Expected value of conditionals and expected utility: a probabilistic account of conditional evaluative constructions[1]

WooJin CHUNG — *New York University*

Abstract. This paper develops a probabilistic account of the Korean conditional evaluative construction that conveys obligation, which is expressed in terms of a conditional and an evaluative predicate. Assuming that conditionals denote the degree of support for the consequent given the antecedent and that evaluative predicates are measure functions that return the utility value of a given world argument, the compositional semantics of the construction at hand gives rise to an expected utility-based semantics of deontic modality. I point out its relevance to decision theory and further offer a principled solution the problem of supererogation.

Keywords: modality, conditionals, expected utility, supererogation.

1. Introduction

This paper proposes that the Korean construction that is used to convey obligation (cf. English *must*, *should*, and *ought*) suggests to maximize expected utility, as proposed in the literature on decision theory (Gibbard and Harper, 1978; Lewis, 1981). Korean does not make use of an auxiliary or a verb to express deontic concepts. Rather, such concepts are expressed in terms of a conditional and an evaluative predicate. For example, English 'ought *p*' effectively translates to 'only if *p*, good' in Korean, as exemplified below:[2]

(1) John-un Aleppo-ey ka-**ya** **toy**-n-ta.
John-TOP Aleppo-to go-**only.if** **GOOD**-PRES-DECL
'John ought to go to Aleppo.'
'(Lit.) Only if John (were to) go to Aleppo, good.'

Adopting Kaufmann's (2017) terminology for similar constructions in Japanese, I will refer to these constructions as *conditional evaluative constructions*. Assuming a weak version of semantic uniformity, Chung (2019) hypothesizes that the above conditional evaluative construction is a transparent version of the corresponding weak modal necessity expressions in other languages. Chung provides a compositional account of the construction based on Kratzer's (1981b) premise semantics and further claims that it offers a principled solution to the Professor Procrastinate puzzle (Jackson and Pargetter, 1986), fully explaining the presence of the order effect pointed out by von Fintel (2012).

While in agreement with Chung (2019) that Korean conditional evaluative constructions can and should receive a compositional account, this paper entertains an entirely different set of assumptions regarding the semantics of conditionals and evaluative predicates: (i) conditionals denote the degree of support for the consequent, given the antecedent and certain relevant facts of the world of evaluation, and (ii) the evaluative predicate *toy* 'GOOD' is a function of worlds that returns the utility value of the world argument. The resulting semantics suggests that the

[1]I gratefully acknowledge comments and suggestions from Sam Alxatib, Chris Barker, Lucas Champollion, Michael Franke, Magdalena Kaufmann, Stefan Kaufmann, Daniel Lassiter, Salvador Mascarenhas, Anna Szabolcsi, Christopher Tancredi, and audiences at SuB 24. All errors are mine.

[2]Korean conditionals do no require a subjunctive morphology to express counterfactuality. It is possible that (1) receives a counterfactual reading or an indicative reading, depending on the context.

Proceedings of Sinn und Bedeutung 24, vol. 1, pp.116–125. Osnabrück University.

expected utility of the prejacent is significantly higher than that of its alternatives. One notable characteristic of the Korean conditional evaluative construction is that it not only compares the expected utility of the prejacent and its alternatives but also takes into consideration the relative degree to which they differ. I suggest that this is the key to resolving the problem of supererogation—the problem of whether a theory of deontic modality based on (comparative) goodness can distinguish duties from supererogatory acts.

In what follows, I will first introduce expected value-based accounts of conditionals. I then offer a measure function analysis of the Korean evaluative predicate *toy* 'GOOD'. I show that a compositional analysis of the Korean conditional evaluative construction in (1) derives an expected utility-based semantics. I will focus on the counterfactual interpretation of the construction, while remaining agnostic on whether an indicative reading is available. I would like to note that the counterfactual reading of the construction relates to causal decision theory, whereas the indicative reading of it makes connection to evidential decision theory.

In the latter half of the paper, I will turn to the problem of supererogation which poses a non-trivial issue for existing theories of deontic modality. I show that the proposed semantics provides room for distinguishing duties from supererogatory acts, while maintaining the cherished relation between *ought* and the conception of goodness.

2. Expected valued-based accounts of conditionals

There are at least two lines of research that hypothesize that the interpretation of a conditional involves expected value calculation. In computer science, Pearl (2000, 2011, 2013) developed a theory of conditionals that involves manipulation of causal graphs, which are used to determine which propositions to additionally condition on in calculating the conditional probability of the consequent given the antecedent. In linguistics and philosophy, there have been attempts to define the assertability or value of a conditional in terms of its corresponding conditional probability (Jeffrey, 1964; Adams, 1965; Stalnaker, 1970; Lewis, 1976; Jeffrey and Edgington, 1991; Kaufmann, 2005; Douven, 2008).[3] In its simplest incarnation, the assertability/value of a conditional is defined as the conditional probability of the consequent given the antecedent. Following Kaufmann (2005), I will additionally incorporate Pearl's (2000) notion of causality which allows the theory to properly handle counterfactuals and embedded conditionals.

I assume that the value of a proposition-denoting expression is either 0 (false) or 1 (true). On the other hand, conditionals denote the degree of support for the consequent, given (i) the antecedent and (ii) facts that are causally independent of the antecedent.[4] Causal independence is interpreted with respect to (i) Φ: a set of causally relevant propositions singled out from the set of all propositions and (ii) $\prec$: a strict partial order which informally read as "affects the expectation of". The pair $\langle \Phi, \prec \rangle$ uniquely determines a causal graph that characterizes the

[3]There is no consensus on whether the value of a conditional is defined in terms of conditional probability. For instance, Stalnaker's (1970) objective was to develop a theory of conditionals that validates the hypothesis the *probability* of a conditional is its corresponding conditional probability.

[4]Kaufmann's original analysis takes a time-related parameter into account. Specifically, the evaluation procedure additionally conditions on facts that are settled at the time of evaluation. To my understanding, the settled facts only affect the probability distribution over worlds, and omitting this parameter amounts to assuming that the probability distribution remains constant throughout time. For simplicity, I will omit this parameter and focus on how the compositional analysis of the Korean conditional evaluative construction connects to the decision theory literature.

causes and effects.

(2) Causal independence (Kaufmann, 2005)
Given a causal structure $\langle \Phi, \prec \rangle$, for all $p, p' \in \Phi$: p' is *causally independent* of p iff $p \prec p'$.

The degree of support is defined in terms of expected value: it is the expected value of the consequent given the antecedent and causally independent facts, as fleshed out in (3). By the definition of expected value, if the consequent is a proposition-denoting expression (i.e., its value is either 0 or 1), then the value of a conditional is a real number within the interval [0, 1]. A value close to 1 represents a high degree of support for the consequent, whereas a value close to 0 indicates a low degree of support. Due to their range, conditionals with a proposition-denoting consequent can be interpreted as the conditional probability of the consequent given the antecedent.

(3) Expected value-based analysis of conditionals (cf. Kaufmann, 2005)
$[\![p \mathrel{\Box\!\!\rightarrow} q]\!] = \mathbb{E}[q \mid p, c_1, ..., c_i] = \sum_j q(w_j) * Pr(\{w_j\} \mid p, c_1, ..., c_i)$
where c_1, ..., c_i are causally independent facts of p and $w_j \in \cap\{p, c_1, ..., c_i\}$

While I have shown that probability measures can be obtained as a special case of expected value calculation (i.e., when the consequent of a conditional is a proposition-denoting expression), the analysis in (3) does not require that the consequent of a conditional is proposition-denoting. In the following section, I will take advantage of this underspecification and claim that conditional evaluative constructions do not denote a probability measure but rather an expected utility one.

3. Deriving an expected utility-based semantics from scratch

As for the analysis of the evaluative predicate *toy* 'GOOD', I propose that it is a measure function that takes a world and returns the degree of goodness (i.e., utility value) of the world argument.[5]

(4) A measure function analysis of *toy* 'GOOD'[6]
$[\![\text{GOOD}]\!] = \lambda w'.\mu_{good}(w')$,
where $\mu_{good}(w')$ returns the utility value of w'

Recall that Korean utilizes an *-(e)ya* 'only if' conditional to convey obligation. I will first leave out the contribution of the *only* component and show how 'if p, then GOOD' is interpreted. The formula in (5) is derived from (3) by simply replacing the consequent q with the evaluative predicate *toy* 'GOOD'. The upshot is that the formula calculates the expected value of the measure function μ_{toy} conditioned on the antecedent p and causally independent facts. In other words, the conditional denotes the expected utility of the counterfactual p-worlds.

(5) $[\![p \mathrel{\Box\!\!\rightarrow} \text{GOOD}]\!] = \mathbb{E}[\mu_{good} \mid p, c_1, ..., c_i] = \sum_j \mu_{good}(w_j) * Pr(\{w_j \mid p, c_1, ..., c_i\})$,
where c_1, ..., c_i are causally independent facts of p and $w_j \in \cap\{p, c_1, ..., c_i\}$

[5]This is reminiscent of how Lassiter (2017) predicates goodness of worlds, although he eventually lifts the domain of assessment from worlds to propositions. What will be shown is that a probabilistic analysis of the Korean conditional evaluative construction replicates Lassiter's lifting operation due to the conditional semantics.

[6]The analysis presented as it is implies that values are not world-dependent and is likely a oversimplification of the matters. But since the main objective of this paper is to highlight the connection between Korean conditional evaluative constructions and decision theory, I will simplify the matters as long as it does not affect the main argument.

I will assume that a thresholding operation (Lassiter, 2017) is performed on (5) to map the degree representation to a bivalent representation. Specifically, if the expected utility of the counterfactual p-worlds is higher than the contextually determined threshold θ_{toy}, we can map it to true (1). If the expected utility is less than or equal to θ_{toy}, we map it to false (0). So if we assume that the thresholding operation is invoked by default, (5) conveys that the counterfactual p-worlds have a sufficiently high expected utility.

As for the threshold θ_{toy}, I assume along with Lassiter (2017) that setting its value is guided by the contextually salient set of alternatives. Specifically, the guideline is the expected utility of the union of the alternatives. So for instance, if the contextually salient set of alternatives is $\{p, \neg p\}$, the corresponding baseline is the expected utility of the union of the counterfactual p-worlds and the counterfactual $\neg p$-worlds. The intuition behind this is that *good* conveys '(somewhat) better than average', and it is also in accordance with how other relatives adjectives (*tall*, *long*, *happy*) behave.

A complete analysis of Korean obligation can be derived from the above conditional semantics by exhaustifying it. As for the semantics of the exhaustification operator, I will posit a simplified variant of Rooth (1992), such that the exhaustification operator takes a proposition and negates each of the alternatives. Applying the exhaustification operator on top of (5) yields the semantics in (6). What is additionally conveyed due to exhaustification is that for every alternative r to p, $[\![r \;\square\!\!\rightarrow \textsc{good}]\!]^w$ returns a value that is less than or equal to the threshold θ_{toy} (i.e., not sufficiently high). The expected utility of the counterfactual r-worlds is not sufficiently high, whereas that of the counterfactual p-worlds is sufficiently high. An intuitive way of rephrasing this is to say that the expected utility of the counterfactual p-worlds *stands out* among the alternatives.

(6) An expected value-based analysis of Korean obligation
$[\![\text{only}]\!]([\![p \;\square\!\!\rightarrow \textsc{good}]\!])$
$= \mathbb{E}[\mu_{good} \mid p, c_1, ..., c_i] > \theta_{toy} \wedge \forall r \in Alt(p) \text{ s.t. } r \neq p : \mathbb{E}[\mu_{good} \mid r, c_1, ..., c_i] \leq \theta_{toy}$,
where c_1, ..., c_i are causally independent facts of p

There is an interesting connection between the above semantics and causal decision theory. Causal decision theory partitions the set of worlds into act-independent states s_i and calculates the expected utility of an act p by summing over the product of (i) the probability that s_i would obtain if p were the case (i.e., $Pr(p \;\square\!\!\rightarrow s_i)$) and (ii) the utility value of the outcome jointly determined by the act p and the state s_i (i.e., $o[p, s_i]$).[7] This amounts to calculating the expected utility of the counterfactual p-worlds.

(7) Causal expected utility of p (Gibbard and Harper, 1978)
$\text{EU}_{\text{CDT}}(p) = \Sigma_i Pr(p \;\square\!\!\rightarrow s_i) * u(o[p, s_i])$

Causal decision theory compares the causal expected utility of each available act to determine the optimal choice. Since (6) conveys that the counterfactual prejacent-worlds have the best expected utility (∵ only the counterfactual p-worlds have the expected utility higher than the threshold), it makes the same recommendation.

There is one important difference, however, due to the thresholding operation. The proposed

[7] Although debatable, I will assume along with Gibbard and Harper (1978) that acts and outcomes are represented as propositions.

semantics requires that the expected utility of the counterfactual prejacent-worlds is considerably higher than that of the counterfactual alternative-to-the-prejacent-worlds. If one of the alternatives has a slightly lower expected utility than the prejacent but the expected utility is still higher than the threshold θ_{toy}, we would not be able to conclude that the prejacent is obligatory. On the other hand, causal decision theory does not impose this requirement. In the following sections, I show that this additional requirement is the key to resolving the problem of supererogation.

4. Supererogation and the good-ought tie-up

Supererogation refers to the class of acts that go "beyond the call of duty" (Heyd, 1982). Such acts are not obligatory, although it would be good to bring them about. Although there is no clear consensus as to what technically qualifies as supererogatory, informally, small acts of favor, acts of heroism, self-sacrifice, politeness, and consideration are mentioned as typical instances of supererogation. Among ethical theorists, discussions on supererogation have centered around what has been called *the paradox of supererogation*: given that supererogatory acts are morally good to bring about, why aren't they required? Some deny the existence of supererogation and dictate that whatever is good is required (Moore, 1948; New, 1974; Feldman, 1986; Pybus, 1982; Crisp, 2013). Others affirm the existence and characterize it in terms of the intrinsic value of beneficent intentions or altruism (Heyd, 1982; Zimmerman, 1996), or in terms of the limited cost-effectiveness of sanctioning the agent who has failed to fulfill a given duty (Richards, 1971; Cohen, 2015).

While much attention has been paid to the problem of supererogation in metaethics, it has received little interest among linguists. However, the potential impact it has on linguistic theories of deontic modality is not negligible. The following scenario is originally due to Lassiter (2017), slightly modified to better manifest the problem:

> *Suppose that John ought to visit an ailing friend, say Mary. Suppose also that it would be even better, when John visits, to cook Mary dinner. Does it then follow that John ought to visit and cook Mary dinner? No doubt in some cases he should, but this is not a semantic fact: it is clearly possible that the following could hold as well.*

(8) There are many things John can do for Mary. He can make a visit, cook dinner, and so on. In view of Mary's well-being...
  a. John should/ought to visit Mary.
  b. Visiting and cooking dinner is better than visiting and not cooking dinner.
  c. However, cooking dinner is strictly optional: it's not the case that you should/ought to visit *and* cook dinner.

In the above scenario, cooking for Mary is beyond what is required, and in this sense it is supererogatory. We do not want to validate an inference from 'ought **visit**' to 'ought **visit** $\wedge$ **cook**'. While previous discussions on supererogation focused on the interpretation of *ought*, we can also observe the same invalid inference pattern for *has to*: even though John has to visit Mary, it does not necessarily follow that John has to visit Mary and cook for her.

Lassiter (2017) points out that the above scenario poses a non-trivial challenge to Lassiter's (2011) account of *ought* based on expected utility (EU), which is briefly summarized in (9).

The formula states that 'ought *p*' is true if and only if the expected utility of *p* is greater than the threshold value θ_{ought} determined by the lexical semantics of *ought* in interaction with the context.

(9) Lassiter's (2011) EU-based account
$[\![\text{ ought }]\!] = \lambda p.\text{EU}(p) > \theta_{ought}$

Given the above definition, 'ought **visit**' asserts that the expected utility of **visit** is greater than θ_{ought}. The problem is that the expected utility of **visit** ∧ **cook** is always greater than or equal to that of **visit** in the given scenario. The proof goes as follows: Since it is assumed that (8b) is true, the expected utility of **visit** ∧ **cook** is greater than that of **visit** ∧ ¬**cook**. And by definition, the expected utility of a disjunction falls between the expected utility of the two disjuncts, so the expected utility of **visit**, which is a disjunction of **visit** ∧ **cook** and **visit** ∧ ¬**cook**, is less than the expected utility of **visit** ∧ **cook**. It then follows that the expected utility of **visit** ∧ **cook** is greater than θ_{ought}, and it is predicted that 'ought **visit** ∧ **cook**' is true.

(10) $[\![\text{ ought } \textbf{visit}]\!] = \text{EU}(\textbf{visit}) > \theta_{ought}$

(11) $\text{EU}(\textbf{visit} \wedge \textbf{cook}) > \text{EU}((\textbf{visit} \wedge \textbf{cook}) \cup (\textbf{visit} \wedge \neg\textbf{cook}))$
$= \text{EU}(\textbf{visit}) > \text{EU}(\textbf{visit} \wedge \neg\textbf{cook})$

(12) $\text{EU}(\textbf{visit} \wedge \textbf{cook}) > \text{EU}(\textbf{visit}) > \theta_{ought}$, therefore 'ought **visit** ∧ **cook** is true.

Due to the above issue, Lassiter (2017) takes a weaker stance and does not offer a truth condition of *ought*. Instead, he lists necessary conditions for an *ought* statement to be true. One of them is Sloman's principle which states that if 'ought *p*' is true, then the expected utility of *p* is greater than the expected utility of each of its alternatives. This condition cannot be a sufficient condition though, because had it been the case, we would be back to the analysis in (9) and would suffer from the problematic scenario.

(13) Sloman's principle
ought $p \rightarrow \text{EU}(p) > [\forall q \in Alt(p) : q \neq p \rightarrow \text{EU}(p) > \text{EU}(q)]$

The standard account of deontic modality due to Kratzer (1981a) also validates the problematic inference under certain realistic assumptions. It is standardly assumed, under the simplifying *Limit Assumption* (Lewis, 1973), that necessity modals assert that the modal prejacent is necessarily true in the best worlds.[8] In the case of deontic modals, the best worlds are identified by two conversational backgrounds supplied by the context, namely *circumstantial modal base* f and *deontic ordering source* g. The former takes a world of evaluation and returns the set of propositions that correspond to the relevant circumstances for the interpretation of the modal, and the latter takes a world of evaluation and returns the set of propositions that correspond to the ideals. Intersecting the former (i.e., $\cap f(w)$) yields the set of relevant worlds, and the deontic ordering source is utilized to select the deontically best worlds within the set. Specifically, $g(w)$ induces an ordering $\leq_{g(w)}$ (informally read as 'at least as good as') such that for any two worlds u and v, $u \leq_{g(w)} v$ if and only if the set of propositions in $g(w)$ that are true in v is a subset of the set of propositions in $g(w)$ that are true in u. The set of deontically best worlds consists of

[8] I will not make a distinction between strong necessity modals (e.g., *must*, *have to*) and weak necessity modals (e.g., *should*, *ought*), since the problem applies to both classes. See von Fintel and Iatridou (2008) and Silk (2018) for relevant discussions.

circumstantially relevant worlds that are at least as good as any of the circumstantially relevant worlds.

(14) Ordering $\leq_{g(w)}$ with respect to $g(w)$
For all $u, v \in W$, $u \leq_{g(w)} v$
iff $\{p : p \in g(w) \wedge p(v) = 1\} \subseteq \{p : p \in g(w) \wedge p(u) = 1\}$

Due to the fact that the standard account evaluates the truth of the prejacent only in the deontically best worlds, the presented scenario would be problematic if such worlds exclusively consisted of **visit** ∧ **cook**-worlds. Given the assumption that 'ought **visit**' is true, the deontically best worlds are **visit**-worlds. Moreover, since we are assuming that **visit** ∧ **cook** is better than **visit** ∧ ¬**cook**, any **visit** ∧ ¬**cook** would be outranked by a **visit** ∧ **cook**-world that minimally differs in the truth of **cook**. Consequently, the deontically best worlds are all **visit** ∧ **cook**-worlds, and it is predicted that 'ought **visit** ∧ **cook**' is true.

The formal argumentation can be given in two steps: (i) show that **cook** is a member of the deontic ordering source $g(w)$ if visiting and cooking is better than visiting and not cooking, and (ii) show that the deontically best worlds are all **visit** ∧ **cook**-worlds, given the assumption 'ought **cook**' is true. The first argument requires the assumption that **cook** does not contradict any of the ideals. While this is not entirely innocuous, I find it to be a plausible circumstance. Moreover, we would want to invalidate the problematic inference in such a circumstance. Now, suppose that **cook** $\notin g(w)$. By the definition of comparative possibility given in (i), for '**visit** ∧ **cook**' to be better than '**visit** ∧ ¬**cook**', there needs to be a **visit** ∧ **cook**-world $u \in \cap f(w)$ such that no **visit** ∧ ¬**cook**-world $v \in \cap f(w)$ is at least as good as u. But this cannot be the case because given the assumption that **cook** $\notin d(w)$, a world that minimally differs from u in that **cook** is false is at least as good as u. This leads to a contradiction, and it is concluded that **cook** $\in g(w)$.

(15) Comparative possibility (Kratzer, 1981a)
A proposition p is more possible than a proposition q in a world w in view of a modal base f and an ordering source g if, and only if, the following conditions are satisfied:
a. $\forall u \in \cap f(w) : [u \in q \rightarrow \exists v \in \cap f(w) : v \leq_{g(w)} u \wedge v \in p]$
b. $\exists u \in \cap f(w) : [u \in p \wedge \neg \exists v \in \cap f(w) : [v \in q \wedge v \leq_{g(w)} u]]$

For the proof of the second point, suppose that there exists a **visit** ∧ ¬**cook**-world, u, which is deontically best. A world v that minimally differs from u with respect to the truth of **cook** is strictly better than u (i.e., $v \leq_{g(w)} u$ and $u \nleq_{g(w)} v$), given that **cook** $\in g(w)$. By the definition of the best worlds, any two best worlds need to be equally good or should be incomparable. Therefore, u is not a deontically best world, and the set of deontically best worlds exclusively consists of **visit** ∧ **cook**-worlds.

One possible objection in favor of the standard account is to say that (8a) and (8b) make use of different ordering sources. What prevents me from further developing this line of thought at the moment is that I have no explanation for why the ordering source shifts so easily even in the presence of an explicit ordering source (i.e., the explicit reference to what John can do for Mary, and prefixing each sentence with 'in view of Mary's well-being'). It would be an interesting and outstanding challenge to pinpoint the source of the shift, and to delineate the shifting process. In this paper, I will focus on how the probabilistic account of the Korean conditional evaluative construction invalidates the problematic inference.

5. Case study

What I intend to show in this section is that certain value assignments *can* render 'John ought to visit Mary' true but falsify 'John ought to visit Mary and cook dinner' at the same time. Note that the issue in Lassiter's (2011) semantics was that *no* value assignment allows this possibility in the scenario depicted in section 4, and that the standard account suffers as well due to the fact that visiting and cooking is better than visiting and not cooking.

Let us begin with assessing 'ought **visit**'. The contextually salient set of alternatives is given as $\{ \textbf{visit}, \neg\textbf{visit} \}$. One possible set of value assignments for the two alternatives is provided in (16). The baseline for the threshold θ_{toy} has been computed under the assumption that the two alternatives are equally probable, but a minor shift in the threshold value would not make a difference. Because the expected utility of the counterfactual **visit**-worlds is greater than θ_{toy} and that of the counterfactual ¬**visit**-worlds is less than or equal to θ_{toy}, 'ought **visit**' is predicted as true.

(16) Possible value assignments:
$\mathbb{E}[\mu_{good} \mid \textbf{visit}, c_1, ..., c_i] = 50$
$\mathbb{E}[\mu_{good} \mid \neg\textbf{visit}, c_1, ..., c_i] = 0$
$\theta_{toy} = \mathbb{E}[\mu_{good} \mid \textbf{visit} \cup \neg\textbf{visit}, c_1, ..., c_i] = 25$

(17) $[\![\text{ ought } \textbf{visit}]\!] = \mathbb{E}[\mu_{good} \mid \textbf{visit}, c_1, ..., c_i] > \theta_{toy} \wedge \mathbb{E}[\mu_{good} \mid \neg\textbf{visit}, c_1, ..., c_i] \leq \theta_{toy}$,
where c_1, ..., c_i are causally independent facts of p

As for 'ought $\textbf{visit} \wedge \textbf{cook}$', I will assume that the contextually salient set of alternatives is $\{ \textbf{visit} \wedge \textbf{cook}, \textbf{visit} \wedge \neg\textbf{cook}, \neg\textbf{visit} \}$. Given that cooking is only a preference, I will assume that **cook** makes less contribution to the expected utility than **visit**. Assuming that cooking and not cooking are equally probable, a possible set of value assignments that is consistent with the one in (16) is given in (18), and 'ought $\textbf{visit} \wedge \textbf{cook}$' receives the analysis in (19) accordingly.

(18) Possible value assignments:
$\mathbb{E}[\mu_{good} \mid \textbf{visit} \wedge \textbf{cook}, c_1, ..., c_i] = 60$
$\mathbb{E}[\mu_{good} \mid \textbf{visit} \wedge \neg\textbf{cook}, c_1, ..., c_i] = 40$
$\mathbb{E}[\mu_{good} \mid \neg\textbf{visit}, c_1, ..., c_i] = 0$
$\theta_{toy} = \mathbb{E}[\mu_{good} \mid (\textbf{visit} \wedge \textbf{cook}) \cup (\textbf{visit} \wedge \neg\textbf{cook}) \cup (\neg\textbf{visit}), c_1, ..., c_i] = 25$

(19) $[\![\text{ ought } \textbf{visit} \wedge \textbf{cook}]\!]$
$= \mathbb{E}[\mu_{good} \mid \textbf{visit} \wedge \textbf{cook}, c_1, ..., c_i] > \theta_{toy} \wedge \mathbb{E}[\mu_{good} \mid \textbf{visit} \wedge \neg\textbf{cook}, c_1, ..., c_i] \leq \theta_{toy}$
$\wedge\, \mathbb{E}[\mu_{good} \mid \neg\textbf{visit}, c_1, ..., c_i] \leq \theta_{toy}$

In this configuration, 'ought $\textbf{visit} \wedge \textbf{cook}$' is false. Although the expected utility of the counterfactual $\textbf{visit} \wedge \textbf{cook}$-worlds is greater than the threshold, the expected utility of the counterfactual $\textbf{visit} \wedge \neg\textbf{cook}$-worlds is also greater than the threshold. Given the proposal, the truth of 'ought $\textbf{visit} \wedge \textbf{cook}$' requires that the the expected utility of the counterfactual $\textbf{visit} \wedge \neg\textbf{cook}$-worlds is less than or equal to the threshold, but this is not the case.

6. Conclusion

A probabilistic account of Korean deontic modality highlights an interesting connection between natural language and decision theory. Furthermore, the semantics adds a proviso that the expected utility of the prejacent-worlds is significantly higher than that of the alternatives. It allows one to distinguish duties from supererogatory acts, thus illustrating the usefulness of expected utility in studying normative discourse. Conversely, it also manifests the usefulness of natural language semantics for decision theory. It makes suggestions as to how the expected utilities should be compared for rational decision making.

As a concluding remark, I would like to emphasize that the proposed account does not invalidate the problematic inference from 'ought **visit**' to 'ought **visit** $\wedge$ **cook**' no matter what. For instance, if cooking contributed to the expected utility as much as visiting, the analysis would predict that 'ought **visit** $\wedge$ **cook**' is in fact true. This is not a weakness, but rather a preferred property. If cooking is so important, we are inclined to conclude that bringing it about is also required. We do want the semantics to validate the inference under such circumstances, but at the same time, the semantics should be flexible enough to invalidate it if cooking is indeed supererogatory.

References

Adams, E. (1965). The Logic of Conditionals. *Inquiry 8*(1-4), 166–197.

Chung, W. (2019). Decomposing Deontic Modality: Evidence from Korean. *Journal of Semantics 36*(4), 665–700.

Cohen, S. (2015). Forced Supererogation. *European Journal of Philosophy 23*(4), 1006–1024.

Crisp, R. (2013). Supererogation and Virtue. *Oxford Studies in Normative Ethics 3*, 13–34.

Douven, I. (2008). The evidential support theory of conditionals. *Synthese 164*(1), 19–44. Publisher: Springer.

Feldman, F. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*, Volume 35. Springer Science & Business Media.

Gibbard, A. and W. L. Harper (1978). Counterfactuals and two kinds of expected utility. In *Foundations and applications of decision theory*, pp. 125–162. Springer.

Heyd, D. (1982). *Supererogation*. Cambridge University Press.

Jackson, F. and R. Pargetter (1986). Oughts, Options, and Actualism. *The Philosophical Review 95*(2), 233–255.

Jeffrey, R. and D. Edgington (1991). Matter-of-fact conditionals. *Proceedings of the Aristotelian Society, Supplementary Volumes 65*, 161–209.

Jeffrey, R. C. (1964). 'If'. *Journal of Philosophy 61*, 702–703.

Kaufmann, M. (2017). What 'may' and 'must' may be in Japanese. In K. Funakoshi, S. Kawahara, and C. Tancredi (Eds.), *Japanese/Korean Linguistics 24*.

Kaufmann, S. (2005). Conditional predictions. *Linguistics and Philosophy 28*(2), 181–231.

Kratzer, A. (1981a). The Notional Category of Modality. In H.-J. Eikmeyer and H. Rieser (Eds.), *Words, Worlds, and Contexts, New Approaches to Word Semantics*, pp. 38–74. Berlin: Walter de Gruyter.

Kratzer, A. (1981b). Partition and Revision: The Semantics of Counterfactuals. *Journal of Philosophical Logic 10*, 201–216.

Lassiter, D. (2011). *Measurement and modality: The scalar basis of modal semantics*. Ph.D. thesis, New York University.

Lassiter, D. (2017). *Graded modality: Qualitative and quantitative perspectives*. Oxford University Press.

Lewis, D. (1973). *Counterfactuals*. Oxford: Oxford University Press.

Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. In *Ifs*, pp. 129–147. Springer.

Lewis, D. (1981). Why Ain'cha Rich?'. *Noûs 15*(3), 377–380.

Moore, G. E. (1948). *Principia ethica*. Cambridge University Press.

New, C. (1974). Saints, Heroes and Utilitarians. *Philosophy 49*(188), 179–189.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press.

Pearl, J. (2011). The algorithmization of counterfactuals. *Annals of Mathematics and Artificial Intelligence 61*(1), 29.

Pearl, J. (2013). Structural Counterfactuals: A Brief Introduction. *Cognitive Science 37*(6), 977–985.

Pybus, E. M. (1982). Saints and Heroes. *Philosophy 57*(220), 193–199.

Richards, D. A. J. (1971). *A Theory of Reasons for Action*. Oxford: Clarendon Press.

Rooth, M. (1992). A Theory of Focus Interpretation. *Natural Language Semantics 1*(1), 75–116.

Silk, A. (2018). Weak and Strong Necessity Modals: On Linguistic Means of Expressing 'A Primitive Concept OUGHT'. In B. Dunaway and D. Plunkett (Eds.), *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*.

Stalnaker, R. C. (1970). Probability and Conditionals. *Philosophy of Science 37*(1), 64–80.

von Fintel, K. (2012). The best we can (expect to) get? Challenges to the classic semantics for deontic modals. In *Central meeting of the American Philosophical Association*, Volume 17.

von Fintel, K. and S. Iatridou (2008). How to say ought in foreign: The composition of weak necessity modals. In *Time and modality*, pp. 115–141. Springer.

Zimmerman, M. J. (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.