Neutrality and Inference Patterns of the Deliberative Ought

Rory Harder

June 29, 2017

University of Toronto
rory.harder@mail.utoronto.ca

Outline

Framework

Framework

Extensional Semantics

- Assign referring expressions type e (names, DPs).
  $[\![Mary]\!] = m$
- Assign predicative expressions type et (verbs, nouns).
  $[\![runs]\!] = \lambda x_e.[x \textit{ runs}]$
- e is a set of individuals. $t = \{1, 0\}$. If σ and σ' are types, then $\sigma\sigma'$ is.
- $[\![every]\!] = \lambda p_{et}.\lambda q_{et}.[\forall x(p(x) \rightarrow q(x))]$.
- Basic semantic composition proceeds via function application (FA), and an expression of type t is derived for sentences.

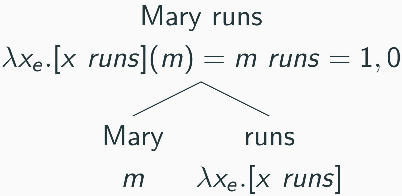

- FA: If α a node with set of daughters $\{\beta, \gamma\}$, and $[\![\beta]\!]$ is a function with $[\![\gamma]\!]$ in its domain, then $[\![\alpha]\!] = [\![\beta]\!]([\![\gamma]\!])$.

Framework

Intensionalize

- Problem: The meaning of a sentence contains more than just its (actual) truth-value, but also its truth-value at all different ways things could be.
- This is made vivid when they are embedded under modals.
- Both *Mary runs* and *2+2=4* are true. But *Necessarily, Mary runs* and *Necessarily, 2+2=4* differ in truth-value.
- Another example: Suppose *You steal from the rich* and *You give to the poor* are true. But surely *You ought to steal from the rich* and *You ought to give to the poor* can differ in truth-value.
- Simple LF: [ought$_i$ [you t$_i$ to steal from the rich]]

- Solution: relativize interpretation function to worlds.
  $[\![Mary]\!] = \lambda w_s.m$
  $[\![runs]\!] = \lambda x_{se}\lambda w_s.[x(w)\ runs\ in\ w]$
  $[\![Mary\ runs]\!] = \lambda w_s.[m\ runs\ in\ w]$
  $[\![Mary\ runs]\!]^{@} = m\ runs\ in\ @$
- Now we have type *st* (propositions) for sentences, which modals can operate on.
- How exactly do modals operate on embedded propositions? What additional information do they require from their contexts of use?

Classic Kratzer Semantics

- Kratzer (1977–): Modals are restricted quantifiers ($\forall, \exists$) over possible worlds, and which are restricted by contextually relevant information.
- All kinds of modals (root, epistemic) have this same *structure*, and their "flavour" is determined by the kind of information that is fed into the structure.
- Epistemics, for instance, take into account the beliefs/knowledge of the speaker (or other relevant entity).

- The deliberative or subjective *Ought* is the focus of this talk.
- Deliberative *Ought* $<$ Deontics $<$ Roots. Priority?
- It is what we use to advise ourselves or others on what to do when faced with a decision problem. It expresses the result of practical reasoning.
- Kratzer does not discuss this subsubtype explicitly, but here is how it would be treated in her semantics.

- The pieces of information we'll require are beliefs about the way the world is, and values about how the world should be.
- Both of these are treated as functions from worlds to sets of propositions.
- A modal base $f(w)$ is the set of ways things are believed to be in w. An ordering source $g(w)$ is the set of good-making features of worlds in w.
- Roughly, *Ought(A)* is true, at a world w, just in case all the worlds in $\bigcap f(w)$ that are top-ranked by $g(w)$ are worlds in which the proposition expressed by A is true. (All the ways you think things could be, which are optimal according to your values, are worlds in which A is true.)

How does the ordering source work?

- $w' \preceq_{g,w} w'' \Leftrightarrow: \{P \in g(w) | w' \in P\} \supseteq \{P \in g(w) | w'' \in P\}$
- $w' \prec_{g,w} w'' \Leftrightarrow: w' \preceq_{g,w} w'' \wedge \neg(w'' \preceq_{g,w} w')$
- $\mathcal{B}_{f,g,w} =: \{w' \in \bigcap f(w) | \neg\exists w'' \in f(w) \ s.t. \ w'' \prec_{g,w} w'\}$

The semantics

- $[\![Ought(A)]\!]^{f,g,w} = 1 \Leftrightarrow \forall w' \in \mathcal{B}_{f,g,w} : [\![A]\!]^{f,g,w'} = 1$
- iff $\mathcal{B}_{f,g,w} \subseteq [\![A]\!]^{f,g}$
- $[\![May(A)]\!]^{f,g,w} = 1 \Leftrightarrow \exists w' \in \mathcal{B}_{f,g,w} : [\![A]\!]^{f,g,w'} = 1$
- iff $\mathcal{B}_{f,g,w} \cap [\![A]\!]^{f,g} \neq \emptyset$
- (I present the lexical entries 'syncategorematically' to avoid the compositional complexity when the interpretation function is further relativized to modal bases and ordering sources.)

EV Semantics

EV Semantics

Motivation

- One of the main ways of assessing semantic theories is to check what inference patterns they (in)validate.
- The inference pattern profile of a semantics is tied up with its predictive power (modulo pragmatic factors).
- To assess this rigorously, we need a definition of truth and entailment (preservation of truth).

- Suppose a context of utterance determines particular values f, g, w (and others we might need). Sentence A is true at context c iff $[\![A]\!]^{f_c,g_c,w_c} = 1$.
- A entails B iff for all c: if $[\![A]\!]^{f_c,g_c,w_c} = 1$, then $[\![B]\!]^{f_c,g_c,w_c} = 1$.
- iff for all c: $[\![A]\!]^{f_c,g_c} \subseteq [\![B]\!]^{f_c,g_c}$.
- iff any situation in which the information expressed by A (at c) is true, is also a situation where the information expressed by B (at c) is true.

Unfortunate Facts about Kratzer's Semantics

- Inheritance: If $A \models B$, then $Ought(A) \models Ought(B)$.
- Conjunctive Inheritance: $Ought(A \wedge B) \models Ought(A)$.
- Disjunctive Inheritance: $Ought(A) \models Ought(A \vee B)$.
- Inheritance holds because if all the best worlds are A-worlds, and all the A-worlds are B-worlds, then all the best worlds are B-worlds.

Why are they Unfortunate?

- Professor Procrastinate: counterexample to conjunctive inheritance (Jackson & Pargetter 1986).
- *Ought(accept and write)*, but not *Ought(accept)*.
- Ross' Puzzle: counterexample to disjunctive inheritance (Ross 1941).
- *Ought(mail this letter)*, but not *Ought(mail it or burn it)*.
- EV semantics, as we will see, invalidates the inheritances in an explanatory way.

EV Semantics

Proposal

- Developed in Goble (1996), Cariani (2008), and Lassiter (2011–).
- Gloss: you ought to do something if its EV is higher than that of its alternatives.
- To determine the EV of a proposition, we need a probability function Pr defined over propositions, and a value function v defined over individual worlds. We also need a background set $\mathcal{A}$ of alternatives, which is a partition over possible worlds.

- The EV formula ‘lifts’ the valuing of worlds to a valuing of sets of worlds (propositions), in a way influenced by the relevant probability function.
- $EV_{Pr,v}(P) = \sum_{w \in P} v(w) Pr(\{w\}|P)$.
- $[\![Ought(A)]\!]^{Pr,v,\mathcal{A},w} = 1 \Leftrightarrow EV_{Pr,v}(\mathbf{A}) >> EV_{Pr,v}(Q)$, for all $Q \in \mathcal{A}$ s.t. $Q \cap \mathbf{A} = \emptyset$.

Fortunate Facts about EV Semantics

- If $EV(P) > EV(Q)$, then $EV(P) > EV(P \cup Q) > EV(Q)$.
- If $EV(P \cap Q) > EV(P \cap \bar{Q})$, then $EV(P \cap Q) > EV(P)$.
- These general facts explain the invalidity of the inheritances, and what's happening in Prof. Procrastinate and Ross' Puzzle.

Neutrality Criticisms

Neutrality Criticisms

The Nature of Neutrality

- Empirical Neutrality: ‘for every coherent practical theory P, if there are circumstances in which P entails a set of deontic sentences S, then the semantic theory ought to imply that S is consistent’ (Cariani 2016).
- Structural Neutrality: The structure of the semantic theory should not encode a particular decision rule, i.e. practical theory; the interpretation of deontic sentences should be sensitive to a variable decision rule parameter (Carr 2015; Charlow 2016, forthcoming).

- Empirical neutrality is motivated by the idea that a semantics should be able to properly predict the judgments of semantically competent speakers. A coherent practical theory should not be ruled out as a matter of the *meaning* of deontic language.
- Structural neutrality is a stronger claim, which is that the structure of the lexical entries for deontic modals should adequately reflect how semantic competence combines with "conceptual competence" to determine normative judgments.

Neutrality Criticisms

Against EV

- EV semantics clearly violates structural neutrality because it writes in a specific decision rule. EV is popular, yet not universal among semantically competent speakers.
- (A decision rule/practical theory is a way of processing information and values to arrive at recommendations/requirements for action).
- Demonstrating a violation of empirical neutrality involves finding a specific set of sentences in a specific circumstance.

- Miner's Puzzle (Forced Choice): 10 miner's are either stuck in shaft A or shaft B (50/50). If you block one shaft, it will save all miner's if they are in that one, but kill all miner's if they are in other. If you were to block neither, only one would die (but guaranteed to). You are forced though to pick one to block.
- (See Kolodny & MacFarlane 2010 for an important discussion of this puzzle).

- Suppose I'm a maximiner, i.e. what ought to be done is action with least bad worst outcome, no matter the probabilities.
- (1) You ought not block A and you ought not block B.
- (2) Even if it were more likely than not the miner's are in A, you ought not block A.
- EV semantics cannot predict consistency of (1) and (2), since the antecedent of (2) only shifts probability function, and holding values fixed this would mean EV of blocking A is higher than EV of blocking B.

- More formally: given Pr reflects 50/50 chance miner's are in A or B, find a v such that $EV(Block\ A) = EV(Block\ B)$. This makes (1) true.
- But if we shift Pr to assign higher chance to miner's being in A, then by EV formula: $EV(Block\ A) > EV(Block\ B)$. This is because the higher valued outcome of blocking A becomes more likely. So (2) cannot also be true.
- It is maybe *possible* to find a v that satisfies both, but it would be incredibly gerrymandered.

Neutrality Criticisms

Against Kratzer

- Structurally, there's an inherent 'weightiness' in EV, which other decision rules do not share.
- Kratzer's semantics, on the other hand, does not accommodate any weightiness, since the semantics is not sensitive to probabilities. (This is part of the reason for the problematic validities.)
- Relatedly, it encodes a form of maximax: if all worlds that are best (satisfy most members of g), and have some non-zero probability, are A-worlds, then $Ought(A)$ is true.

- Miner's Puzzle (High Probability, No Forced Choice): We're 99% sure they're in shaft A. We're able to pick neither.
- (3) We ought to block shaft A.
- Some of the best worlds are still block-B worlds, i.e. the ones where they are also in B. So Kratzer's semantics predicts (3) is false.
- $g = \{\{w |$ at least 10 miners saved in $w\}, \{w |$ at least 9 miners saved in $w\}, \ldots\}$.

- It may be possible to find a proper g, but, again, it would be gerrymandered. And it runs together first-order and rational preferences (the EV loophole would too).

Generalized Kratzer Semantics

- Developed in Carr (2015) and Charlow (2016, forthcoming). I present Carr's.
- Keep the modal base f, and add a probability function and value function (like EV semantics), and additionally a decision rule parameter δ.
- A decision rule is a function from decision problems $\pi = (f, Pr, v)$ to sets of propositions: those good-making features of worlds given the relevant information and values.

- $w' \preceq_{\delta,\pi,w} w'' \Leftrightarrow: \{P \in \delta(\pi) | w' \in P\} \supseteq \{P \in \delta(\pi) | w'' \in P\}$
- $\mathcal{B}_{\delta,\pi,w} =: \{w' \in \bigcap f(w) | \neg\exists w'' \in f(w) \ s.t. \ w'' \prec_{\delta,\pi,w} w'\}$
- $[\![Ought(A)]\!]^{\delta,\pi,w} = 1 \Leftrightarrow \forall w' \in \mathcal{B}_{\delta,\pi,w} : [\![A]\!]^{\delta,\pi,w'} = 1$
- iff $\mathcal{B}_{\delta,\pi,w} \subseteq [\![A]\!]^{\delta,\pi}$

- So we've generalized the Kratzer structure in such a way that the ordering of possibilities that constructs the best-worlds-set can take into account probabilities and cardinal values (but also optionally not).
- EV: $\delta(\pi) = \{\{w|\ Pr\text{-weighted } v\text{-value is maximized in } w\}\}$.
- Maximax: $\delta'(\pi) = \{\{w|\ v(w) \text{ is higher than } v \text{ of all other relevant } w'\}\}$.

REV Semantics

REV Semantics

Motivation

- Keeping the ‘superset of the best worlds’ structure in the semantics violates structural neutrality, no matter how complicated the construction of the best-worlds-set is.
- It retains a form of maximax: if the best worlds are all in the proposition expressed by A, that is sufficient for the truth of $Ought(A)$.
- This can be borne out empirically by noting how this Kratzer structure means the semantics validates the Inheritances, and then building a specific case around that.

- Miner's Puzzle (original): 50/50 chance they are in either shaft. You can pick to block one, or neither.
- I am an EV reasoner.
- (1) You ought to block neither.
- (2) It's not the case that you ought to block neither or shaft A. (You simply ought to block neither!)
- The generalized Kratzer semantics cannot predict consistency of (1) and (2), since by inheritance, (1) entails 'You ought to block neither or shaft A'.

- Solution: instead of starting with Kratzer semantics, and trying to make it versatile enough to capture, e.g., EV reasoning, let's start with EV and generalize it to capture other forms of reasoning (maximin, maximax).

REV Semantics

Proposal

- REV is developed in Buchak (2013), but not as a semantic proposal.
- In addition to the Pr, v, and $\mathcal{A}$ from EV semantics, we need a ‘risk’ parameter r.
- $r : [0, 1] \rightarrow [0, 1]$, and it modifies how the probabilities and values are aggregated to determine the REV of propositions. Hence, different values of r can be understood as different decision rules.

How to Calculate the REV of a Proposition

- $P = \{w_1, \ldots, w_n\}$.
- Stipulate $v(w_0) = 0$ and $v(w_0) \leq v(w_1) \leq \ldots \leq v(w_n)$.
- Consider sequence:
  $v(w_1) - v(w_0), v(w_2) - v(w_1), \ldots, v(w_n) - v(w_{n-1})$.
- $REV_{Pr,v,r}(P) = \sum_{j=1}^{n} r\Big(\sum_{i=j}^{n} Pr(\{w_i\}|P)\Big)(v(w_j) - v(w_{j-1}))$

- $[\![Ought(A)]\!]^{Pr,v,\mathcal{A},r,w} = 1 \Leftrightarrow REV_{Pr,v,r}(\mathbf{A}) \geq REV_{Pr,v,r}(Q)$, for all $Q \in \mathcal{A}$.
- Genuine EV: $r(x) = x$.
- Maximax: $r(x) = \{1 \text{ if } x > 0, 0 \text{ if } x = 0\}$.
- Maximin: $r(x) = \{0 \text{ if } x < 1, 1 \text{ if } x = 1\}$.
- $r(P_{w_1 \to w_n} = 1)v(w_1) + r(P_{w_2 \to w_n})(v(w_2) - v(w_1)) + r(P_{w_3 \to w_n})(v(w_3) - v(w_2)) + \ldots + r(P_{w_n})(v(w_n) - v(w_{n-1}))$

- Claims: The REV semantics satisfies the two kinds of neutralities better than the generalized Kratzer semantics.
- In Q&A, I can explain in more detail how the REV formula works, and how it's predictions compare with the generalized Kratzer semantics.
- Thanks!

References

Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.

Cariani, F. (2008). *The Semantics of 'Ought' and the Unity of Modal Discourse*. PhD Dissertation.

Cariani, F. (2016). Deontic modals and probabilities: One theory to rule them all? In Nate Charlow & Matthew Chrisman (eds.), *Deontic Modality*, Oxford University Press.

Carr, J. (2015). Subjective *Ought*. *Ergo 2*(27). 678-710.

Charlow, N. (2016). Decision theory: Yes! truth conditions: No! In Nate Charlow & Matthew Chrisman (eds.), *Deontic Modality*, Oxford University Press.

Charlow, N. (forthcoming). Decision-theoretic relativity in deontic modality. *Linguistics and Philosophy*.

Goble, L. (1996). Utilitarian deontic logic. *Philosophical Studies 82*. 317-357.

Heim, I. & Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell Publishers.

Jackson, F. & Pargetter, R. (1986). Oughts, options, and actualism. *The Philosophical Review 95*(2). 233-255.

Kolodny, N. & MacFarlane, J. (2010). Ifs and oughts. *The Journal of Philosophy, 107*(3). 115-143.

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

Lassiter, D. (2011). *Measurement and modality: The scalar basis of modal semantics*. PhD Dissertation.

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