Free choice permission in STIT

FREDERIK VAN DE PUTTE[1]

Abstract: We argue for a new approach to free choice permission in the context of a-temporal **STIT** logic. According to our analysis, an agent has a free choice permission w.r.t. two propositions φ and ψ iff (a) the agent is permitted to see to $\varphi \wedge \neg\psi$ and (b) the agent is permitted to see to $\psi \wedge \neg\varphi$. The primitive notion of permission we use is the dual of one of Horty's operators for "ought to do" from (Horty, 2001). We argue that the approach improves on existing proposals in various ways.

Keywords: Free choice permission, Deontic logic, STIT logic, Multi-agency

1 Introduction

Jane goes to the fruit buffet and asks the waiter which pieces of fruit she can take. The waiter replies: "you may take an apple or a pear". We take this statement to be equivalent to "Jane is free to choose between taking an apple and a pear" or to "Jane has a free choice between taking an apple and taking a pear". Let us represent such statements formally as $\mathsf{F}_j(apple, pear)$, where $apple, pear$ are states of affairs (that Jane takes an apple, resp. a pear) and j is shorthand for Jane. In this case, we say that the agent (i.c. Jane) has a free choice permission (henceforth FCP) w.r.t. the two options $apple, pear$. Likewise, an agent may have a FCP w.r.t. three, four, or more options.

Our aim in this paper is to propose a new semantics for the operator F_j, one that arguably improves on existing accounts. We focus on the binary case for the sake of simplicity; however, all our observations and our own proposal generalize readily to a finite number $n \geq 2$ of choices.

Mind that we take FCP here to concern a *normative* claim, i.e. it is not simply a descriptive claim about what the agent can choose (as in "you can

[1]Research for this paper was funded by the Flemish Research Foundation (FWO-Vlaanderen). The basic ideas for this paper were obtained during a research stay in Bayreuth (July 2016), which was co-funded by the PIOTR project (RO 4548/4-1). Many thanks to Olivier Roy, Igor Sedlar, and Nathan Wood for stimulating discussions, as well as more concrete suggestions.

choose to kill your wife or not, but there is no reason whatsoever to kill her"), but it is a claim about what the agent is *permitted* to choose. In the remainder, we moreover presuppose that we are working with a *single* normative system that grounds claims concerning FCP.

Another very important point is that FCP, as we think of it, concerns a given agent and a concrete deliberative context for that agent, i.e. a situation in which that agent has to make a choice among a number of options that are "real". By the latter we mean that choosing one such option entails that the option is effectively realized. So our problem is: given that the agent has to make a choice, when can we say that the agent is permitted to choose between φ and ψ?

Perhaps it makes sense to claim that one (everyone, someone, ...) is free to order a pizza or spaghetti at piazza San Marco in Venice, but these are not two options that present themselves to you at this very moment (unless of course you happen to be reading this paper at San Marco). So in this sense *you* do not have the FCP *here and now* to order a pizza or spaghetti at San Marco. We explicitly leave such forms of "free choice permission" out of the picture in this paper.[2]

To fully appreciate this point, consider the following example from Sedlár (2016): suppose that according to the laws of your country, you are permitted to travel abroad and you are permitted to vote. These two actions are clearly distinct, meaning that you can do one without the other. Still, it seems to make little sense to say that you have a free choice between traveling abroad and voting. The two are simply not related, and hence we do not think in terms of making a choice between either.

Suppose however that you are in a specific situation in which traveling abroad is one of your options and voting is another. For instance, you may have to decide here and now between going to a conference, thus missing out on the elections, and staying home so that you can vote; there happens to be no way you can combine both. Assume moreover that not voting is permitted in your country, and likewise, that you are not obliged to go to the conference either. There is no other feasible option that would be way

[2] One might argue that this claim really communicates that *once* you are at San Marco, you have the free choice between pizza and spaghetti, and (perhaps) moreover that you are permitted to go to San Marco in the first place. To model such claims, one could extend our current analysis, e.g. by adding a possibility operator which ranges over alternative deliberative contexts and (perhaps) also a similarly wide-ranging permission operator. But even if one disagrees, it still makes sense to focus on the concrete here-and-now FCP as we do, before trying to tackle other (probably more complicated) notions.

better than each of these options. Then, on the approach we advocate here, it *does* make sense to say you are free to choose between voting and traveling abroad.

In other words, whether we can speak of a genuine choice between two options is a matter of context. In our account, the options of agents are made explicit within the models, whence we can explicate necessary and sufficient conditions for a FCP. This means in turn that FCP does not only depend on the norms that apply in the context, but also to the available choices of the agent — we return to this point below.

2 Some general observations

To get our analysis off the ground, let us make a number of observations about FCP and the way it relates to permission in natural language. Some of these are not original, but we collect them here for ease of reference.[3]

(I) $\mathsf{F}_\alpha(\varphi, \psi)$ does not entail that every way of making φ, resp. ψ true, is permitted.[4] In other words, $\mathsf{F}_\alpha(\varphi, \psi)$ does not entail that either φ or ψ is "sufficient for permissibility".[5]

To understand this claim, suppose that whoever takes a pear or an apple should also take a napkin (since the apples and the pears are very juicy). Jane may well take a pear or an apple without taking a napkin. That mere fact by no means entails that she is no longer free to choose between taking an apple and taking a pear. In fact, "you are free to choose between an apple and a pear, but do take a napkin in case you take either" sounds perfectly consistent. It would be strange to say that in such cases, you are really only free to choose between taking an apple and a napkin, or taking a pear and a napkin.

[3]Although we call them "observations" and although they are clearly inspired by natural language, these points should not be taken as purely descriptive facts of the matter. Rather, they are salient properties of what we consider a useful and natural concept of free choice permission, one that allows us to report on the permissibility of an agent's choices in a given situation.

[4]A similar remark is made by Giordani and Canavotto (2016, p. 89): "we are ordinarily allowed to choose between alternative actions even if there are ways of performing such actions that lead to a violation of the law".

[5]See (Van De Putte, 2016) for a formal investigation of this notion of sufficiency.

(II) A free choice permission w.r.t. φ and ψ – relative to a given body of norms – entails that both choosing φ and choosing ψ are permitted here and now, relative to the same body of norms.

"You are free to choose between taking an apple or a pear, but actually all apples must remain untouched" sounds self-contradictory – at least, if we interpret both parts of the sentence as expressing information about one and the same normative system. When someone makes such a claim, we will automatically infer that the speaker wants to distinguish between two normative systems (e.g. the law regarding apples and pears on the one hand, and the rules of the house regarding apples and pears on the other).

(III) That φ and ψ are permitted does not entail that there is a FCP.

There are at least two ways this can be argued against:

(iii.a) φ may entail ψ. It may be permitted to take a piece of fruit, and it may also be permitted to take an apple, but it doesn't make sense to infer from this that you are free to choose between taking a piece of fruit and taking an apple.

(iii.b) It may be obligatory that φ whenever ψ. For instance, it may be obligatory that you take an apple whenever you take a pear. In that case, even if both are permissible, one cannot infer that you are free to choose between taking an apple and taking a pear.

One may argue that if φ and ψ are mutually exclusive (and i.c. when $\psi = \neg\varphi$), then the permissibility of both entails that there is a FCP. Here, our stance is more subtle: we will argue (in Sections 4 and 5) that this depends essentially on the notion of permissibility one is using.

(IV) FCP is always relative to an (a group of) agent(s): it is always the permission *of an agent* to choose between a number of things.

It makes no sense to say there is a free choice between X and Y, unless one refers (implicitly) to one or more agents that are permitted to make this choice. Also, FCPs of one agent α need not coincide with FCPs of another agent β. Jack may be free to choose whether his back door is left open at night, but his neighbour Daniel clearly does not have a FCP w.r.t. this same state of affairs. This distinguishes FCP from the concept of permission that is at stake in sentences like "it is permitted that your car is parked here", where it does not matter who brings it about that the car is parked.

(V) $\mathsf{F}_\alpha(\varphi, \psi)$ entails that α can in fact choose between φ and ψ, that both are "live options" for α.

"You are free to choose between winning or losing the lottery, but there is no way you can choose among either" seems to make little sense – again, when it is taken to report a permission that applies here and now. In this sense, FCP is a mixed notion, since it presupposes not only normative claims, but also claims about the possible choices of agents.

3 Existing Accounts of FCP

In the deontic logic literature, FCP is usually considered problematic in view of our Observation **(II)**. That is, when FCP is formalized by $\mathsf{P}(\varphi \vee \psi)$, one cannot accomodate this observation within Standard Deontic Logic (**SDL**).[6] Adding a corresponding axiom schema would trivialize the logic, whence one is bound to look for alternative accounts of permission or disjunction in order to handle such inferences.

Mind however that permissions of disjunctions do not behave uniformly in natural language. Often they do behave like FCPs; but often they don't:

(a) If we negate a sentence of the form "it is permitted that φ or ψ", then that permission behaves like a disjunction of permissions after all. That is, "it is not permitted that you take an apple or a pear" usually means "it is not permitted that you take an apple *and* it is not permitted that you take a pear". This is not at all the same as saying that "you are not free to choose between having an apple or a pear" – perhaps someone else is to decide, or perhaps you can only take an apple.[7]

(b) Embedded permissions over a disjunction are also sometimes interpreted as disjunctions of permissions. That is, "You may take an apple or a pear, but ask the waiter which of both" usually communicates uncertainty about the norms at hand, rather than a FCP.

Note also that, even if **SDL** does not validate the inference from $\mathsf{P}(\varphi \vee \psi)$ to $\mathsf{P}\varphi \wedge \mathsf{P}\psi$, it does validate other seemingly plausible principles. For

[6] **SDL** is just the normal modal logic **KD**, with $\Box$ read as "it is obligatory that" and $\Diamond$ as "it is permitted that".

[7] Thanks to Malte Willer for drawing our attention to this problem (during a Q&A session at the DEON2016 conference).

instance, from "it is permitted to take an apple and a pear", a competent language speaker will infer that "it is permitted to take an apple". Rejecting all such inferences, simply in order to be able to validate the inference from $\mathsf{P}(\varphi \vee \psi)$ to $\mathsf{P}\varphi \wedge \mathsf{P}\psi$, can hardly be seen as a proper solution to the problem of free choice permission.

This is not in itself an argument against reducing FCP to a combination of *some* concept of permission and *some* concept of disjunction. It does show that one cannot simply defend such an account by referring to the behavior of permission in natural language. Quite to the contrary: having various formal operators of permission around seems inevitable, if one is to explain the various ways utterances such as "may", "can", "it is permitted that" and the like behave and interact with "not", "and", and "or" in natural language.

But there is more. Even when intended solely as a formal account of FCP, the existing proposals are insufficient. To argue for this in full would require that we go over every such proposal, spell out the relevant formalities, and argue why it gives counterintuitive properties for F. For reasons of space we cannot do this here; we will however point out in brief terms what our main worries are.

First, one general weakness of existing accounts is that they do not make it explicit which agent has a given FCP. Perhaps these approaches can be enriched so that one obtains an agent-relative concept of FCP, which accomodates observations **(IV)** and **(V)**; but this remains to be done. Mind that, even if one abstracts from the agent in question – as this is often done in deontic logic –, one should still be able to distinguish between states of affairs that the agent can see to, and states of affairs that are beyond its abilities; at least if one agrees with observation **(V)**.

Second, it is shown in (Hansson, 2013) that any intensional account of FCP in terms of unary permission and disjunction will give extremely counterintuitive results, due to the validity of replacement of classical equivalents (RE). By (RE), we can e.g. infer from the fact that it is permitted to either take an apple or not, $\mathsf{P}(a \vee \neg a)$, that it is also permitted to either take a napkin or not, $\mathsf{P}(n \vee \neg n)$. But the free choice interpretation of the former statement clearly need not imply the free choice interpretation of the latter.

Hansson's observation still leaves room for alternative accounts. One of them is to translate FCP straight into a *conjunction* of permissions: $\mathsf{F}(\varphi, \psi) =_{\mathsf{df}} \mathsf{P}\varphi \wedge \mathsf{P}\psi$ — let us call this the *conjunctive account* of FCP. This is in line with Makinson (1984), who claims that the so-called paradox of FCP only arises due to a mistranslation of natural language into the deontic formalism.

However, in view of our Observation **(III)**, this simple version of the conjunctive account does not allow us to fully characterize FCP. In Section 4 we will consider more refined conjunctive accounts. There, it will become clear that also these variants fall short of capturing FCP, when based on the standard concepts of permission taken from the literature. In Section 5 we will however show that, when P is itself an agent-dependent notion of permission, one variant of the conjunctive account does work.

Another alternative – which is suggested by Hansson himself – is to treat FCP as a primitive, binary operator. We will not have much to say about this approach here; it may well be promising as a "minimalistic" account of FCP, omitting explicit operators for agency and ability.[8] Still, as we show in the remainder of this paper, for the specific notion of FCP in a deliberative context it *is* possible to reduce FCP to a suitable combination of unary permission and classical connectives. Such an account moreover has the advantage that it clearly links the various well-known notions of permission to FCP, thereby providing an explanation of (some of) its behavior.

Finally, one may question the very idea that a logic for FCP should be intensional and hence closed under (RE). That is, one may advocate a hyperintensional account of permission, following Anglberger, Faroldi, and Korbmacher (2016) and Fine (2016). For instance, Fine works with a distinct set of "ideal" actions, where actions are to be understood as states (ordered according to part-whole). $\mathsf{P}\varphi$ is then true iff every action that is in compliance with φ is contained in some ideal action. Since the set of actions in compliance with $\varphi \vee \psi$ equals the union of the set of actions in compliance with φ and the set of actions in compliance with ψ, we immediately get $\mathsf{P}(\varphi \vee \psi) \equiv (\mathsf{P}\varphi \wedge \mathsf{P}\psi)$. This account also validates the intuitive principle that $\mathsf{P}(\varphi \wedge \psi) \vdash \mathsf{P}\varphi$.

However, the problem remains here that $\mathsf{P}(\varphi \vee \psi)$ and $\mathsf{P}\varphi \wedge \mathsf{P}\psi$ are equivalent, which runs counter to our Observation **(III)** if we take $\mathsf{P}(\varphi \vee \psi)$ as the definiens of $\mathsf{F}(\varphi, \psi)$.[9] It is also not clear how "actions" are to be individuated in this framework, when applied to the simple cases that we

[8]The only work in this direction that we are aware of is the unpublished manuscript (Sedlár, 2016). Here, Hansson's suggestion is worked out into a full formal system with operators for permission that have an arbitrary arity. One of the problems of this formal system is that it does not account for our Observation **(V)**. We decided to leave a full discussion of this for another occasion.

[9]One may counter this objection by defining $\mathsf{F}(\varphi, \psi)$ as $\mathsf{P}((\varphi \wedge \neg\psi) \vee (\psi \wedge \neg\varphi))$, but in the context of the hyperintensional logics from (Fine, 2016), negation seems too weak to fully express mutual exclusiveness of the involved states of affairs.

present in Section 4. Even if this does not amount to a wholesale rejection of the hyperintensional account of FCP, it does show that some more work needs to be done in order to obtain a full explication of FCP.

4 Permissions and Choice in STIT logic

In the remainder, we will focus on a very simple notion of a deliberative context. It is our conviction that we should at least try to get our interpretation of the language right for these simple models; once we are there, we can start looking at more complex structures.

The models we obtain are simple versions of the BT+AC-frames known from (Belnap, Perloff, Xu, & Bartha, 2001) and (Horty, 2001), stripped from temporal aspects, and enriched with a deontic component. They can also be seen as multi-agent, one-shot games. The deontic component consists in specifying which outcomes of the game are permissible (normatively ok, acceptable, legal) and which are not. Intuitively speaking, one may think of those outcomes as satisfying all norms that apply in the context at hand.[10]

Language Let $\mathcal{S} = \{p, q, \ldots\}$ be a set of sentential variables and $Agt = \{\alpha_1, \ldots, \alpha_n\}$ be a set of n distinct agents. Our formal language is defined by the following BNF, where $p \in \mathcal{S}$ and $\alpha \in Agt$:

$$\mathcal{W} := p \mid \neg\langle\mathcal{W}\rangle \mid \langle\mathcal{W}\rangle \vee \langle\mathcal{W}\rangle \mid \Diamond\langle\mathcal{W}\rangle \mid \mathsf{P}\langle\mathcal{W}\rangle \mid \mathsf{P}^s\langle\mathcal{W}\rangle \mid [\alpha]\langle\mathcal{W}\rangle$$

The connectives $\wedge, \supset, \equiv$ are defined in the standard way. $\Diamond$ is an existential modality (read as "it is possible that"); P represents weak permission and P^s represents "strong permission" (also known as "deontic sufficiency"). $[\alpha]\varphi$ expresses that "α's choice guarantees that φ is the case".[11]

Frames and Models We assume a set W of possible worlds, and a non-empty set $P \subseteq W$ of permissible worlds. W represents the *modal base*, i.e. the set of all possible outcomes, regardless of what the agents choose. To model choices, we moreover need for each agent α a partition $\mathcal{C}_\alpha$ of W into *choice cells* $X \subseteq W$. *Frames* are thus triples of the type $\langle W, P, \{\mathcal{C}_\alpha\}_{\alpha \in Agt}\rangle$.

[10] Accordingly, the operators for weak and strong permission defined below should be seen as weak, resp. strong counterparts of "must" or "obligatory", rather than of "ought". See (McNamara, 1996) for a discussion of these various modalities and their behavior in natural language.

[11] $[\alpha]$ is also referred to as the *Chellas stit* in the literature on STIT logic. See e.g. (Horty, 2001) for a discussion of this and other operators for agency in STIT.

Where $w \in W$ and $\alpha \in Agt$, we let $\mathcal{C}_\alpha(w)$ denote the unique $X \in \mathcal{C}_\alpha$ such that $w \in X$, i.e. the choice that is actually taken by α at w.

We assume the condition of *independence of agents*: whenever $X_1 \in \mathcal{C}_{\alpha_1}$ and ... and $X_n \in \mathcal{C}_{\alpha_n}$ and each of the α_i ($i \leq n$) are distinct, then $X_1 \cap \ldots \cap X_n \neq \emptyset$. This condition ensures that no given agent α_i can prevent another agent α_j from making any of its choices.

We say that a *model* is a quadruple $M = \langle W, P, \{\mathcal{C}_\alpha\}_{\alpha \in Agt}, V\rangle$, where $\langle W, P, \{\mathcal{C}_\alpha\}_{\alpha \in Agt}\rangle$ is a frame and $V : \mathcal{S} \to \wp(W)$ is a valuation function. The semantic clauses for the modal operators are as follows (where $w \in W$):

(SC1) $M, w \models \Diamond\varphi$ iff there is a $w' \in W$ such that $M, w' \models \varphi$
(SC2) $M, w \models [\alpha]\varphi$ iff $M, w' \models \varphi$ for all $w' \in \mathcal{C}_\alpha(w)$
(SC3) $M, w \models \mathsf{P}\varphi$ iff there is a $w' \in P$ such that $M, w' \models \varphi$
(SC4) $M, w \models \mathsf{P}^s\varphi$ iff, for all $w' \in W$ such that $M, w' \models \varphi$, $w' \in P$

That α is able to enforce φ can be expressed by $\Diamond[\alpha]\varphi$ – see (Horty, 2001, Section 2.3) for an elaborate discussion of this approach to the logic of ability. We can express and abbreviate the claim that "α has a choice between φ and ψ" as follows:

$$\mathsf{C}_\alpha(\varphi, \psi) =_{\mathsf{df}} \Diamond[\alpha](\varphi \wedge \neg\psi) \wedge \Diamond[\alpha](\psi \wedge \neg\varphi)$$

Defining FCP Given these operators, various ways to define F suggest themselves, as shown in Table 1.

(a)	$\mathsf{P}\varphi \wedge \mathsf{P}\psi$	(f)	$\mathsf{P}^s(\mathsf{C}_\alpha(\varphi, \psi))$
(b)	$\mathsf{P}(\varphi \wedge \neg\psi) \wedge \mathsf{P}(\psi \wedge \neg\varphi)$	(g)	$\mathsf{P}^s([\alpha](\varphi \wedge \neg\psi) \vee ([\alpha](\psi \wedge \neg\varphi))$
(c)	$\mathsf{C}_\alpha(\varphi, \psi) \wedge \mathsf{P}(\varphi \wedge \neg\psi) \wedge \mathsf{P}(\psi \wedge \neg\varphi)$	(h)	$\mathsf{P}^s((\varphi \wedge \neg\psi) \vee (\psi \wedge \neg\varphi))$
(d)	$\mathsf{P}[\alpha](\varphi \wedge \neg\psi) \wedge \mathsf{P}[\alpha](\psi \wedge \neg\varphi)$	(i)	$\mathsf{P}^s(\varphi \vee \psi)$
(e)	$\mathsf{P}(\mathsf{C}_\alpha(\varphi, \psi))$	(j)	$\mathsf{P}^s(\varphi \vee \psi) \wedge \mathsf{C}_\alpha(\varphi, \psi)$

Table 1: Some possible ways to model FCP

However, neither of these formalizations are adequate w.r.t. each of our observations **(I)-(V)**, as we argue in the remainder of this section.

Ad (a)-(d). Note first that the definitions in (a) to (d) are of increasing strength. Here, we will show that even (d) does not suffice to speak of free choice.

Example 1 *α and β both have two choices: either take an apple, or not take it. There are two apples in the fruit basket, so they can in fact both take one. It is permitted that one apple is taken; however, it is not permitted that both apples are taken. (Perhaps some third party should still have the option of taking an apple, upon arrival.) Note that $\mathsf{P}[\alpha]apple_\alpha$ and $\mathsf{P}[\alpha]\neg apple_\alpha$ are both true (here, $apple_\alpha$ stands for "α gets an apple"). Indeed, in case β does not take the apple and α does take the apple, we end up in a permitted state. However, it seems incorrect to say that α has a free choice in this scenario, normatively speaking. Indeed, the only choice α can make that ensures that we end up in an acceptable world, is by not taking the apple. One may even argue that in this case, if α really wants to take an apple, then α should discuss this with β so that they can coordinate their choices.*[12]

The example shows that weak permission, in combination with agency and/or the notion of choice as we formalized them here, cannot adequately account for our intuitions regarding FCP. In fact, they point to a more general problem with any such reduction: that it is permitted *for an agent α* to see to it that such-and-such is the case, cannot be reduced to what the agent does or chooses in any of the permissible states. This observation mirrors Horty's observations from Horty (2001, Section 3.4.2); we will return to it in the next section. Mind that this feature of FCP is strongly linked to our observation **(IV)**.

Ad (e). Likewise, one cannot just model FCP in terms of the weak permissibility of choices. Obviously, this cannot be done in the present framework, since $\mathsf{C}_\alpha(\varphi, \psi)$ and $\mathsf{P}(\mathsf{C}_\alpha(\varphi, \psi))$ are equivalent in it.[13] But even in purely informal terms, such an analysis does not seem to make sense.

To see why, suppose α is permitted to carry a loaded gun when walking on the street. Suppose moreover that α runs into an innocent person β. Let us assume that at this point, one of α's options is to kill β (suppose, for the sake of argument, that α is perfectly aware that he can do this, that it is just a matter of taking the gun, aiming it, and pulling the trigger). Clearly, α does not have a FCP to kill β or not; killing β is simply impermissible. However,

[12] Alternatively, one might feel that *conditional on* β's not taking an apple, α does have a FCP, or that in the worlds where β actually takes no apple, α has this FCP. We will not go into this argument in detail, but briefly point out in Section 5.2 how such conditional or world-relative FCPs can be modeled according to our approach.

[13] This is an immediate consequence of the fact that the abilities of α are the same at every point in a model.

that does not all of a sudden make it impermissible for α to carry a loaded gun, and hence to *have the choice* to kill β or not.

Ad (f)-(j). Note that among these definitions, (g) is the weakest.[14] We now show that even (g) is not necessary in order to have a FCP. The example can be seen as a precise illustration of our earlier observation that FCP does not imply sufficiency for permissibility (cf. Observation **(I)**).

Example 2 *α has a choice between four states: take nothing, take an apple and a napkin, take an apple and no napkin, take a napkin and no apple. We assume moreover that there is exactly one state corresponding to each of these choices. The only states that are permissible are the ones where α takes nothing, or where α takes an apple and a napkin. It seems intuitive to claim in this example that α is free to choose between taking an apple or not. Likewise, "you are free to choose between taking an apple and a pear, but in case you take a pear, you have to take a napkin as well" seems perfectly allright.*

5 Free Choice Permission

5.1 Permitted to see to

The above examples and discussion suggest an alternative account of FCP that is directly based on deontic STIT logic (Horty, 2001). In this paper we present a simplified, a-temporal version of it. We first define a deontic preference relation on arbitrary sets of states $X, X' \subseteq W$. That is, let $X \preceq X'$ iff for all $w \in X$, $w' \in X'$, w is permissible or w' is impermissible.[15]

Where α is given, we now define a more fine-grained preference relation between between choices X, X' of α. To do so, we need to introduce the auxiliary definition of all choices of a group of agents $A \subseteq Agt$. That is, where A consists of exactly n distinct agents $\alpha_1, \ldots, \alpha_n$, we let $\mathcal{C}_A =_{\mathsf{df}} \{X_1 \cap \ldots \cap X_n \mid X_1 \in \mathcal{C}_{\alpha_1}, \ldots, X_n \in \mathcal{C}_{\alpha_n}\}$. Note that, in view of the *independence of agents*-condition, $\mathcal{C}_A$ is a partition of W. As before, we let $\mathcal{C}_A(w)$ denote the (unique) member X of $\mathcal{C}_A$ such that $X \in A$".

We define:

[14] To see why, note that in general, if φ is stronger than ψ, then $\mathsf{P}^s\psi$ implies $\mathsf{P}^s\varphi$.

[15] If we use a more refined semantics, e.g. one where each world is assigned a value within the interval $[0, 1]$, then obviously also the preference relation $\preceq$ can be further refined. See (Horty, 2001) for how this would work.

$X \sqsubseteq X'$ iff for every $Y \in \mathcal{C}_{Agt-\{\alpha\}}, X \cap Y \preceq X' \cap Y$

In other words, option X is at least as good as option X' iff whatever the other agents do, the world we end up with when X is chosen will always be at least as good as the world we end up with when X' is chosen.

Let $\sqsubset$ be the strict counterpart of $\sqsubseteq$, defined in the usual way. One may now define a set of permissible choices of α as follows:

$X \in \mathcal{C}_\alpha$ is *permissible* iff there is no $X' \in \mathcal{C}_\alpha$ such that $X' \sqsubset X$

Note that if $\mathcal{C}_\alpha$ is finite, then there will be always at least one permissible $X \in \mathcal{C}_\alpha$ (relative to a given frame). In general there might be infinite descending chains of "ever better actions". In such cases one can refine the definition of permissible choices as follows: X is permissible iff there is no other choice X' that is itself $\sqsubset$-minimal within $\mathcal{C}_\alpha$, and for which $X \sqsubset X'$.[16]

We now introduce a new operator P_α for "α is permitted to see to it that", with the following semantic clause:

(SC5) $M, w \models \mathsf{P}_\alpha\varphi$ iff there is a permissible choice X of α such that $X \subseteq |\varphi|^M$.

Thus, $\mathsf{P}_\alpha\varphi$ means that there is a choice X of α which implies that φ is guaranteed, and there is no choice Y of α that is strictly better than X.

5.2 Relation to Horty's Analysis of Ought to Do

Using Horty's ought-operator O_α for "dominance act utilitarianism" (Chapter 4 of (Horty, 2001)), one can define our P_α as $\neg\mathsf{O}_\alpha\neg[\alpha]$. That is, $\mathsf{P}_\alpha\varphi$ is true iff there is some "minimal" (in our terms, strongly permissible) choice that guarantees φ, iff it is not the case that every optimal choice of α is such that it is not the case that α sees to it that φ.

Following Horty, one may also define variants of these operators. We will not go into detail here but just explain the main idea for two of them. First, one can consider conditional variants. In our Example 1, a conditional operator allows one to express such things as "given that α does not take an apple, β is permitted to take an apple." This can be done by defining the preference relation over β's choices in terms of their intersection with the

[16] See e.g. (Van De Putte & Straßer, 2014) for an elaborate discussion of such constructions that deal with non-smooth preference relations.

truth set of the antecedent of the conditional, i.c. that α does not take an apple. See (Horty, 2001, Chapter 5, Section 1) for the formal details.

Second, one can define operators for permission based on Horty's analysis of "orthodox act utilitariansm" (Horty, 2001, Chapter 5, Section 4). The main difference here is that what the agent is permitted to do depends on the specific point of evaluation $w \in W$. That is, we compare two actions X, X' of α by looking at the intersection of both with the set $\mathcal{C}_{Agt-\{\alpha\}}(w)$. So for instance, looking again at Example 1, in the worlds where β takes no apple, α is permitted to take an apple.

5.3 Free Choice Permission

So far we have only considered a unary operator that allows us to express what an agent α is permitted to do. Once there, we apply a by now familiar trick in order to express (binary) free choice permission:

$$\mathsf{F}_\alpha(\varphi, \psi) =_{\mathsf{df}} \mathsf{P}_\alpha(\varphi \wedge \neg\psi) \wedge \mathsf{P}_\alpha(\psi \wedge \neg\varphi)$$

In words, you have the free choice between φ and ψ (in a particular deliberative situation) if and only if you are permitted to see to it that $\varphi \wedge \neg\psi$ and you are permitted to see to it that $\psi \wedge \neg\varphi$.

This implies that you are able to guarantee each of these states of affairs, and hence also that they are distinct. It does however not imply that φ and ψ are disjoint, or that their conjunction is not permitted.

It can be easily checked that this definition accomodates each of our observations **(I)**-**(V)**:

*Ad **(I)*** Immediate in view of the existential quantification in the semantic clause for P_α. There may well be another *impermissible* choice X of α which also guarantees either φ or ψ.

*Ad **(II)*** Suppose that $\mathsf{F}_\alpha(\varphi, \psi)$ holds in a model; we prove that $\mathsf{P}\varphi$. By the definition of F_α, $\mathsf{P}_\alpha(\varphi \wedge \neg\psi)$. Hence by the semantic clause for $\mathsf{P}\alpha$, there is a permissible action X of α such that $X \subseteq \|\varphi\|^M$. We know that $X \cap P \neq \emptyset$ – otherwise X would not be permissible[17] – and hence $\|\varphi\|^M \cap P \neq \emptyset$. But this means exactly that $\mathsf{P}\varphi$ holds in the model.

*Ad **(III)*** In the model that corresponds to Example 1, $\mathsf{P}(apple_\alpha)$ and $\mathsf{P}(\neg apple_\alpha)$ are both true, but $\mathsf{F}_\alpha(apple_\alpha, \neg apple_\alpha)$ fails, since $\mathsf{P}_\alpha(apple_\alpha)$ fails.

[17]For this step, it is crucial that $P \neq \emptyset$.

*Ad **(IV)*** Immediate in view of the way F_α is defined. Note that we can easily accomodate cases where $\mathsf{F}_\alpha(\varphi,\psi)$ holds, but $\mathsf{F}_\beta(\varphi,\psi)$ fails.[18]
*Ad **(V)*** Immediate in view of the definition of $\mathsf{F}_\alpha(\varphi,\psi)$ and the semantic clauses of P_α, $\Diamond$, and $[\alpha]$.

So although F implies weak permission, the latter does not allow us to define F, in line with our examples from Section 4. Interestingly, we have the following relation between FCP and strong permission:

$$\mathsf{P}^s(\varphi \wedge \neg\psi), \mathsf{P}^s(\psi \wedge \neg\varphi), \mathsf{C}_\alpha(\varphi,\psi) \vdash \mathsf{F}_\alpha(\varphi,\psi) \tag{1}$$

The converse implication fails, as it should in view of our earlier discussion. So although FCP does not reduce to sufficiency for permissibility, the former can still be derived from the latter whenever the agent has the corresponding abilities.

6 Conclusion and outlook

In the preceding, we argued for a new approach to free choice permission, in which the agents and their choices are represented explicitly within the models and formal language. We showed that this approach matches each of our basic observations concerning FCP, in contrast to existing formal accounts.

A lot of open issues remain; we mention just three here. First and foremost, it remains to be seen if and how the resulting logic (including the operators for weak and strong permission, P and P^s) can be axiomatized.[19] Second, we mentioned some variants of Horty's deontic STIT logic, whose application to FCP is in need of further consideration. Third and last, we deliberately restricted the focus to a very specific and simple notion of FCP in this paper; one next step is to ask whether we can also get a grip on more "loose" or complex types of free choice permission which do not refer to one specific deliberative context.

[18] In fact, whenever $\alpha \neq \beta$, $\mathsf{F}_\alpha(\varphi,\psi) \supset \neg\mathsf{F}_\beta(\varphi,\psi)$ is a theorem of our logic. This follows from the fact that $\mathsf{P}_\alpha\varphi$ entails $\Diamond[\alpha]\varphi$ and the independence of agents-condition, cf. Section 4.

[19] As we noted above, the fragment of our logic without P and P^s coincides with the a-temporal fragment of (one of) the logic(s) of *ought to do* in (Horty, 2001). An axiomatization for these logics can be found in (Murakami, 2005).

References

Anglberger, A. J., Faroldi, F., & Korbmacher, J. (2016). An exact truthmaker semantics for permission and obligation. In O. Roy, A. Tamminga, & M. Willer (Eds.), *Deontic logic and normative systems* (pp. 16–31). College Publications.

Belnap, N., Perloff, M., Xu, M., & Bartha, P. (2001). *Facing the future: Agents and choice in our indeterminist world*. Oxford University Press.

Fine, K. (2016). Compliance and Command I and II. *Unpublished manuscript*. (`https://nyu.academia.edu/KitFine`)

Giordani, A., & Canavotto, I. (2016). Basic action deontic logic. In O. Roy, A. Tamminga, & M. Willer (Eds.), *Deontic logic and normative systems* (pp. 80–92). College Publications.

Hansson, S. O. (2013). The varieties of permission. In D. Gabbay, J. Horty, X. Parent, R. van de Meyden, & L. van der Torre (Eds.), *Handbook of deontic logic and normative systems* (Vol. 1, pp. 195–240). College Publications.

Horty, J. F. (2001). *Agency and deontic logic*. Oxford University Press.

Makinson, D. (1984). Stenius' Approach to Disjunctive Permission. *Theoria*, *50*(2-3), 138–147.

McNamara, P. (1996). Must I do what I ought? (or will the least I can do do?). In M. A. Brown & J. Carmo (Eds.), *Deontic logic, agency and normative systems* (p. 154-173). Springer London.

Murakami, Y. (2005). Utilitarian deontic logic. In *Advances in modal logic* (Vol. 5, p. 211-230).

Sedlár, I. (2016). Generalized permissions and free choice. *Unpublished manuscript*.

Van De Putte, F. (2016). That will do: Logics of deontic necessity and sufficiency. *Erkenntnis*. (in print, published online)

Van De Putte, F., & Straßer, C. (2014). Preferential semantics using non-smooth preference relations. *Journal of Philosophical Logic*, *43*(5), 903–942.

Frederik Van De Putte
Ghent University
Belgium
E-mail: `frederik.vandeputte@ugent.be`