

Context Probabilism

Seth Yalcin

University of California, Berkeley, USA
yalcin@berkeley.edu

Abstract. We investigate a basic probabilistic dynamic semantics for a fragment containing conditionals, probability operators, modals, and attitude verbs, with the aim of shedding light on the prospects for adding probabilistic structure to models of the conversational common ground.

Keywords: dynamic semantics, common ground, conditionals, epistemic modals, attitudes, probability operators, probability.

1 Introduction

Our story begins with a core notion of pragmatics:

Definition 1. It is COMMON GROUND that ϕ in a group if all members presuppose that ϕ , and it is common knowledge in the group that all members presuppose that ϕ .

Broadly following Stalnaker and Lewis building on Grice, we take it that the mutually understood proximal rational aim of speech acts is generally to change the common ground of the conversation—to update, if you like, the conversational scoreboard. Our more distal communicative objectives (of transferring knowledge, raising questions, convincing, misleading, etc.) are generally achieved by way of changing what is presupposed. Presupposition is a public attitude, in the sense that the presuppositional states of a set of interlocutors are correct just in case they match, having the same content and structure. Contexts in which agents are not presupposing the same things are defective, though the defect may never emerge explicitly. In theorizing we largely restrict attention to the nondefective case.

We do not assume that the common ground coincides with common belief. Generally, it does not. What one comes away from a conversation actually believing depends on subtleties about trust and authority, and about what the mutually understood import of the communicative exchange was. Such matters depends on the particular goals and interests of the interlocutors, and do not fall within the scope of linguistic pragmatics *per se*.

A formal pragmatics investigates (*inter alia*) possible models for the common ground, and investigates possible ways the common ground might be updated by agents who share knowledge of the semantics and pragmatic rules of a language. A natural course is to begin with a relatively unstructured model of the common

ground, adding structure just as required by the phenomena. Thus we might begin by taking the common ground to be an unstructured set of propositions, the presupposed propositions. Assertion (for instance) could then be modeled as a sort of proposal to add a proposition to the common ground. Still more austere we might, following Stalnaker, take the common ground to be a set of possible worlds, the possible worlds not yet eliminated by what is presupposed (what Stalnaker calls a *context set*). On this model propositions are identified with sets of possible worlds, and the propositions which are common ground are taken to be those true throughout the context set. The set of presupposed propositions is thus assumed to be closed under entailment. The effect of successful assertion, on Stalnaker’s view, is to eliminate possibilities incompatible with the proposition asserted from the common ground.

Various dynamic semantic systems can be seen as adding structure to a broadly Stalnakerian conception of the common ground. The file change semantics of [Heim 1982] famously added assignment functions and discourse referents; [Veltman 1996] added expectation patterns; [Hulstijn 1997] added partitions; and so on. My interest in the paper is in the question of adding probabilistic structure. I call the thesis I want to investigate *context probabilism*:

Context probabilism. The common ground of a conversation may have probabilistic structure.

I will suggest this thesis (broached initially in [Yalcin 2005]) has at least the following three motivations: first, adopting it paves the way for a plausible semantics and pragmatics of probability operators (*probably, likely, is as likely as*, etc.); second, it enables us, together with other assumptions, to capture a nontrivial connection between indicative conditionals and the corresponding conditional probabilities; third, it permits a model of linguistic communication congenial to a broadly Bayesian perspective. Further motivations tied to particular linguistic phenomena will accrue as we proceed.

There are a variety of conceivable ways of implementing context probabilism in a semantic-pragmatic theory. I elect in this paper to approach the issue from the perspective of a relatively ‘textbook’ dynamic semantic system for a propositional language, asking how that textbook system might be conservatively extended to incorporate context probabilism. We begin with a review of the textbook system.

2 A textbook propositional dynamic semantics

Assume a propositional language \mathcal{L} containing negation, conjunction, epistemic *might* (\diamond), an indicative conditional operator (\rightarrow), a belief operator (B), a knowledge operator (K) and a presupposition operator (∂).

Definition 2. A MODEL $\langle \mathcal{W}, \mathcal{I} \rangle$ is a pair of a space of worlds \mathcal{W} and a valuation function \mathcal{I} mapping propositional letters of \mathcal{L} to subsets s of \mathcal{W} (INFORMATION STATES).

Definition 3. A MODEL WITH ATTITUDES is a triple of a model \mathcal{M} together with a doxastic accessibility relation \mathcal{B}^w and an epistemic accessibility relation \mathcal{K}^w (construing these as functions from worlds to sets of worlds within $\mathcal{W}_{\mathcal{M}}$).

Definition 4. AN UPDATE FUNCTION $[\cdot]$ is a function from wffs of \mathcal{L} to functions from information states to information states (in some given model with attitudes) subject to the following constraints, where α is any propositional letter, ϕ and ψ are any wffs:

$$\begin{aligned}
s[\alpha] &= s \cap \mathcal{I}(\alpha) && \\
s[\neg\phi] &= s - s[\phi] && \text{[Heim 1983]} \\
s[\phi \wedge \psi] &= s[\phi][\psi] && \text{[Heim 1982], [Stalnaker 1974]} \\
s[\Diamond\phi] &= s \text{ iff } s[\phi] \neq \emptyset, \text{ else } \emptyset && \text{[Veltman 1996]} \\
s[\phi \rightarrow \psi] &= s \text{ iff } s[\phi] = s[\phi][\psi], \text{ else } \emptyset && \text{[Gillies 2004]} \\
s[B\phi] &= \{w \in s : \mathcal{B}^w[\phi] = \mathcal{B}^w\} && \text{[Heim 1992]} \\
s[\partial\phi] &= s \text{ iff } s[\phi] = s; \text{ else undefined} && \text{[Beaver 2001]} \\
s[K\phi] &= \{w \in s[\partial\phi] : \mathcal{K}^w[\phi] = \mathcal{K}^w\} &&
\end{aligned}$$

Definition 5. s ACCEPTS ϕ iff $s[\phi] = s$.

Definition 6. $\phi_1, \dots, \phi_n \models \psi$ iff $\forall s$: if s accepts each of ϕ_1, \dots, ϕ_n , s accepts ψ .

This supplies our starting point. I assume familiarity with ideas reflected in the above semantics. See the works cited above for introduction and some initial motivation (see also [Yalcin 2007]).

3 Discontents of the textbook semantics

We should like to extend this system so as to make room for probability operators. As it is, it is difficult to see how this might be achieved if, as seems intuitive, we have it that $\Delta\phi \models \Diamond\phi$ (abbreviating *probably* with Δ). For if $\Diamond\phi$ performs a ‘test’ on a state of information, so too should $\Delta\phi$; but what test would be suitable? It seems clear we need information states to incorporate further structure, structure which the test corresponding to $\Delta\phi$ can be sensitive to. We should also like to extend this system so as to supply some intelligible connection between the acceptability of an indicative conditional and its corresponding conditional probability, given empirical evidence that the two are highly correlated (see [Douven and Verbrugge 2010] and references therein).

Aside from these limitations of coverage, the textbook semantics also suffers from the following two quite nontrivial predictive failures, failures our probabilistic upgrade will eventually overcome.

Negation problem. According to the semantics, (1) \models (2):

1. It is not the case that John believes it might be raining. ($\neg B\Diamond\phi$)
2. John believes that it is not raining. ($B\neg\phi$)

This is not plausible in general. One can fail to believe it might be that ϕ without believing $\neg\phi$. The failure might, for instance, be the result of not being appropriately sensitive to the question whether ϕ . (This worry is discussed in detail in [Yalcin 2011]; see also [Willer 2010].)

Defective logic of knowledge and belief. Suppose it is raining. From this, it of course follows that it is compatible with what John knows that it is raining. Hence it follows, according to the semantics above, that John knows it might be raining. This result is absurd.

We get a yet greater absurdity when we add the plausible assumption that knowledge entails belief (understood as $K\phi \models B\phi$). When we add this assumption, we get the result that belief in falsehoods is not possible. As an example, suppose (for *reductio*) that it is raining, but that John believes it is not raining. By the reasoning above, John knows it might be raining. By the assumption that knowledge entails belief, John believes it might be raining. But this precludes, given the semantics, the possibility that John believes it is not raining.¹

4 Sharp context probabilism

To make room for probability operators and conditional probabilities, let us enrich the structure of the common ground. An obvious first idea is to embrace:

Sharp context probabilism. A conversational common ground can be modeled as a probability space.

We let probabilities spaces of a certain sort displace our earlier, non-probabilistic information states:

Definition 7. A SHARP INFORMATION STATE i is a pair of a set s , $s \subseteq W$ (call it the DOMAIN of i) and a probability function Pr on the elements of some Boolean algebra A of subsets of W such that: (i.) $Pr(s) = 1$; and (ii.) $Pr(p \cup q) = Pr(p) + Pr(q)$, for all disjoint p, q in A .

In essence, we are simply equipping our earlier information states with probability measures. Reinterpret \mathcal{B}^w and \mathcal{K}^w so that they correspond to sharp information states (given a world). Let our update function now be defined on sharp information states as follows:²

Definition 8. A PROBABILISTIC UPDATE FUNCTION $[\cdot]$ is a function from wffs of \mathcal{L} to functions from information states to information states subject to the following constraints, where α is any propositional letter, ϕ and ψ are any wffs:

$$\begin{aligned} i[\alpha] &= \langle s_i \cap \mathcal{I}(\alpha), Pr_i(x|s_i \cap \mathcal{I}(\alpha)) \rangle \\ i[\neg\phi] &= \langle s_i - s_{i[\phi]}, Pr_i(x|s_i - s_{i[\phi]}) \rangle \\ i[\phi \wedge \psi] &= i[\phi][\psi] \\ i[\diamond\phi] &= i \text{ iff } s_{i[\phi]} \neq \emptyset, \text{ else } \langle \emptyset, Pr_i(x|\emptyset) \rangle \end{aligned}$$

¹ I am indebted here to conversations with Stephen Yablo and Rohan Prince.

² Thanks here to Justin Bledin for essential suggestions.

$$\begin{aligned}
i[\Delta\phi] &= i \text{ iff } Pr_i(s_{i[\phi]}) > .5, \text{ else } \langle \emptyset, Pr_i(x|\emptyset) \rangle \\
i[\phi \rightarrow \psi] &= i \text{ iff } i[\phi] = i[\phi][\psi], \text{ else } \langle \emptyset, Pr_i(x|\emptyset) \rangle \\
i[B\phi] &= \langle \{w \in s_i : \mathcal{B}^w[\phi] = \mathcal{B}^w\}, Pr_i(x|\{w \in s_i : \mathcal{B}^w[\phi] = \mathcal{B}^w\}) \rangle \\
i[\partial\phi] &= i \text{ iff } i[\phi] = i; \text{ else undefined} \\
i[K\phi] &= \langle \{w \in s_{i[\partial\phi]} : \mathcal{K}^w[\phi] = \mathcal{K}^w\}, Pr_i(x|\{w \in s_{i[\partial\phi]} : \mathcal{K}^w[\phi] = \mathcal{K}^w\}) \rangle
\end{aligned}$$

Definition 9. A sharp information state i ACCEPTS ϕ iff $i[\phi] = i$.

Definition 10. $\phi_1, \dots, \phi_n \models \psi$ iff $\forall i$: if i accepts each of ϕ_1, \dots, ϕ_n , i accepts ψ .

(Definition 10 now supplants Definition 6.) In general, updates to the common ground involving ordinary factual information proceed by eliminating worlds from the domain of the context (as before) *and* by conditionalizing the probability measure. Probability operators now perform tests, akin to epistemic modals and conditionals. Whereas \diamond serves to perform a test on the domain of a sharp information state, Δ serves to perform a test on the measure. It is obvious how to add *is as likely as* (\succeq) to this system, representing it, too, as a test.

The present account addresses the two limitations of coverage noted for the textbook account. For the connection between indicatives and conditional probabilities, we can note that, for instance,

3. John believes that if the door is ajar, Bob is probably in his office.

is in the present framework a matter of whether John's credence in the proposition that Bob is in his office, conditional on the proposition that the door is ajar, is sufficiently high (above .5). In general, the system tells us that an indicative conditional is accepted relative to a sharp information state i just in case, roughly, the consequent is accepted relative to the state resulting from conditionalizing i on the antecedent.

As for the semantics of probability operators, the system delivers an array of desirable entailment predictions. I will just list a few: (i) $\Box\phi \models \Delta\phi \models \diamond\phi$; (ii) $\phi \models \Delta\phi$; (iii) $\neg\diamond\phi \models \neg\Delta\phi$; (iv) $\{\phi \rightarrow \psi, \Delta\phi\} \models \Delta\psi$; (v) $\{\phi \rightarrow \psi, \neg\Delta\psi\} \models \neg\Delta\phi$; (vi) $\Delta\phi \wedge \neg\phi \models \perp$. Adding \succeq , we could include (vii) $\phi \rightarrow \psi \models \psi \succeq \phi$; (viii) $\{\phi \succeq \psi, \Delta\psi\} \models \Delta\phi$.

It is natural to ask whether the resources of probability theory are strictly required to cover these and other inference patterns. For some further relevant discussion, see [Yalcin 2010], [Lassiter 2011]. Whether or not such resources are strictly required, the system is anyway perspicuous, and it offers a relatively conservative extension of the textbook picture. Where probabilities are not at issue, the additional probabilistic structure can safely be ignored, and one can work with the ordinary textbook account.

Indeed, the present system is *too* conservative an extension of the textbook picture, for it inherits the two predictive failures described in the last section 3 above: as the reader can verify, it also suffers from the negation problem, and it also yields a defective logic for knowledge and belief.

To these problems we can add a more basic difficulty: the representation of the common ground the present system delivers is too rich. The common ground

is supposed to reflect the information that participants in conversation mutually presuppose. But the present system *requires* each interlocutor to assign all the propositions which reflect live possibilities in the conversation a probability, *regardless* of whether or not that probability has come up in conversation. This makes it impossible to ‘read off’, given an agent’s state of presupposition, which probability values are those which the agent has explicitly coordinated on, vitiating to a significant degree the very idea of the common ground. It also means, given the semantics in place, that for every open proposition p , one is *required* to presuppose either that p is likely, or that p is not likely. But it is just false that one is required to presuppose one of these two things for every open proposition—especially when nothing has been said on the matter.

5 Blunt context probabilism

The discontents of sharp context probabilism can be addressed by retreating to a weaker, but still probabilistic, representation of the common ground, and by making a supervaluationist move within the semantics. We drop the thesis of sharp context probabilism in favor of:

Blunt context probabilism. A conversational common ground can be modeled as a set of probability spaces.

We introduce the notion of a blunt information state:

Definition 11. A BLUNT INFORMATION STATE is a set of sharp information states.

We assume that a state of presupposition, hence a common ground, is representable as a blunt information state; and we assume the same for states of belief and of knowledge. Correspondingly, we reinterpret \mathcal{B}^w and \mathcal{K}^w in our model so that they correspond to blunt information states (given a world).

The idea of representing a state of belief as a set of probability spaces (or measures) has been investigated extensively in the formal epistemology literature. The idea of representing a state of knowledge as a set of probability spaces has been less discussed, but it is perfectly intelligible; it is moreover recommended by the considerations reviewed in [Moss 2011].

Turning back to semantics, we need not start over, redefining every clause now for blunt information states. Rather, we can simply take our probabilistic update function as defined above, and extend it to blunt states by adding the following clause: for all blunt information states I :³

$$I[\phi] = \{i \in I : i[\phi] = i\}$$

To update a blunt information state I with a sentence ϕ , eliminate all those $i \in I$ which don’t accept ϕ . We extend the notion of acceptance to blunt information

³ Compare the structurally analogous approach of [Willer 2010].

states, and we redefine consequence in terms of preservation of acceptance with respect to blunt states:

Definition 12. Blunt state I ACCEPTS ϕ iff $I[\phi] = I$.

Definition 13. $\phi_1, \dots, \phi_n \models \psi$ iff $\forall I$: if I accepts each of ϕ_1, \dots, ϕ_n , I accepts ψ .

6 Problems solved

Blunt context probabilism, together with the dynamic semantics given above, solves all our problems.

First, the negation problem is solved: (1) $\not\models$ (2). Likewise, (4) $\not\models$ (5):

4. It is not the case that John believes that it is probably raining. ($\neg B\Delta\phi$)
5. John believes that it is not likely to be raining. ($B\neg\Delta\phi$)

For John to fail to believe it is probably raining, his doxastic state need only leave open some sharp information state which associates the rain outcome with some probability less than or equal to .5. By contrast, for John to believe it is not likely to be raining, *all* of the sharp information states compatible with his doxastic state must be such as to associate the rain outcome with some probability less than or equal to .5.

Second, we are no longer driven to absurdities concerning the relation between belief and knowledge. Given the truth of some factual ψ , it does *not* follow that $K\Diamond\psi$; hence we cannot use it to infer, by the schema $K\phi \models B\phi$, that $B\Diamond\psi$; hence it doesn't follow, merely from the assumption of ψ , that $\neg B\neg\psi$. In short, false belief is once again compatible with $K\phi \models B\phi$. In general, we can allow for a picture which comports with the one we find in familiar possible worlds models of belief and knowledge. In that picture, one's belief worlds form a subset of one's knowledge worlds; one believes everything one knows, and (typically) more besides. Likewise, we can say that one's blunt doxastic state forms a subset of one's blunt epistemic state.

Third, blunt context probabilism delivers a much more plausible representation of the common ground than its sharp cousin. One need not presuppose probabilities for propositions that one has no reason to believe one's interlocutor is presupposing. And it is not the case that for every open proposition p , one is required to either presuppose that p is likely, or that p is not likely. One might leave the whole interval of probability values for p open.

There is a further advantage of the system worth mentioning. The textbook semantics for \Diamond and \rightarrow faces a problem explaining how exactly epistemic modal and conditional utterances can be informative, despite the fact that they do not serve to add information to the common ground in virtue of their context change potentials. Our sharp probabilistic dynamic semantics has the same problem—indeed, if anything, it makes the problem worse, as it adds sentences of the form $\Delta\phi$ to the list of problematic sentences. By contrast, the semantics based around blunt information states has no such explanatory burden. Sentences of

the form $\diamond\phi$, $\phi \rightarrow \psi$, and $\Delta\phi$ all have the potential to change, without destroying, the common ground. Characteristically their role is to eliminate sharp information states from the common ground, just like ordinary factual claims. These sentences differ from ordinary factual claims insofar as they may eliminate information states as a function of global properties of their domains, or as a function of their probability measures *per se*.

7 Probability conditions

Our system in effect replaces the traditional notion of a truth-condition (a condition on possible worlds) with the more general notion of a probability condition (a condition on probability spaces). Anything you can do with truth-conditions, you can do with probability conditions; but as the system illustrates, with probability conditions one can do more besides. In addition to eliminating possibilities in conversation, we can shift the admissible probabilities over the possibilities, perhaps without eliminating any possibilities.

The conception of information this picture recommends is not as radical as it may appear. On the contrary, it is not far from the conception already found in standard information theory. For in information theory, the amount of information a signal carries is not just a function of the possibilities it eliminates. It is also a function of how it shifts the probabilities over the open possibilities. Information is a fundamentally probabilistic notion.

8 Dynamicness

A semantics which associates sentences with probability conditions might in principle be given statically or dynamically. We have adopted a dynamic formulation, but it can be asked: is such a formulation strictly required to achieve exactly the update effects the system delivers?

We are in effect asking whether the update system supplied by our blunt dynamic semantics is “fundamentally dynamic.” We should clarify this question. The general notion of an update system may be defined as follows:

Definition 14. An UPDATE SYSTEM is a triple $\langle L, C, [\cdot] \rangle$, where L is a set of sentences, C is a set of contexts, and $[\cdot]$ is a function from L to a set of unary operations on C .

We can then define the notion of a *static* update system by reference to a highly general notion of a *static semantics*, as follows:

Definition 15. A STATIC SEMANTICS is a triple $\langle L, W, \llbracket \cdot \rrbracket \rangle$, where L is a set of sentences, W is a set of points, and $\llbracket \cdot \rrbracket$ is an interpretation function, with $\llbracket \cdot \rrbracket : L \rightarrow \mathcal{P}(W)$.

Definition 16. An update system $\langle L, C, [\cdot] \rangle$ is STATIC if and only if there exists a static semantics $\langle L, W, \llbracket \cdot \rrbracket \rangle$ and a one-to-one function f from C to $\mathcal{P}(W)$ such that for all $c \in C$ and $s \in L$: $f(c) \cap \llbracket s \rrbracket = f(c[S])$.

It is fair to say this reflects a fairly standard sense of ‘static’.

What makes for staticness so conceived? [Rothschild and Yalcin 2012] answer this question by proving the following theorem:

Theorem (static representability). An update system $\langle L, C, [\cdot] \rangle$ is static iff for all $s \in L$ and $c \in C$: (i) $c[s][s] = c[s]$ (*idempotence*); (ii) $c[s][s'] = c[s'][s]$ (*commutativity*).

Given this result, it is easy to see that the blunt dynamic system is static, because it is easy to see that it is idempotent and commutative. To make it plain, we can rewrite the extension of the semantics to blunt information states as:

$$6. I[\phi] = I \cap \{i \in \mathbb{I} : i[\phi] = i\}$$

(Where \mathbb{I} is the space of all sharp information states for the given model.) So in a relatively deep sense, the blunt dynamic system is only superficially dynamic.

My central interest here has not been to adjudicate between dynamic versus static implementations of the above ideas. Perhaps ultimately a static formulation (see, for instance, [Yalcin 2007], [Yalcin 2012a]) is to be preferred. Or perhaps the (superficially) dynamic statement of the system is to be preferred, not because an overtly static formulation is impossible, but simply because the dynamic statement is relatively elegant and perspicuous, or has other explanatory virtues. The matter deserves further investigation.

I wish to close by noting one way of reintroducing bonafide dynamicness. Consider the idea of rewriting the extension of the semantics to blunt states as follows:⁴

$$7. I[\phi] = \{i : \text{for some } i' \in I : i'[\phi] = i \text{ and } i \neq \langle \emptyset, Pr_i(x|\emptyset) \rangle\}$$

Here we update a blunt state I , not by tossing out those $i \in I$ which fail to accept ϕ , but rather by updating each i with ϕ (and by setting aside any updates that result in the null context).

The proposal (7.) allows for commutativity failures, and (hence) gives rise to a non-static update system. For instance, it would not be true in general that $I[\neg\phi][\diamond\phi] = I[\diamond\phi][\neg\phi]$. The update corresponding to the former order will always yield the empty set; not so the latter order. Likewise, it is not generally the case that $I[\neg\phi][\Delta\phi] = I[\Delta\phi][\neg\phi]$.

It is not at all obvious that this failure of commutativity is an empirical virtue. On the contrary, (6.) makes it easier to understand these data:

8. It is probably raining. # It is not raining.

9. # If it is probably raining, then if it is not raining, we don't need an umbrella.

(See [Yalcin 2012b] for some additional discussion.) In support of (7.), however, consider a conversation accepting $\diamond\phi$ and $\diamond\neg\phi$. Doesn't this context make explicit provision for the possibility of eventually incorporating the information that ϕ ? But such a context would be crashed by ϕ , according to (6.).

A conservative compromise would be to take (6.) as the primary rule, letting (7.) reflect the appropriate recourse where (6.) would yield the empty set.

⁴ Compare again [Willer 2010].

References

- [Beaver 2001] David Beaver. *Presupposition and Assertion in Dynamic Semantics*. Studies in Logic, Language and Information. CSLI Publications, 2001.
- [Douven and Verbrugge 2010] Igor Douven and Sara Verbrugge. The Adams family. *Cognition*, 117(3):302–318, 2010.
- [Gillies 2004] Anthony Gillies. Epistemic conditionals and conditional epistemics. *Nous*, 38:585–616, 2004.
- [Heim 1982] Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, 1982.
- [Heim 1983] Irene Heim. On the projection problem for presuppositions. In D. Flickinger et. al. (eds) *Proceedings of the Second West Coast Conference on Formal Linguistics*, 114-25. Stanford University Press.
- [Heim 1992] Irene Heim. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3): 183-221.
- [Hulstijn 1997] Joris Hulstijn. Structured information states. In A. Benz and G. Jäger, editors, *Proceedings of MunDial97*. University of Munich.
- [Lassiter 2011] Daniel Lassiter. *Measurement and Modality: The Scalar Basis of Modal Semantics*. PhD thesis, PhD dissertation, New York University, 2011.
- [Lewis 1976] David K. Lewis. Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85:297–315, 1976.
- [Moss 2011] Sarah Moss. Epistemology formalized. Unpublished manuscript, University of Michigan.
- [Rothschild and Yalcin 2012] Daniel Rothschild and Seth Yalcin. Dynamics. Unpublished manuscript, Oxford University and the University of California, Berkeley.
- [Stalnaker 1974] Robert Stalnaker. Pragmatic presuppositions. In M. K. Munitz and P. Unger, editors, *Semantics and Philosophy*. NYU Press, 1974.
- [Stalnaker 1978] Robert Stalnaker. Assertion. In P. Cole, editor, *Syntax and Semantics 9: Pragmatics*. Academic Press, 1978.
- [Stalnaker 1984] Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, 1984.
- [Stalnaker 2002] Robert Stalnaker. Common ground. *Linguistics and Philosophy*, 25: 701–21, 2002.
- [Veltman 1996] Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–61, 1996.
- [Willer 2010] Malte Willer. *Modality in Flux*. PhD thesis, University of Texas, Austin, 2010.
- [Yalcin 2005] Seth Yalcin. Epistemic modals. In J. Gajewski, V. Haquard, B. Nickel, and S. Yalcin, editors, *New Work on Modality*, pages 231-72. MIT Working Papers in Linguistics vol. 51, 2005.
- [Yalcin 2007] Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, Nov. 2007.
- [Yalcin 2010] Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–37, 2010.
- [Yalcin 2011] Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 295–332. Oxford University Press, 2011.
- [Yalcin 2012a] Seth Yalcin. A counterexample to Modus Tollens. *Journal of Philosophical Logic*, forthcoming 2012.
- [Yalcin 2012b] Seth Yalcin. Dynamic semantics. in G. Russell and D. Fara, editors, *Routledge Encyclopedia of the Philosophy of Language*, forthcoming 2012.