

Deontic Modality and the Semantics of Choice

Melissa Fusco

August 26, 2015

Abstract

I propose a unified solution to two puzzles: Ross’s puzzle (the apparent failure of $\text{OUGHT } \phi$ to entail $\text{OUGHT}(\phi \text{ OR } \psi)$) and free choice permission (the apparent fact that $\text{MAY}(\phi \text{ OR } \psi)$ entails both $\text{MAY } \phi$ and $\text{MAY } \psi$). I begin with a pair of cases from the decision theory literature illustrating the phenomenon of *act dependence*, where what an agent ought to do depends on what she does. The notion of permissibility distilled from these cases forms the basis for my analysis of MAY and OUGHT . This framework is then combined with a generalization of the classical semantics for disjunction—equivalent to Boolean disjunction on the diagonal, but with a different two-dimensional character—that explains the puzzling facts in terms of semantic consequence.

1 Two Puzzles

Suppose I say to you

- (1) You may have the gin or the whiskey.
 $\text{MAY}(G \text{ OR } W)$

As you help yourself to the latter, I cry, “Stop! You can’t have whiskey!”¹

It seems that I have contradicted myself. For how, given (1), could your action have been impermissible? That we hear sentences like (1) as communicating that you may have the gin, and you may have the whiskey, is the puzzle, or paradox, of *free choice permission* (von Wright, 1969; Kamp, 1973).

Free choice permission is usually glossed as a felt entailment from a narrow-scope disjunction under MAY to a wide scope conjunction:

$$\text{(FC)} \quad \text{MAY}(\phi \text{ OR } \psi) \Rightarrow (\text{MAY } \phi) \wedge (\text{MAY } \psi)$$

More can be added to this characterization. Most speakers have the strong intuition that while the permission in (1) communicates that one may choose the gin and one may choose the whiskey, it emphatically fails to communicate the permissibility of the corresponding narrow-scope conjunction, which would allow you to have both:

¹I use double-quotes to mark assertions (or other speech acts). For readability, and where it causes no confusion, I will be loose with use and mention.

- (2) You may have the gin *and* the whiskey.
 $\text{MAY}(G \wedge W)$

... and perhaps even entails that (2) is false. This feature goes by the name *exclusivity*:

$$\text{(Exclusivity)} \quad \text{MAY}(\phi \text{ OR } \psi) \not\Rightarrow \text{MAY}(\phi \wedge \psi)$$

This observation is embedded in the scholarship on free choice permission.² The tension between (1) and (2) is usually glossed as a scalar implicature.³

Here is a second puzzle. Consider a similar situation, in which I say to you

- (3) You ought to post the letter.
 $\text{OUGHT } P$

You reason as follows.

- (4) I ought to post the letter.
 $\text{OUGHT } P$

Therefore,

- (5) I ought to post the letter or burn it.⁴
 $\text{OUGHT}(P \text{ OR } B)$

Again, it appears that something has gone wrong. For (4) does not seem to entail (5). That this entailment is classically valid, but feels intuitively *invalid*, is Ross's Puzzle (Ross, 1941):

$$\text{(R)} \quad \text{OUGHT}(\phi) \not\Rightarrow \text{OUGHT}(\phi \text{ OR } \psi)$$

A straightforward way of trying to explain what is wrong with the inference from (4) to (5) is that the introduced disjunct (here, *burning the letter*) is impermissible. Recent work on Ross's puzzle (Cariani, 2011; Lassiter, 2011b,a) has pursued this intuitive route, investigating various ways of working out the thought that a Ross sentence $\text{OUGHT}(\phi \text{ OR}$

²Barker (2010) writes that FC sentences like (1) "never" guarantee conjunctions like (2). The failure of (1) to entail (2) is also assumed by Simons (2005), who calls it a "consensus in the literature." Danny Fox, an implicature theorist, makes his psychological explanation of free choice permission *dependent* on a hearer's rejection of (2) (Fox, 2007). Fox's official position thus includes an endorsement of the stronger inference we may call *and-false*:

$$\text{(And-False)} \quad \text{MAY}(\phi \text{ OR } \psi) \Rightarrow \neg \text{MAY}(\phi \wedge \psi)$$

However, Fox expresses some reservations about this inference—and thus the fact that his explanation of free choice permission relies on it—in the last section of his paper. See Fox (2007, pg. 35-36).

³Both (FC) and (Exclusivity) have been analyzed as scalar implicatures, the former of an unusual, post-Gricean kind, in the wake of Kratzer & Shimoyama 2002 (see especially Fox (2007) and Chierchia (2006)). I present an argument analyses of (FC) in the Kratzer and Shimoyama vein in Fusco (2014).

⁴This example is Ross's original.

ψ) entails that there must be something to be said, deontically speaking, for each of the embedded disjuncts.

$$(R+) \text{ OUGHT}(\phi \text{ OR } \psi) \Rightarrow (\text{MAY } \phi) \wedge (\text{MAY } \psi)$$

This looks like a unification with free choice permission: disjunction under both OUGHT and MAY carry what we can call *an entailment to disjunct permissibility*.⁵

A simple variation on Ross's puzzle, due to Geoffrey Sayre-McCord, suggests that there is yet more to say about Ross sentences. Sayre-McCord observes that, according to the hypothesis that disjunction introduction is blocked in the scope of OUGHT *just in case* the introduced disjunct is impermissible, it would follow that, if ψ is known to be permissible, then

$$\text{OUGHT } \phi$$

should entail

$$\text{OUGHT}(\phi \text{ OR } \psi)$$

in context.

But, Sayre-McCord argues, "OUGHT(ϕ OR ψ) cannot legitimately be inferred from OUGHT ϕ even when ψ is perfectly permissible" (Sayre-McCord, 1986, pg. 189). His example is as follows: suppose it is taken for granted that it is perfectly fine for Ralph to go to the movies;

- (6) Ralph may go to the movies.
MAY M

is true. It is also true that

- (7) Ralph ought to pay back his loan.
OUGHT L

Still,

- (8) Ralph ought to pay back his loan or go to the movies.
OUGHT(M OR L)

sounds wrong; it does not seem to follow from (6) and (7). Call this

$$(SM) \text{ MAY}(\psi) \wedge \text{OUGHT}(\phi) \not\Rightarrow \text{OUGHT}(\phi \text{ OR } \psi)$$

⁵This apparent connection between free choice and Ross is noted by, amongst others, von Fintel (2012), Cariani (2011, pg. 20-21), and von Wright himself (von Wright, 1969, pg. 22).

Sometimes, disjunction introduction in the scope of a deontic modal fails *even when* the introduced disjunct is permissible.

Sayre-McCord's observation seems to show that there is no paraphrase for the 'or' in a Ross sentence in terms of the *unconditional* deontic status of the disjuncts expressible in the object language—that is, in terms of the deontic status each disjunct has, considered independently of the other one.

$$\begin{aligned}
 \text{OUGHT}(\phi \text{ OR } \psi) &\not\equiv \text{OUGHT } \phi \wedge \text{OUGHT } \psi \\
 &\text{both obligatory—too strong} \\
 &\not\equiv \text{MAY } \phi \wedge \text{MAY } \psi \\
 &\text{both (merely) permissible—too weak} \\
 &\not\equiv (\text{OUGHT } \phi \wedge \text{MAY } \psi) \vee (\text{OUGHT } \psi \wedge \text{MAY } \phi) \\
 &\text{one of each—(SM)}
 \end{aligned}$$

These examples suggest we should switch tracks, and try to directly characterize situations that *are* appropriately described by sentences of the form $\text{OUGHT}(\phi \text{ OR } \psi)$. Suppose that I have come down with a cold this morning and am too sick to host my housewarming party tomorrow night. It seems appropriate to say:

(9) I ought to cancel or postpone the party.

If (9) is true, then canceling the party has a certain deontic status (that of being obligatory), *provided that I do not postpone*, and postponing the party has a certain deontic status (that of being obligatory), *provided that I do not cancel*. That nothing at all is said by (9) about the situation where I both postpone and cancel is shown by the fact that (9) can be true, even though the conjunction, *postpone and cancel*, has no positive normative status at all: it would be rude to my guests to postpone the party only to cancel it later, and it's impossible to cancel a party and *then* to postpone it.

I propose, then, to add to the data associated with Ross sentences like (9): the deontic status of the act described by each disjunct depends on whether the other one is performed. Starting with a Ross sentence as a premise, an agent can reason with future-directed conditionals like this:

(Conditionals-O)

I ought to do $(\phi \text{ OR } \psi) \Rightarrow$ If I do not do ϕ , I ought to do ψ ;

I ought to do $(\phi \text{ OR } \psi) \Rightarrow$ If I do not do ψ , I ought to do ϕ .

In addition being a promising paraphrase of (9), (Conditionals-O) helps us to precisify our discomfort with Sayre-McCord's defective (8), where disjunction introduction in the scope of OUGHT is blocked, even though the introduced disjunct is permissible. Going to the movies is strictly optional—it's just *false* that movie-going achieves the status of an obligation in cases where Ralph fails to pay back his loan. So, according to

(Conditionals-O), (8) is in one respect too strong. It is also, in another respect, too weak, which we can see from looking at the second conditional licensed by the schema. In the situation Sayre-McCord describes, paying back the loan is not something Ralph ought to do *on the condition that he doesn't go to the movies*; it's just something he ought to do.

Another nice feature of (Conditionals-O) is that it holds out the promise of a unification with free choice permission—a better one, since it is more empirically adequate from the OUGHT-side. For consider the property that is formally parallel to (Conditionals-O), with MAY substituted for OUGHT:

(Conditionals-M)

I may do $(\phi \text{ OR } \psi) \Rightarrow$ If I do not do ϕ , I may do ψ ;

I may do $(\phi \text{ OR } \psi) \Rightarrow$ If I do not do ψ , I may do ϕ .

(Conditionals-M) is just a re-framing of free choice permission that incorporates the proviso of (Exclusivity). If I give you permission to have the gin or the whiskey, then having the gin has a certain deontic status (that of being permissible), *provided that you do not also take the whiskey*. And vice-versa. As in the OUGHT case, it is certainly *compatible* with what I said that having both gin and whiskey is permissible. But the truth of the free choice premise, $\text{MAY}(\phi \text{ OR } \psi)$, does not require this; $\text{MAY}(\phi \text{ OR } \psi)$ may be true even when doing ϕ and ψ is not permitted at all.

I won't argue further for (Conditionals-M) here. But I think the intuitive appeal of (Conditionals-M)—especially in light of its formal similarity to (Conditionals-O)—is clear enough to make the two inferences a basis for an exploration of the semantic behavior of 'or' in these modal environments.⁶ Table 1 gives an interim summary of the intuitions that constitute our data.

May:

(Failure of 'or' intro)	$\text{MAY } \phi \not\Rightarrow \text{MAY}(\phi \text{ OR } \psi)$
(FC)	$\text{MAY}(\phi \text{ OR } \psi) \Rightarrow \text{MAY } \phi \wedge \text{MAY } \psi$
(Conditionals-M)	$\text{MAY}(\phi \text{ OR } \psi) \Rightarrow$ if $\neg\phi$, then $\text{MAY } \psi$

Ought:

(Failure of 'or' intro)	$\text{OUGHT } \phi \not\Rightarrow \text{OUGHT}(\phi \text{ OR } \psi)$
(R+)	$\text{OUGHT}(\phi \text{ OR } \psi) \Rightarrow \text{MAY } \phi \wedge \text{MAY } \psi$
(Conditionals-O)	$\text{OUGHT}(\phi \text{ OR } \psi) \Rightarrow$ if $\neg\phi$, then $\text{OUGHT } \psi$

Table 1: Data for OUGHT and MAY.

(Conditionals-M) and (Conditionals-O) suggest that the deontic status of the acts described by the disjuncts depends on what else the agent does. If so, they are instances of

⁶On this point, it is worth noting that (Conditionals-M) follows from (FC) on Kratzer (1981, 1991b)'s semantics for the deontic modals and the indicative conditional, a package considered orthodox in the current literature on the basis of a wealth of independent data.

act dependence: the normative status of each disjunct is future-contingent—in particular, contingent on what the agent decides vis-à-vis the other disjunct. Act dependence seems to be part of the data of free choice permission and Ross’s puzzle.

In this paper, I present an account on which act dependence is also one half of a two-part semantic explanation of the data. I begin by choosing a model theory adequate to modeling act dependence, where obligation depends on what is chosen. I introduce concepts of obligation and permissibility in this framework inspired by decision theoretic work on the concept of ratifiability (Jeffrey, 1983), and impose on these models a two-dimensional semantic framework familiar from Davies & Humberstone (1980). In §4, I introduce an entry for disjunction, which is equivalent to Boolean disjunction outside the scope of modal operators, but has a different two-dimensional character. I then use these ingredients to validate the patterns in Table 1.

2 Proto-Semantics with Kripke Frames

In order to model cases where what an agent ought to do depends on what she chooses to do, we will need models capable of representing (i) multiple candidates for how the agent will act, and (ii) multiple candidates for what she ought to do.

Let an *agentive Kripke frame* be a tuple $\mathcal{M} = \langle W, R \rangle$ consisting of a universe of worlds W and a binary accessibility relation R on those worlds. In our frames, two worlds w and w' in W are distinct just in case some choice the agent can make is different in each of them: a world is a maximally decided course of choices.⁷ A *modal base* $s \subseteq W$ (which I will sometimes call a *choice situation*) is a less specific possibility, leaving some future choices undecided. It represents the position from which the agent chooses.

In the service of (i), we assume that the agent represents all actions within the range of her practical abilities as contingent with respect to s . A simple sentence p is *circumstantially possible at s* if it is true at some world(s) in s ,⁸ and *settled-true at s* if it is true at all worlds in s .

Definition 1 (Circumstantial Possibility). p is *circumstantially possible at s* iff there is some world $v \in s$ such that p is true at v .

Definition 2 (Settled-truth). p is *settled-true at s* iff for all worlds $w' \in s$, p is true at w' .

The notion of settled-truth in Definition 2 is our approximation of truth at a situation of choice. For example, it is settled-true at some situation s that I will lose the game I am playing if and only if it is true at every world compatible with my practical abilities in that situation—true *no matter what I choose*—that I will lose the game. Conversely, it is

⁷Our worlds are reminiscent of elements of choice partitions familiar from the work of Belnap, Perloff, and Xu (2001). Note, however, that my gloss on worlds as maximal courses of action idealizes away from the fact (emphasized by Belnap et. al) that we intuitively are able to choose between progressively finer options as time passes. In this respect, our worlds $w \in W$ might be better seen as cells of the *minimal partition refinement* of all the choice partitions an agent faces in one of Belnap et. al’s models.

⁸I take the term ‘circumstantial’ from the influential discussion in Kratzer (1981, §5).

circumstantially possible that I will *win* the game if there is some world $w \in s$ where it is true that I win the game.⁹

The R relation is a *deontic accessibility relation*: two worlds w and w' are R -related just in case w' is deontically ideal by the lights of w .¹⁰ Following standard deontic logic ('SDL'¹¹), we gloss world-centric obligation and its dual, permissibility, in terms of the R relation:

Definition 3 (Obligation and Permissibility at Worlds).

(3-a) p is **obligatory at** w iff for any world v such that wRv , p is true at v .

(3-b) p is **permissible at** w iff there is some world v such that wRv and p is true at v .

Following Kratzer (1981) and SDL, I assume the R relation is realistic and serial in s :

Definition 4 (Realistic). R is **realistic in** s iff for any w and v , if $w \in s$ and wRv , then $v \in s$.

Definition 5 (Serial). R is **serial in** s iff for all $w \in s$, there is some v such that wRv .

Definitions 4 and 5 describe what it is for a world to see another world as ideal, when s is taken to circumscribe the range of practically available options: Definition 4 says a world in s sees another world only if that other world is also circumstantially possible,¹² and Definition 5 says that no world is nihilistic, seeing nothing as ideal.

In service of goal (ii) above, our models allow that two worlds w and v , both in s , may be R -related to different outcomes: there can be a genuine variety of perspectives, amongst future-contingent states, regarding what is deontically ideal. To illustrate, here is a case from the decision-theory literature.

Nice Choices at the Spa. Aromatherapy [= p] or body-wrap [= q]—which is it to be? You believe that, whichever you choose, you will be very glad you chose it. Mid-aromatherapy, the aromatherapy will seem self-evidently superior [to the body-wrap]. Mid-body-wrap, the body-wrap will seem self-evidently superior [to the aromatherapy]. (Hare & Hedden, 2015, pg. 3)

Figure 1 shows a Kripke model for a case like (Nice Choices), with arrows to represent the R relation.

⁹ Our proto-semantics for atomics is therefore *trivalent* (or *gappy*); it is similar to the supervaluationist semantics for branching time in Thomason (1970), in which future contingents may be neither (settled- nor super-) true nor (settled- nor super-) false.

¹⁰The thought that deontic relations hold between *choosable* points (rather than anything finer-grained) is suggested by MacFarlane & Kolodny (2010)'s attractively named principle *Ought Implies Can Choose* (pg 132, footnote 28). See also MacFarlane (2013, §11.4) and Charlow (2013).

¹¹See, for example, the introduction in McNamara (2010).

¹²Suppose, for example, that it is true at every world in s that Bill gets mugged; it is not circumstantially possible to prevent the mugging. Then worlds in s will 'see' worlds where Bill is aided as ideal—even though it would have been better, relative to some enlarged s^+ , if Bill had never been mugged at all.




Figure 1: A situation s_{nice} representing (Nice Choices).

In (Nice Choices), what you ought to do depends on what you *choose* to do. There is a sense in which you can't go wrong: if you pick aromatherapy, you'll find yourself at w_1 , relative to which aromatherapy is the unique ideal option (since $w_1 R w_1$ and not $w_1 R w_2$.) So you'll be satisfied with what you did. But if you pick the body-wrap, you'll find yourself at w_2 , relative to which the body-wrap is uniquely ideal (since $w_2 R w_2$ and not $w_2 R w_1$), so again, you'll be satisfied with what you did. And you know all this before you choose.

Notice that, even though there is something nice in (Nice Choices) about both p and q , p and q are *deontic contraries*—no post-choice perspective that sees p as ideal sees q as ideal, and vice-versa. So there is a sense in which, if you choose (say) to bring about p , the “niceness” of q immediately evaporates: looking back, you *won't* see q as ideal. This is the difference between (Nice Choices) and a model that simply represents p and q as *equally ideal*, from every point of view.

Definition 3 uses the R relation to tell us what is permissible and obligatory at each individual world in s_{nice} . But we want to know what is permissible and obligatory from the *global* perspective the agent occupies *before* she chooses—what is obligatory and permissible relative to the whole of s . It is s , after all, that represents the point of view *from which* she makes her choice.

A simple answer to both questions would be that each normative notion, obligation and permissibility, scales up to s by imposing the conditions in Definition 3 on each individual world in s , just as we said that a fact-describing statement like “I will lose the game” is settled-true at s when it is true at each individual world in s . Call these null hypotheses Postulates 1 and 2:

Postulate 1. p is *obligatory at* s iff p is obligatory at every world in s .

Postulate 2. p is *permissible at* s iff p is permissible at every world in s .

Postulate 1 looks right to describe the obligation-facts in (Nice Choices). In (Nice Choices), neither p nor q is obligatory at every world, so neither is obligatory relative to the whole modal base. Given that an agent has the freedom to choose between worlds, it does not seem that anything *less* than p 's being obligatory at *every* world in s could be sufficient for p 's being obligatory at s : if there is even *one* world in s where p is *not* obligatory, the agent could choose that world, and thereby “escape” the obligation.

By contrast, Postulate 2 is less secure. The instinct that you “can't go wrong” in (Nice Choices) remains unexplained on Postulate 2, since it says that (for example) p is permis-

sible at s only if p is permissible by the lights of *every* world in s . Hence by Postulate 2, p is not permissible at s_{nice} , since it is not permissible at w_2 . Likewise, q is not permissible at s_{nice} , since it is not permissible at w_1 . Whereas instinct holds that *both* of p and q are choiceworthy in (Nice Choices), Postulate 2 tells us that *neither* is. Scaling up the intuitive notion of permissibility from w to s by the null hypothesis thus misses something about the case.

While I do not think Postulate 2 is the correct notion of permissibility at s , it is useful to have a name for the condition it attempts to impose on modal bases, for I think a better notion can be defined in terms of it. Let us call the notion associated with Postulate 2 *admissibility*:

Definition 6. p is *admissible at* s iff p is permissible at every world in s .

While it is a substantive hypothesis—Postulate 2—that p is permissible at s just in case p is permissible at every individual world in s , it is true by definition that p is *admissible* at s just in case this condition holds.

We can also consider admissibility at subsets s' in s , representing post-choice contexts:

Definition 7. p is *admissible at nonempty* s' in s iff p is permissible at $\langle s, w \rangle$ for every w in s' .

The admissibility of p at an s' in s is *persistent* permissibility—permissibility that is inescapable by the lights of one's future choices.¹³

(Nice Choices) highlights cases of preestablished harmony between whether an act is performed and whether that act is *post-choice admissible*. In act-dependent frames, cases of *disharmony* between whether an act is performed and whether it has this status are also possible. Consider:

Death. You live in Damascus and learn that Death is coming to collect your soul. Death always follows his predetermined schedule and Death never misses his quarry. Should you flee to Aleppo [= p]? You are confident that, if you flee to Aleppo, Death will be there. But if you stay in Damascus [= q], Death will be there too. (Gibbard & Harper, 1978)

Nasty Choices at the Spa. Abdominal-acupuncture [= p] or bee-sting-therapy [= q]*—*which is it to be? Whichever you choose, you will wish that you had chosen the other. (Hare & Hedden, 2015, pg. 13)

In (Nasty Choices), what the agent ought to do once again depends on what she chooses to do. But now her freedom has become a curse: even though every world in s

¹³ Once p is admissible at s' in s , then any $s'' \subseteq s'$ will be such that p is admissible at s'' in s . The retained relativity to s is for a technical reason: $\langle s', w \rangle$ may fail to be a well-defined point of evaluation if R fails to be serial or realistic (Definitions 4 and 5) in the contracted state.




Figure 2: A situation s_{nasty} representing (Nasty Choices).

sees *some* option as ideal, no option sees *itself* as ideal. There is no way the agent can make a choice and be satisfied by the deontic perspective she will occupy *after* she chooses.

Like the predicament in (Nice Choices), the pickle the agent is in in (Nasty Choices) is possible because we are considering what is choiceworthy at a relatively indeterminate state. In typical ‘pointed’ applications of Kripke frames—where permissibility claims are evaluated at fully determinate worlds—the seriality of R in s (Definition 3) is sufficient to guarantee that there is always some intuitively permissible option available, since the seriality of R guarantees that from the point of view of any *world*, something always counts as ideal. It is only from the s -centric perspective that the seriality of R is no longer enough to capture this intuitive thought, since it is from the s -centric perspective that preestablished disharmonies like (Nasty Choices) are possible.

This is why cases like (Nice Choices) and (Nasty Choices) are prominent in the literature on rational choice. The thought is that they bring out a novel feature that choiceworthy acts must have: any choice that is ideal for you must be ideal for you *on the assumption that you perform it*. Richard Jeffrey called the property in question *ratifiability*:

A ratifiable decision is a decision to perform an act of maximum estimated desirability relative to the probability matrix the agent thinks he would have if he finally decided to perform that act...*Maxim*. Make ratifiable decisions. To put it romantically: ‘Choose for the person you expect to be when you have chosen’ (Jeffrey, 1983, pg.16).

What is ratifiability in a Kripke frame? We shall say that an option, p , is ratifiable at s if the agent can contract s to some s' such that p is both settled-true and admissible at s' in s : this corresponds to an agent’s being able to conditionalize on performing p in such a way that p is deontically ideal from her post-choice standpoint (ideal to her “future self,” to use Jeffrey’s phrase.)¹⁴

Definition 8 (Ratifiability). p is *ratifiable at* s iff there is some nonempty $s' \subseteq s$ such that:

- (i) p is settled-true at s' in s ;¹⁵

¹⁴Jeffrey’s norm is usually glossed in terms of subjective credences, but a more objective gloss on this talk of “conditionalization” is also available. As the agent acts, she constrains the course of history (whether she knows it or not); we leave open the possibility that the deontic status of an act exhibits objective, *causal* dependence on how history unfolds.

¹⁵That is, p is true at $\langle s, w \rangle$ for each $w \in s'$. See Footnote 13.

(ii) p is admissible at s' in s .

Cashing out admissibility in terms of the R -relation, this condition's holding at s is equivalent to its holding at a single world $v \in s$.¹⁶ So Definition 8 simplifies to:

Definition 8 (Ratifiability, simplified). p is **ratifiable at** s iff there is some world $v \in s$ such that:

- (i) p is true at v ;
- (ii) $\exists v' \in s$ such that vRv' and p is true at v' .

In terms of an agent's ability to navigate between world-relative obligations: p is ratifiable at s just in case it is possible, practically speaking, to choose a p -world that sees a p -world. In place of Postulate 2, I advance

Postulate 3. p is **permissible at** s iff p is ratifiable at s .

Postulate 3 differs from Postulate 2 because the ratifiability of p relative to s is distinct from the admissibility of p relative to s . I claim that it is better at accounting for the pretheoretical notion of *permissibility*, and it is the notion I will use, in the coming sections, to model the semantics of MAY. (Nice Choices) and (Nasty Choices) make the case: in (Nice Choices), neither p nor q is s -admissible, but both p and q are ratifiable. To the extent to which we feel that the preestablished harmony between act and status in (Nice Choices) renders both options *permissible* relative to the agent's undecided state, our intuitions are pegging the notion of permissibility to ratifiability, not to admissibility.

Turning our attention to (Nasty Choices), we see that, at every world in s_{nasty} , some nontrivial proposition (either p , or q) is admissible, according to that world. So relative to s_{nasty} as a whole, it is settled-true that

(10) Some option is *admissible*.

But while (10) is settled-true at s_{nasty} ,

(11) Some option is *ratifiable*.

is settled-false. To the extent to which Nasty cases strike us as hopeless—to the extent to which we feel there is nothing you *may* do in a Nasty case—our intuitions are once again tracking ratifiability, not admissibility.

To endorse the hypothesis that permissibility is ratifiability, relative to an undecided state s , is to implement the hindsight-directed point of view recommended by our cases: to abandon the prospective (and often indeterminate) question, “Is p admissible?” in favor of the retrospective (often more determinate) question, “If I do p , will I *have done* what is admissible?”

¹⁶In the left-to-right direction, just take $s' = \{v\}$. In the right-to-left direction: by nonemptiness of s' , there is at least one world $w \in s$ which satisfies the conjunction of (i) and (ii); let v be this world w .

Turning to the package of s -obligation and s -ratifiability together—Postulates 1 and 3, united—we see an important difference in perspective. By Postulate 1, world-centric obligation is capable of serving as the basic concept of obligation: global conditions on s are derived from it by imposing that concept on each individual world in s . By contrast, Postulate 3 tells us that permissibility can only be fully understood by taking the global perspective as basic: an option a is ratifiable just in case it is *choosably* both true and admissible at s , but whether that conjunction is *choosable* depends irreducibly on the whole of s : it is not a distributive property of individual worlds. Since global features are not built up pointwise out of local features, they are not, in general, *persistent*: the fact that s has some global feature does not entail that arbitrary subsets of s do. So Postulate 3 predicts that permissibility-facts may fail to endure as s becomes more and more determinate—for example, as the agent executes a series of acts. That is a good fit with the data. For example, it is a good fit with the intuition that you may take the gin and you may take the whiskey, but these acts are not permissible *come what may*: the latter permission does not persist if you exercise the former one.

Given this, I introduce a toy language. The operator ‘ O ’, for OUGHT, tracks obligation, as defined according to Postulate 1. The operator ‘ M ’, for MAY, tracks permissibility, as defined according to Postulate 3. For expressive completeness, I also introduce an operator ‘ \Diamond ’ for admissibility, as defined according to Definition 6, and ‘ \blacklozenge ’ for circumstantial possibility, as defined according to Definition 1. Summing up what we have so far:

Propositional fragment:

$$\begin{aligned} a \text{ is true at } \langle s, w \rangle & \text{ iff } w \text{ is an } a\text{-world} \\ \neg p \text{ is true at } \langle s, w \rangle & \text{ iff } p \text{ is not true at } \langle s, w \rangle \\ p \wedge q \text{ is true at } \langle s, w \rangle & \text{ iff } p \text{ is true at } \langle s, w \rangle \text{ and } q \text{ is true at } \langle s, w \rangle \end{aligned}$$

Circumstantial Modality:

$$\blacklozenge p \text{ is true at } \langle s, w \rangle \text{ iff } \exists v \in s: p \text{ is true at } \langle s, v \rangle$$

Deontic Modality:

$$\begin{aligned} (\text{Obligation}) \quad Op \text{ is true at } \langle s, w \rangle & \text{ iff } \forall w' \in s: \text{ if } wRw', \text{ then } p \text{ is true at } \langle s, w' \rangle \\ (\text{Admissibility}) \quad \Diamond p \text{ is true at } \langle s, w \rangle & \text{ iff } \exists w' \in s: wRw' \text{ and } p \text{ is true at } \langle s, w' \rangle \\ (\text{Ratifiability}) \quad Mp \text{ is true at } \langle s, w \rangle & \text{ iff } \exists v \in s \text{ such that (i) } p \text{ is true at } \langle s, v \rangle, \text{ and (ii)} \\ & \quad \Diamond p \text{ is true at } \langle s, v \rangle \end{aligned}$$

Notice that the truth of Mp at $\langle s, w \rangle$ is independent of w , the world-parameter of the point of evaluation. This encodes the global nature of ratifiability—its dependence only on s .


Since we are considering truth at a context to be akin to settled-truth at s , our attendant notion of semantic consequence is the preservation of settled-truth at s . Hence $\phi \vDash \psi$ iff for any situation s , the truth of ϕ at every world in s guarantees the truth of ψ at every world in s .

3 Cases Revisited

Armed with our postulate that MAY tracks ratifiability, we can sketch two patterns of reasoning in act-dependent models. The first involves OUGHT. The second involves only MAY. Suppose that, for live options p and q :


Pattern A:

1. If I make p true, then OUGHT p is true. (Premise)
2. If I make q true, then OUGHT q is true. (Premise)
3. So, if I make p true, p is admissible. (Transitivity of consequence)¹⁷
4. And, if I make q true, q is admissible. (Transitivity of consequence)
5. So, I can make p both true and admissible. (from 3)
6. And, I can make q both true and admissible. (from 4)
7. So, MAY p is true. (5, Ratifiability)
8. And MAY q is also true. (6, Ratifiability)



Pattern B:

1. If I make p true, then MAY p is true. (Premise)
2. If I make q true, then MAY q is true. (Premise)
3. So, if I make p true, I can make p both true and admissible. (1, Ratifiability)
4. And if I make q true, I can make q both true and admissible. (2, Ratifiability)
5. So I can make p both true and admissible. (from 3)¹⁸
6. And I can make q both true and admissible. (from 4)
7. So MAY p is true. (from 5)
8. And MAY q is true. (from 6)



I will argue that Pattern A is the underlying form of the inference that gives rise to

¹⁷Since Op entails $\Diamond p$ in s whenever R is realistic and serial in s (Definitions 4 and 5).

¹⁸There is a state where p is done admissibly in s iff there is a state where p is done admissibly in the region of s throughout which p is true, so the truth of (3) is sufficient to guarantee the truth of (5); the transition is analogous to the inference from “Amongst the F ’s, some F ’s are G ” to “Some F ’s are G ”, when F is nonempty.

(R+)—the strengthened version of Ross’s puzzle—and that Pattern B is the underlying form of the inference that gives rise to free choice permission. The missing ingredient is the semantics of sentential disjunction, which takes us from $\text{OUGHT}(p \text{ OR } q)$ to Premises 1-2 of Pattern A—as illustrated in the branching diagram on the left—and from $\text{MAY}(p \text{ OR } q)$ to Premises 1-2 of Pattern B—as illustrated in the branching diagram on the right. The logic of ratifiability will take care of the rest.

Before I plunge ahead, let me say a little in defense of this strategy. If disjunction can bridge the gap we’ve framed, we would have a story about how free choice permission sentences and Ross sentences could impose conditions on s that approximate the structure of Nice Cases—the structure, that is, of preestablished harmony between whether an act is performed and whether it has some positive normative status. Since Nice Cases exemplify the kind of impersistent permissibility we seem to get as *outputs*, when free choice permission sentences and Ross sentences are taken as *inputs*, this seems like a promising route to pursue. Moreover, the approach seems not unfeasible, since both patterns characterize the contribution of the premises’ common factor—disjunction—in the same way.

To go forward, we need to get from (e.g.)


$$\text{OUGHT}(p \text{ OR } q)$$

being settled-true at s , to the validity of these transitions:

$$\frac{p}{\text{OUGHT } p} \qquad \frac{q}{\text{OUGHT } q}$$

That means we need to set up a dependency between which of p or q the agent actually brings about, and the semantic contribution of the sentence embedded under the modal. We must achieve this despite the fact that the modal operator takes a proposition—a *condition* on possible worlds—as an argument. To set up a dependence between what is true in the actual state and the *propositional content* of a sentence like $(p \text{ OR } q)$, we will require a two-dimensional semantics, in the style of Davies & Humberstone (1980). Our full valuation function will therefore recursively define truth at points of evaluation which are *triples* $\langle s, y, x \rangle$ consisting of a modal base s and a *pair* of worlds y and x in s .

Two-dimensional semantics is motivated by the thought that the pretheoretical notion of sentential truth can be upgraded to a relation in three parts: what it is for a sentence ϕ to be true at w is for the proposition ϕ expresses *at* w to be true *in* w . We can also ask whether that very same proposition might have been true in some other world, v : in asking this question, we hold w fixed in its role as the *world-as-actual*, which contributes to the proposition expressed by ϕ ; but we let v play the role of the *world of evaluation*: the world *in which* that proposition’s truth is evaluated.



In the case of disjunction, we want the proposition expressed by a disjunction to depend on which disjunct is true in w . We can then ask whether *that very disjunct* is true at some other, more ideal world v .

The connection we are trying to make can be put like this. In the previous section, we examined act dependence—the idea that the deontic status of an act may vary across different worlds, between which the agent can choose. The move to two-dimensional semantics endeavors to mirror this state of affairs by implementing *semantic* act dependence—the idea that a *sentence* may express different *propositions* in different worlds, amongst which the agent can choose. Concretely, when considering an actual world w where only p is true, we want ‘ p OR q ’ to be equivalent to p ; and in an actual world w where only q is true, we want ‘ p OR q ’ to be equivalent to q .

I sketch the required treatment of disjunction in the following section. We will then return to the protosemantics of §2 and upgrade it to two dimensions, bringing the rest of the system along for the ride.

4 Disjuncts as Truthmakers

The semantics for disjunction I lay out in this section develops the idea that any true disjunction is *witnessed* by one of its disjuncts, in the same way that a true existential statement in first-order logic is witnessed by an individual. We will hold, further, that the *interpretation* of a disjunction ‘ p OR q ’ depends on the way in which the disjunction is witnessed. The resulting view parallels, in a “static” framework, the dynamic semantics for indefinite noun phrases pioneered by Irene Heim (Heim, 1982).¹⁹

We begin by identifying by a disjunction’s true disjuncts. Let Alt_w (where $w \in W$) be a function that takes two sentences as arguments, returning the sentences from amongst those two that are true in w , whenever there are any. For completeness, I add the condition that Alt_w returns both disjuncts in the case where both are false—the intuition being that, since there are no truthmakers for a false disjunction, there is nothing to break the symmetry between the two inputs.

Definition 9 (The Alt_w function). Alt_w is a function that takes a pair of sentences ϕ and ψ as arguments. It returns the set containing all and only the true-in- w sentences in

¹⁹Heim’s work concerns so-called “donkey anaphora,” whereby sentences like

- (i) If MacGregor owns a donkey, he beats it.

give rise to truth-conditions like

- (ii) $\forall x(\text{MacGregor owns } x \rightarrow \text{MacGregor beats } x)$.

The intuitive connection between the projects is this. In addition to being truth-apt, the antecedent of (i), *MacGregor owns a donkey*, can be made true in different ways: that is, by different donkeys. For each donkey that makes the antecedent true, the semantics generates a conclusion about that very donkey (that MacGregor beats *it*.) Hence, the *interpretation* of the pronoun ‘it’ must depend on the way in which the antecedent is made true. This requires a finer semantic grain than merely whether the antecedent is true—we must track *how* it is made true. Further inspirations and applications of the view, including connections to disjunctive questions, are explored in Fusco (2015).

$\{\phi, \psi\}$, if there are any, and returns $\{\phi, \psi\}$ otherwise.

Here is Alt_w for the four world-types that correspond to the standard truth table for p and q :²⁰

	p	q	$Alt_w(p, q)$
w_1	T	T	$\{p, q\}$
w_2	T	F	$\{p\}$
w_3	F	T	$\{q\}$
w_4	F	F	$\{p, q\}$

Notice that the truth-at- w of one member of $Alt_w(p, q)$ is both necessary and sufficient for the truth of ‘ p OR q ’ at w . So to recover the classical truth-conditions of disjunction, it suffices to quantify existentially over $Alt_w(p, q)$.

Definition 10 (Protosemantics for ‘or’). ‘ p OR q ’ is true at $\langle s, w \rangle$ iff there is some $\alpha \in Alt_w(p, q)$ such that α is true at $\langle s, w \rangle$.

Given our definition of Alt_w , Definition 10 is equivalent to the classical truth-conditions of unembedded disjunction. However, because w appears on both sides of the “is true at” on the right hand side, this restatement opens up the possibility of seeing w as contributing to the proposition expressed by the disjunction. The Alt_w function allows its parameterizing world, w , to associate a disjunction with a unique truthmaker: p , if p is true and q false in w ; and q , if q is true and p false in w . So using Alt_w , we can make ‘ p OR q ’ equivalent to ‘ p ’ in worlds where only p is true, and make ‘ p OR q ’ equivalent to ‘ q ’ in worlds where only q is true.

Definition 10 tells us when ‘ p OR q ’ is true at a given world w . This underspecifies a two-dimensional entry for disjunction, because it does not specify *in virtue of what role*—world-as-actual or world of evaluation— w parameterizes the Alt function. Using x , the world of evaluation, will get us Boolean disjunction even in embedded environments; using y , the world-as-actual, will get us the disjunct that is true in the actual world, no matter which world the disjunction is being evaluated in. Let us therefore stipulate that it is y , the world-as-actual, that parameterizes Alt .

Postulate 4 (Two-dimensional ‘or’). $s, y, x \models (p \text{ OR } q)$ iff there is some $\alpha \in Alt_y(p, q)$ such that $s, y, x \models \alpha$.

When the disjunction in Postulate 4 is embedded under a deontic modal like OUGHT, the modal shifts the world of evaluation, but not the world-as-actual. That means any dependence the embedded sentence displays on the world-as-actual remains anchored to that world. So at any point of evaluation where only p is (actually) true, OUGHT(p OR q)

²⁰Here, I use ‘ p ’ and ‘ q ’ as metalanguage variables over atomics.

is equivalent to OUGHT p , and at any point where only q (actually) is true, OUGHT(p OR q) is equivalent to OUGHT q . Likewise, at points where only p is true, MAY(p OR q) is equivalent to MAY p , and at points where only q is true, MAY(p OR q) is equivalent to MAY q . These are the two sets of transitions in the premises of Patterns A and B in §3.²¹

Postulate 4 is the last postulate we will need to explain free choice permission and Ross’s puzzle. But in order to really see how the object language works, we need to bring the whole semantics into two dimensions.

5 Two Dimensions: Deontic Modality

Deontic modals are classical “one-dimensional” modals, shifting only the world of evaluation, x , and not the world-as-actual, y : their interpretation holds fixed the proposition expressed by ϕ in evaluating the question, *ought it to be the case that ϕ ?*

Our old definition of obligation at $\langle s, w \rangle$, Definition 3-a, was

p is obligatory at $\langle s, w \rangle$ iff $\forall v$ such that wRv , p is true at $\langle s, v \rangle$.

which is settled-true at s iff it holds at every world in s . In our full system, sentences ϕ express propositions only relative to some choice of world-as-actual. Using ‘ y ’ for this world, Definition 3-a upgrades to

ϕ is obligatory at $\langle s, y, x \rangle$ iff $\forall x'$ such that xRx' , the proposition ϕ expresses at y is true at $\langle s, x' \rangle$.

So our new semantic entry is

Definition 11 (OUGHT). $s, y, x \models O\phi$ iff $\forall x' \in s$ such that xRx' : $s, y, x' \models \phi$.

This is just our old OUGHT, with an free y parameter added. The OUGHT in Definition 11 is local twice over: its truth at $\langle s, y, x \rangle$ depends on which worlds x is R -related to, and what proposition ϕ expresses, relative to y . So when $O\phi$ occurs unembedded, $O\phi$ is settled-true at s just in case, at every world $w \in s$, the proposition ϕ expresses, *at w* , is obligatory, *in w* .

What of Ratifiability, our companion notion to deontic obligation? In our protosemantics in §2, we endorsed this notion of a proposition’s being ratifiable (Definition 8):

p is ratifiable at $\langle s, w \rangle$ iff $\exists v \in s$ such that

- (i) p is true in $\langle s, v \rangle$;
- (ii) p is admissible in $\langle s, v \rangle$.

²¹That is, assuming that conditionalizing s on p reduces s to the (or a) *largest* subset $s|p$ throughout which p is true, rather than some smaller subset throughout which e.g. *both* p and q are true. For an explicit connection between this idea and the semantics of an object-language indicative conditional, see, for example, Kratzer (1991a), Yalcin (2007, pg. 998), and MacFarlane & Kolodny (2010, pg. 136).

once again, in our full system, sentences ϕ express determinate propositions only relative to some choice of world-as-actual. Using ‘ y ’ for this world, Definition 8 upgrades to

Definition 12 (MAY). $s, y, x \models M\phi$ iff $\exists v \in s$ such that (i) $s, y, v \models \phi$ and (ii) $\exists v' \in s$ such that vRv' and $s, y, v' \models \phi$.

Once again, this is just our old MAY, with a free y -parameter added. Our semantic entry for MAY therefore combines the *global* character of the underlying notion of propositional ratifiability (independence of $M\phi$ from the x -parameter) with the *local* character of propositional content (dependence of $M\phi$ on the y -parameter). The derived settled-truth conditions for unembedded $M\phi$, where the same world plays both roles, can be glossed like this: $M\phi$ is settled-true at s just in case there is some world v in s such that the proposition ϕ expresses, *at* v , is admissible, *in* v .²²

This is a nice result. The two-dimensional semantics for disjunction, when combined with the underlying inferential properties of ratifiability, will give us proofs of (FC) and (R+)—the inference to the permissibility of each disjunct—for both MAY and OUGHT (see Appendix, Theorems 2 and 3). When a sentence of the form MAY(p OR q) is settled-true at s , we can describe the agent’s situation metaphorically as follows. To be permitted to do p OR q , where p OR q is an disjunction of future contingent actions, is like being issued a *ticket of permission*, bearing p and q , that can be valued in two different ways. If the agent sees to it that exactly one of the disjuncts is true, it is the unique interpretation of the ticket. Looking back from her post-choice perspective, she will see that she has done what is permissible. However, where there is no true disjunct, or too many of them, all that follows is that the ticket has some value or other amongst the possibilities provided by the disjuncts. That is enough to guarantee that, if the agent does neither p nor q , she refrained from *something* she had permission to do, and that, if she does more than one of those things, she did *more* than she had explicit permission to do. From this latter fact, exclusivity follows.

Things are similar in the OUGHT case. Here, the agent is issued a ticket *obligating* her to do p OR q . The one she picks values the propositional content of the disjunction, putting her in a state where she has done what she was obligated to do; so long as she chooses to perform only one of the disjuncts, she will have discharged, at her post-choice context, the content the obligation picks out at that context. But if she does both p and q , she does *more* than she was obligated to do. Since the obligation statement has nothing to say about the other disjunct, it is entirely possible that it was *not* obligatory—or, indeed, even permissible. This captures the other set of exclusivity intuitions we had in §1, about the case where I ought to cancel or postpone the party: while each option is permissible, their conjunction may well be *impermissible*.

6 Two Dimensions: Consequence

Our target notion of consequence in the protosemantics of §2 was the preservation of settled-truth, or truth at every world in s . The move to a two-dimensional semantics

²²More technically: $M\phi$ is settled-true at s iff $\exists v \in s$ such that: (i) $s, v, v \models \phi$ and (ii) $\exists v' \in s$: vRv' and $s, v, v' \models \phi$.

presents us with a further choice: whether we seek the preservation of truth at points where $x = y$, or more generally at any ‘two-dimensional’ point. Following tradition, I will say that it is *diagonal consequence* that most closely approximates intuitive consequence relations between natural language sentences.²³ This reflects the motivation we began with: the idea that ϕ ’s being true at w unpacks (and repacks) into the notion that the proposition expressed by ϕ at w is true *in* w . ψ is a diagonal consequence of ϕ just in case, for any s , if ϕ is settled-true at every point $\langle y, x \rangle$ in s such that $y = x$, then so is ψ . That gives us a consequence relation that preserves diagonal settled-truth:

Definition 13 (Settled-Diagonal (SD) Consequence). $\phi \vDash \psi$ iff for any s : if $s, w, w \vDash \phi$ for all $w \in s$, then $s, w, w \vDash \psi$ for all $w \in s$.

Let us take a closer look at disjunction from the point of view of SD consequence. In the style made familiar by Stalnaker (1978), we can capture the complete semantic profile of ‘ p OR q ’ with a two-dimensional matrix, once again letting w_1 - w_4 represent the standard four lines of a truth-table for atomic p and q .²⁴

	w_1	w_2	w_3	w_4
w_1	T	T	T	F
w_2	T	T	F	F
w_3	T	F	T	F
w_4	T	T	T	F

Figure 3: 2D matrix for ‘ p OR q ’

In w_2 , where only p is true, ‘ p OR q ’ is equivalent to p . In w_3 , where only q is true, ‘ p OR q ’ is equivalent to q . The diagonal of the matrix witnesses the T-T-T-F truth-conditions of Boolean disjunction. With diagonal consequence on the table, the status of classical propositional logic is simple to state. Diagonal consequence begins by looking at the truth-values of sentences at diagonal points. Within the propositional fragment of the language, there are no (one-dimensional) modal operators, so nothing ever moves us *off* the diagonal. This is what lies behind

Theorem 1 (Classicality). For any ϕ in the nonmodal fragment of L , $\vDash \phi$ iff ϕ is a theorem of classical logic.

The matrix also shows why ‘OR’-introduction is not valid off the diagonal—for example, once we start looking at modally embedded environments. For example, it does not

²³See, for example, the corresponding notion of validity in Kaplan (1989, pg. 547), and the notion of *real world validity* in Davies & Humberstone (1980).

²⁴The y -axis of the matrix represents world-types in their role as world-as-actual, and the x -axis represents world-types in their role as world of evaluation. Hence, reading across a row will give one the proposition expressed by the sentence, when the y -axis world plays the role of the world-as-actual. I emphasize that Stalnaker’s notion of a propositional concept is a pragmatic one, while I use the matrix here to model the semantics. See Stalnaker (1978, pg. 81) for discussion.

preserve truth to infer from the proposition expressed by ‘ p ’ to the proposition expressed by ‘ p OR q ’ when y is held fixed at w_3 . In w_3 , ‘ p OR q ’ is equivalent to ‘ q ’. So the inference ‘ p , therefore (p OR q)’ is equivalent at w_3 to the inference ‘ p , therefore q ’, which is clearly *not* valid. This is what lies behind

Fact (Failure of ‘OR’ introduction). $O\phi \not\equiv O(\phi \text{ OR } \psi)$ and $M\phi \not\equiv M(\phi \text{ OR } \psi)$.

This observation targets the intuitive badness of inferences like “you may post the letter; therefore, you may post the letter or burn it.” The problem is not (in the first place) that such transitions are misleading, infelicitous, or uncooperative. They are just plain invalid: they can take us from a premise that is settled-true at s to a premise that is settled-false at s . This is because we have a logic in which (ϕ or ψ) is a diagonal, but not an unrestricted, consequence of ϕ : disjunction introduction will pattern with cases in which it is valid to introduce a disjunct *outside* the scope of an upward-entailing intensional operator O , but not *inside* its scope.²⁵ The status of disjunction introduction—the inference from ϕ to (ϕ or ψ)—will be that of an *a priori contingent* inference, in the sense of Evans (1977). It is like one’s knowledge of the truth of the sentence

(12) I am here now.
IHN

Since (12) is true at all diagonal points, conjoining it with any sentence will preserve truth at a diagonal point. We might call an inference rule that reflects this fact ‘ \wedge *IHN*’-Introduction: from any ϕ , conclude ($\phi \wedge$ *IHN*).

‘ \wedge <i>IHN</i> ’ Introduction	$\frac{\phi}{\phi \wedge \textit{IHN}}$	$\frac{O(\phi)}{O(\phi \wedge \textit{IHN})}$
	valid	invalid

For example, if $2+2 = 4$, then $2+2=4$ and I am here now; but from the fact that it is (metaphysically) necessary that $2+2=4$, it does not follow that it is (metaphysically) necessary that ($2+2=4$ and I am here now), since it is not metaphysically necessary that I am here now.

Disjunction introduction—an ‘or ψ ’ rule—works the same way: if I am mailing the letter, it follows that I am mailing it or burning it; but from the fact that I *ought* to mail the letter, it does not follow that I ought to mail it or burn it.

7 Conclusions and Horizons

This paper began with two puzzles: one about deontic permissibility, and one about deontic obligation. Our initial step towards their joint solution was indirect: we began by

²⁵ An operator M is *upward-entailing* just in case preserves the direction of (here, settled-diagonal) consequence: if $\phi \vDash \psi$, then $M(\phi) \vDash M(\psi)$.

‘or ψ ’ Introduction

$$\frac{\phi}{\phi \text{ or } \psi}$$

valid

$$\frac{O(\phi)}{O(\phi \text{ or } \psi)}$$

invalid

looking at cases of act dependence, formalizing them in agentive Kripke frames. Our first semantic move was to postulate that object-language MAY tracks the concept of permissibility developed for these cases. Our operator is thus constrained by the language-independent intuitions we have about choiceworthiness, in cases where our future selves disagree. Our second move was a two-dimensional, “truthmaker” account of disjunction, which allowed the object language to mirror the dependence of norms on acts by making the *propositional content* of a deontic claim dependent on what becomes actual.

The combination of the two moves allows the stronger-than-classical conclusions of free choice permission to follow, and blocks the inference in Ross’s puzzle. It also predicts the positive entailment properties of disjunction under OUGHT in terms of MAY (the pattern we called (R+)). Assuming there is an object-language conditional ‘ \rightsquigarrow ’ which is related to \models by the Deduction Theorem, we can add (Conditionals-M): $M(\phi \text{ OR } \psi) \models \neg\phi \rightsquigarrow M\psi$, and (Conditionals-O): $O(\phi \text{ OR } \psi) \models \neg\phi \rightsquigarrow O\psi$.²⁶ The full system preserves the inviolability of the role classical logic plays in our reasoning, while unifying the puzzles in terms of their common factor.

The thought behind leveraging language-independent intuitions about choiceworthiness *before* turning back to the object language was to show that our analysis of MAY has independent appeal. Having done what is admissible, in the context you occupy once you have done it, has claim to being the primary way permissibility constrains action; it is comparable to the claim that the primary goal of assertion is to express a proposition that is true, in the context you occupy once you begin to speak. These are not quirks of English, but reflections of the nature of assertion and action.

The slice of data we have confronted focuses tightly on deontic modality and its conceptual underpinnings. A question to consider by way of conclusion is this: does the appeal to ratifiability make the account of the phenomena defended in this paper too narrow?²⁷ Free choice permission, in particular, is often studied as a part of the general phenomenon of *free choice*, which involves disjunction under modals, including epistemic modals, circumstantial modals, and generics:

- (13) It might be raining or snowing.
(Santorio & Romoli, 2015, pg. 9)

²⁶A conditional \rightsquigarrow obeys the Deduction Theorem just in case $\Delta, \phi \models \psi$ entails $\Delta \models \phi \rightsquigarrow \psi$. Lacking space to defend this (plausible) feature of the natural language indicative conditional, I put (Conditionals-M) and (Conditionals-O) in a separate category in Table 2.

²⁷I thank an anonymous referee for encouraging me to consider this question, and for pointing me to examples of the type of (14).

May:	
(Failure of ‘or’ intro)	$M\phi \not\models M(\phi \text{ OR } \psi)$
(FC)	$M(\phi \text{ OR } \psi), \blacklozenge\phi, \blacklozenge\psi \models M\phi \wedge M\psi$
Ought:	
(Failure of ‘or’ intro)	$O\phi \not\models O(\phi \text{ OR } \psi)$
(R+)	$O(\phi \text{ OR } \psi), \blacklozenge\phi, \blacklozenge\psi \models M\phi \wedge M\psi$
Propositional Fragment:	
(Classicality)	for ϕ in the propositional fragment of L , $\models \phi$ iff ϕ is a theorem of classical logic.
Conditional Inferences	
(given Deduction Theorem:)	
(Conditionals-M)	$M(\phi \text{ OR } \psi) \models \text{if } \neg\phi, \text{ then } M\psi$
(Conditionals-O)	$O(\phi \text{ OR } \psi) \models \text{if } \neg\phi, \text{ then } O\psi$

Table 2: Semantic entailments.

- (14) Petunias or hydrangeas can grow here.
(after (Kratzer, 1981, pg. 28))
- (15) Elephants live in Africa or Asia.
(Nickel, 2010, pg. 480)

Whether this is a shortcoming of the account defended in this paper, I think, depends on two questions. First, is general free choice a reason to think the slice of data in Table 2 is not cut along a natural axis? I am not sure this conclusion is warranted. While free choice permission was the first of our named puzzles, and free choice does seem to be a feature of other modal flavors, our second puzzle, Ross’s puzzle, appears to have particularly strong force in the deontic case: compare the inference “you ought to post the letter, therefore, you ought to post it or burn it” to “Bill must [epistemic] be in the garden; therefore, he must [epistemic] be in the garden or in the attic.”²⁸ The combination of free choice permission and Ross’s puzzle indeed seems unique to deontic modality, and that is a reason for favoring an account that focuses specifically on deontic modality, and the relationship between deontic OUGHT and deontic MAY.

Second, it may be possible to adapt the account of free choice permission—free choice in the deontic case—to the other modal flavors in sentences like (13)-(15). Whether this can be done is a complex question. It should be noted that the concept of ratifiability—the key to our semantic Postulate 3—is not confined to traditional, or “practical” decision theory. It crops up in epistemic utility theory—decision theory for epistemic acts, such as assigning credences to propositions. Greaves (2013) states a typical case:

²⁸ Ross’s original (Ross, 1941) concerns imperatives; for OUGHT, see, inter alia, von Fintel (2012) and Cariani (2011).

Leap. Bob stands on the brink of a chasm, summoning up the courage to try and leap across it. Confidence helps him in such situations: specifically, for any value x between 0 and 1, if Bob attempted to leap across the chasm while having degree of belief x that he would succeed, his chance of success would then be x . (Greaves, 2013, pg. 2)

The question for an epistemic utility theorist is what credence it is rational for Bob to assign to success in this case. Our semantic theory does suggest an answer to a related question about the practical case: one should be entirely confident that an action is permissible if it is admissible, given that one does it; this is, given Postulate 3, an *analytic* feature of deontic MAY.

Abstracting from the fine-grained values associated with credences, the quantity of interest in an epistemic utility theoretic case like (Leap) is:

Definition 14 (Epistemic Ratifiability). p is *epistemically ratifiable* iff it is epistemically possible, given that one believes it.

A belief that he can leap across the chasm is, in this sense, epistemically ratifiable for Bob. Is it plausible to think that this property would be sufficient for p (for example, the proposition that *Bob's leap will succeed*) to be epistemically possible (for example, for Bob), full stop?

First, Bob's situation does seem reminiscent of the free choice case; there is an intuition that Bob cannot really go wrong in picking a credence for p in (Leap). This is much like the intuition that, when I give you permission to have coffee or tea, you cannot go *wrong* in picking the coffee, and cannot go *wrong* in picking the tea; it is most like the intuition recorded by (FC). So perhaps epistemic ratifiability is sufficient for epistemic possibility—of the kind recorded by the object-language epistemic modal MIGHT—just as I argued that deontic ratifiability is sufficient for deontic possibility, of the kind recorded by the object-language deontic modal MAY. As Greaves' discussion highlights, however, there are differences between the epistemic and practical cases that should raise doubts about any easy parallelisms.²⁹ And so the issue, I think, is best left to future research.³⁰

References

Barker, Chris (2010). "Free choice permission as resource-sensitive reasoning." *Semantics and Pragmatics*, 3: pp. 1–38.

Cariani, Fabrizio (2011). "Ought' and Resolution Semantics." *Noûs*.

Charlow, Nate (2013). "What we know and what to do." *Synthese*, 190: pp. 2291–2323.

²⁹ See, for example, Greaves' discussion of "epistemic bribes" (*op. cit.*, pg. 35-36).

³⁰ Thanks to John MacFarlane, Seth Yalcin, Wes Holliday, Lara Buchak, Alex Kocurek, Stephen Darwall, and two anonymous referees at *Philosophers' Imprint* for helpful feedback on this paper. For stimulating discussion, I am also indebted to audiences at UC Berkeley, UC Santa Cruz, and Magdalen College, Oxford.

- Chierchia, Gennaro (2006). “Broaden Your Views: Implicatures of Domain Widening and the ‘Logicity’ of Language.” *Linguistic Inquiry*, 37(4).
- Davies, Martin, and Lloyd Humberstone (1980). “Two Notions of Necessity.” *Philosophical Studies*, 38(1): pp. 1–30.
- Evans, Gareth (1977). “Reference and Contingency.” *The Monist*, 62: pp. 161–189.
- von Fintel, Kai (2012). “The best we can (expect to) get?: Challenges to the classic semantics for deontic modals.” Central APA session on Deontic Modals.
- Fox, Danny (2007). “Free Choice Disjunction and the Theory of Scalar Implicatures.” In U. Sauerland, and P. Stateva (eds.) *Presupposition and Implicature in Compositional Semantics*, New York, Palgrave Macmillan.
- Fusco, Melissa (2014). “Free Choice Permission and the Counterfactuals of Pragmatics.” *Linguistics and Philosophy*, 37(4): pp. 275–290.
- Fusco, Melissa (2015). *Deontic Disjunction*. Ph.D. thesis, University of California, Berkeley. In progress.
- Gibbard, Allan, and William Harper (1978). “Counterfactuals and Two Kinds of Expected Utility.” In C. A. Hooker, J. J. Leach, and E. F. McClennen (eds.) *Foundations and Applications of Decision Theory, Vol 1*, Dordrecht: D. Reidel.
- Greaves, Hilary (2013). “Epistemic Decision Theory.” *Mind*.
- Hare, Caspar, and Brian Hedden (2015). “Self-Reinforcing and Self-Frustrating Decisions.” *Noûs*. Forthcoming.
- Heim, Irene (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- Jeffrey, Richard (1983). *The Logic of Decision*. University of Chicago Press.
- Kamp, Hans (1973). “Free Choice Permission.” *Proceedings of the Aristotelian Society, New Series*, 74: pp. 57–74.
- Kaplan, David (1989). “Demonstratives.” In J. Almog, J. Perry, and H. Wettstein (eds.) *Themes from Kaplan*, Oxford University Press.
- Kratzer, Angelika (1981). “The Notional Category of Modality.” In H.-J. Eikmeyer, and H. Rieser (eds.) *Words, Worlds, and Context*, de Gruyter.
- Kratzer, Angelika (1991a). “Conditionals.” In A. von Stechow, and D. Wunderlich (eds.) *Handbuch Semantik*, Berlin: de Gruyter.
- Kratzer, Angelika (1991b). “Modality.” In A. von Stechow, and D. Wunderlich (eds.) *Semantics: An International Handbook for Contemporary Research*, Berlin: de Gruyter.

- Kratzer, Angelika, and Junko Shimoyama (2002). "Indeterminate Pronouns: The view from Japanese." *Proceedings of the Third Tokyo Conference on Psycholinguistics*.
- Lassiter, Daniel (2011a). *Measurement and Modality: The Scalar Basis of Modal Semantics*. Ph.D. thesis, New York University.
- Lassiter, Daniel (2011b). "Nouwen's Puzzle and a Scalar Semantics for Obligations, Needs, and Desires." In *Proceedings of SALT 21*, pp. 694–711.
- MacFarlane, John (2013). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press.
- MacFarlane, John, and Niko Kolodny (2010). "Ifs and Oughts." *Journal of Philosophy*, 107: pp. 115–143.
- McNamara, Paul (2010). "Deontic Logic." In *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/fall2010/entries/logic-deontic/>.
- Nickel, Bernhard (2010). "Generically Free Choice." *Linguistics and Philosophy*, 33(6): pp. 479–512.
- Ross, Alf (1941). "Imperatives in Logic." *Theoria*, 7(1): pp. 53–71.
- Santorio, Paolo, and Jacopo Romoli (2015). "Probabilistic *might* and the free choice effect." [Paolosantorio.net](http://paolosantorio.net).
- Sayre-McCord, Geoffrey (1986). "Deontic Logic and the Priority of Moral Theory." *Noûs*, 20(2).
- Simons, Mandy (2005). "Dividing Things Up: The Semantics of *Or* and the Modal/*Or* Interaction." *Natural Language Semantics*, 13: pp. 271–316.
- Stalnaker, Robert (1978). "Assertion." In *Context and Content*, Oxford University Press, pp. 78–95. In Stalnaker (1999).
- Stalnaker, Robert (1999). *Context and Content*. Oxford University Press.
- Thomason, Richmond (1970). "Indeterminist Time and Truth-Value Gaps." *Theoria*, pp. 264–281.
- von Wright, G. H. (1969). *An Essay on Deontic Logic and the General Theory of Action*. North Holland: Amsterdam.
- Yalcin, Seth (2007). "Epistemic Modals." *Mind*, 116: pp. 983–1026.