper, the participants of the Dianoia Institute of Philosophy Language Workshop, and two anonymous referees for this journal.

References

Abreu Zavaleta, M. (2019). Weak speech reports. *Philosophical Studies 176*(8), 2139–2166.

Aher, M. (2012). Free choice in deontic inquisitive semantics (dis). In M. Aloni, V. Kimmelman, F. Roelofsen, G. W. Sassoon, K. Schulz, and M. Westera (Eds.), *Logic, Language and Meaning*, Berlin, Heidelberg, pp. 22–31. Springer Berlin Heidelberg.

Aloni, M. (2007). Free choice, modals, and imperatives. *Natural Language Semantics 15*(1), 65–94.

Aloni, M. (2018). Fc disjunction in state-based semantics.

Aloni, M. and I. Ciardelli (2013). A logical account of free choice imperatives. In M. F. M. Aloni and F. Roelofsen (Eds.), *The Dynamic, Inquisitive, and Visionary Life of ϕ, ?ϕ , and $\Diamond\phi$: a Festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman*, pp. 1–17.

Alonso-Ovalle, L. (2006). *Disjunction in Alternative Semantics*. PhD Dissertation, University of Massachusetts Amherst.

Barker, C. (2010). Free choice permission as resource-sensitive reasoning. *Semantics and Pragmatics 3*(10), 1–38.

Blumberg, K. and B. Holguín (2019). Embedded attitudes. *Journal of Semantics 36*(3), 377–406.

Cariani, F. (2013). *Ought* and resolution semantics. *Noûs 47*(3), 534–558.

Chemla, E. (2009). Universal implicatures and free choice effects: experimental data. *Semantics and Pragmatics 2*, 1–33.

Chemla, E. and L. Bott (2014). Processing inferences at the semantics/pragmatics frontier: disjunctions and free choice. *Cognition 130*(3), 380–396.

Ciardelli, I., J. Groenendijk, and F. Roelofsen (2018). *Inquisitive Semantics*. Oxford University Press.

Ciardelli, I. and F. Roelofsen (2011). Inquisitive logic. *Journal of Philosophical Logic 40*(1), 55–94.

Ciardelli, I. and F. Roelofsen (2017). Hurford's constraint, the semantics of disjunction, and the nature of alternatives. *Natural Language Semantics 25*(3), 199–222.

Ciardelli, I., F. Roelofsen, and N. Theiler (2017). Composing alternatives. *Linguistics and Philosophy 40*(1), 1–36.

Ciardelli, I., L. Zhang, and L. Champollion (2018). Two switches in the theory of counterfactuals: A study of truth conditionality and minimal change. *Linguistics and Philosophy* (6).

Fara, D. G. (2013). Specifying desires. *Noûs 47*(2), 250–272.

Fine, K. (2017a). A theory of truthmaker content i: Conjunction, disjunction and negation. *Journal of Philosophical Logic 46*(6), 625–674.

Fine, K. (2017b). A theory of truthmaker content ii: Subject-matter, common content, remainder and ground. *Journal of Philosophical Logic 46*(6), 675–702.

Fox, D. (2007). Free choice and the theory of scalar implicatures* mit,.

Fusco, M. (2015). Deontic modality and the semantics of choice. *Philosophers' Imprint 15*.

Gazdar, G. (1979). *Pragmatics: Implicature, presupposition, and logical form*. Academic Press.

Goldstein, S. (2019). Free choice and homogeneity. *Semantics and Pragmatics 12*, 1–48.

Gotzner, N., J. Romoli, and P. Santorio (2020). Choice and prohibition in non-monotonic contexts. *Natural Language Semantics 28*(2), 141–174.

Groenendijk, J. and F. Roelofsen (2010). Radical inquisitive semantics.

Hurford, J. R. (1974). Exclusive or inclusive disjunction. *Foundations of Language 11*(3), 409–411.

Kamp, H. (1973). Free choice permission. *Proceedings of the Aristotelian Society 74*(1), 57–74.

Katzir, R. and R. Singh (2013). Hurford disjunctions: embedded exhaustification and structural economy. *Sinn und Bedeutung 18*, 201–216.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy 1*(3), 337–355.

Kratzer, A. (1991). Modality. In A. von Stechow & Dieter Wunderlich (Ed.), *Semantics: An International Handbook of Contemporary Research*, pp. 639–650. Berlin: de Gruyter.

Kratzer, A. (2012a). Conditionals. In *Modals and Conditionals*, pp. 86–108. Oxford University Press.

Kratzer, A. (2012b). The notional category of modality. In *Modals and Conditionals*, pp. 27–69. Oxford University Press.

Kratzer, A. and J. Shimoyama (2002). Indeterminate pronouns: The view from japanese. *Proceedings of the Third Tokyo Conference on Psycholinguistics*.

Kripke, S. A. (1963). Semantical considerations on modal logic. *Acta Philosophica Fennica 16*(1963), 83–94.

Križ, M. (2015). *Aspects of Homogeneity in the Semantics of Natural Language*. PhD Thesis, University of Vienna.

Križ, M. (2016). Homogeneity, non-maximality, andall. *Journal of Semantics 33*(3), 493–539.

Marty, P., J. Romoli, Y. Sudo, and R. Breheny (ms.). Negative free choice.

Menéndez Benito, P. (2005). *The Grammar of Choice*. Ph.D. thesis, University of Massachusetts, Amherst.

Menéndez Benito, P. (2010). On universal free choice items. *Natural Language Semantics 18*(1), 33–64.

Meyer, M.-C. (2013). *Ignorance and Grammar*. PhD Thesis, MIT.

Meyer, M.-C. (2014). Deriving hurford's constraint. *Proceedings of SALT 24*, 577–596.

Nickel, B. (2010). Generically free choice. *Linguistics and Philosophy 33*(6), 479–512.

Nygren, K. (2019). Supercover semantics for deontic action logic. *Journal of Logic, Language and Information 28*(3), 427–458.

Partee, B. and M. Rooth (1983). Generalized conjunction and type ambiguity. In C. S. Rainer Bauerle and A. von Stechow (Eds.), *Meaning, Use, and Interpretation of Language*, pp. 361–383.

Portner, P. (2009). *Modality*. Oxford University Press.

Roelofsen, F. (2013). Algebraic foundations for the semantic treatment of inquisitive content. *Synthese 190*(S1), 1–24.

Romoli, J. and P. Santorio (2017). Probability and implicatures: A unified account of the scalar effects of disjunction under modals. *Semantics and Pragmatics 10*(3), 1–61.

Romoli, J. and P. Santorio (2019). Filtering free choice. *Semantics and Pragmatics 12*(12), 1–27.

Ross, A. (1941). Imperatives and logic. *Philosophy of Science 11*(1), 30–46.

Sayre-McCord, G. (1986). Deontic logic and the priority of moral theory. *Noûs 20*(2), 179–197.

Schwarzschild, R. (1993). Plurals, presuppositions and the sources of distributivity. *Natural Language Semantics 2*(3), 201–248.

Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy 24*(5), 597–619.

Simons, M. (2005). Dividing things up: The semantics of or and the modal/or interaction. *Natural Language Semantics 13*(3), 271–316.

Starr, W. B. (2016). Expressing permission. *Semantics and Linguistic Theory 26*, 325–349.

van Ditmarsch, H., W. van Der Hoek, and B. Kooi (2008). *Dynamic Epistemic Logic*. Springer.

von Fintel, K. (2012). The best we can (expect to) get? challenges to the classic semantics for deontic modals.

von Wright, G. H. (1968). *An Essay in Deontic Logic and the General Theory of Action*. Amsterdam: North-Holland Pub. Co.

Wedgwood, R. (2006). The meaning of 'ought'. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics: Volume 1*, pp. 127–160. Clarendon Press.

Willer, M. (2018). Simplifying with free choice. *Topoi 37*(3), 379–392.

Willer, M. (2021). Two puzzles about ability can. *Linguistics and Philosophy 44*(3), 551–586.

Yablo, S. (2014). *Aboutness*. Princeton University Press.

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

Zimmermann, T. E. (2006). Monotonicity in opaque verbs. *Linguistics and Philosophy 29*(6), 715–761.