Context, Consistency, and Non-Contradiction[*]

Patrick Skeels

University of California, Davis

Synopsis: Dynamic semantics violates numerous classical laws, including Non-Contradiction. Proponents of dynamic semantics have offered no explanation for this behavior, and some critics consider this result to be strong evidence against the tenability of the dynamic program. I defend and explain failures of Non-Contradiction by comparing dynamic semantics and classical, truth-conditional semantics in terms of their idealizing assumptions. I demonstrate that dynamic semantics rejects context fixity, an idealizing assumption that truth-conditional semantics typically adopts. I then argue that any semantics which rejects context fixity should, by the classical semanticist's own lights, violate Non-Contradiction under certain circumstances. I then demonstrate that dynamic semantics violates Non-Contradiction in all and only those circumstances. I subsequently appeal to this insight to vindicate some of dynamic semantics' more controversial predictions. I close by suggesting that discussion of idealizing assumptions, common in the sciences, is similarly crucial to fruitful discussion in natural language semantics.

Introduction

Dynamic semantics for epistemic modals is non-classical in a variety of respects. In addition to violating all of the structural rules commonly associated with classical entailment, a laundry list of classically valid inference patterns are invalidated. Consider an instance of each below.

(1) **Commutativity**: $\varphi \wedge \psi \Rightarrow \psi \wedge \varphi$.

(2) **Monotonicity**: If $\Gamma \Rightarrow \varphi$ then $\Gamma \cup \Gamma' \Rightarrow \varphi$.

Commutativity and Monotonicity both hold in classical logic, but fail in dynamic semantics. These failures enjoy somewhat privileged status among the classical rules violated, since they are motivated by specific features of natural language. Commutativity fails due to putatively asymmetric behaviors of natural language conjunction. Failures of monotonicity capture the fact that utterances of epistemically modalized sentences are infelicitous in contexts where the negation of the prejacent is established.

Not all deviations from classical logic receive such direct explanations. In fact, most receive none at all.

(3) **Non-Contradiction**: $\varphi \wedge \neg\varphi$ is inconsistent.

Like Commutativity and Monotonicity, Non-Contradiction holds in classical logic but does not hold in dynamic semantics. Unlike Commutativity and Monotonicity, violations of Non-Contradiction, while recognized in the literature, are not explicitly argued for, nor are any motivating cases from natural language proffered. It thus remains unclear whether, and if so, why, we should accept failures of Non-Contradiction in a semantics for natural language. A proper explanation for the rejection of Non-Contradiction thus requires that the dynamic semanticist offer some positive motivation for its rejection. One might be tempted to search for natural language cases that support this rejection. In particular, some sentence of the form $\varphi \wedge \neg\varphi$ that is felicitously utterable would provide justification in similar style to those provided for violations of Commutativity and Monotonicity. Unfortunately, such cases are not likely forthcoming. While dynamic semantics treats certain instances of $\varphi \wedge \neg\varphi$ as consistent, it treats all instances of $\varphi \wedge \neg\varphi$ as incoherent, (Mandelkern, 2020, pp.388). Dynamic semantics recognizes multiple entailment relations as well as their inderdefinable notions of consistency. Coherence is one such notion. Dynamic semanticists regularly appeal to both incoherence and inconsistency to predict instances of infelicity while theorizing in a single semantic framework. Sentences that are consistent, but incoherent, are often taken to be infelicitous, but marginally assertible. Such sentences typically sound a bit strange, but can be assertible in limited contexts, unlike wholesale contradictions. This is to say that properly felicitous utterances of sentences which violate Non-Contradiction seem like a bridge too far, since by the dynamic semanticist's own lights, such sentences will only be marginally assertible at best. Thus, failures of Non-Contradiction are expected to be infelicitous, to some degree, through independent means, and dynamic theorists will be unable to uncontroversially motivate violations of Non-Contradiction in the same way that they have motivated violations of Commutativity and Monotonicity. The dynamic semanticist will be unable to rely on sentence glosses alone to convincingly motivate rejections of Non-Contradiction, and separate justification must be provided.

Even if the dynamic semanticist is able to provide alternative motivatation for the rejection of Non-Contradiction, they face a separate, albeit related, challenge. Mandelkern (2020) argues that instances of $\varphi \wedge \neg\varphi$ that dynamic semantics treats as consistent are intuitively contradictory. Moreover, utterances of their natural language correlates strike many as infelicitous. However, the dynamic strategy mentioned above treats sentences that are consistent, yet incoherent, as marginally assertible in certain contexts, which would mean that violations of Non-Contradiction should also be marginally assertible. This prediction is controversial. Thus, even if positive theoretical motivation for the rejection of Non-Contradiction is provided, dynamic semanticists must still vindicate the system's more controversial predictions concerning natural language.

Failures of Non-Contradiction may be surprising and unintuitive. Moreover,

the lack of explicit motivation coupled with controversial and, heretofore, undefended predictions concerning natural language can be considered significant marks against the program. Matters are worsened by the fact that violations of Non-Contradiction are not alone, and that much of the non-classical behavior of dynamic semantics faces similar explanatory problems.[1] As such, the proponent of dynamic semantics owes a substantial explanatory debt, and it is far from obvious what adequate payment should look like. The present paper sketches the contours of a satisfying explanation and aims to make a substantial down payment. To do this, I argue that failures of Non-Contradiction can be explained and justified once dynamic semantics is characterized in terms of its *idealizing assumptions*.

I begin by articulating an idealizing assumption called *context fixity*, which goes back, at least, to Kaplan and is operative in canonical accounts of context sensitivity. Kaplan observed that spatial and temporal shifts in context that occur mid-sentence would result in immediate failures of numerous classical laws. To avoid this, Kaplan famously argued that semantics concerns occurrences, rather than utterances, and proposed that we evaluate sentences and entire arguments at single contexts, that we might maintain classically valid inference patterns. I then demonstrate that dynamic semantics rejects context fixity by allowing updates, rather than temporal or spatial shifts, to alter context. This allows for dynamic semantics to reject context fixity without resorting to theorizing in terms of utterances. I then consider some results concerning precisely when dynamic semantics predicts failures of Non-Contradiction and demonstrate that these failures always and only occur as a result of the kind of context shift Kaplan aimed to avoid. Thus, dynamic semantics makes predictions that are perfectly in line with the intuitions that motivate traditional contextualist approaches to truth-conditional semantics in non-ideal environments. I conclude that dynamic semantics is justified in rejecting Non-Contradiction despite the fact that this rejection is not motivated by any felicitously utterable sentences. I speculate that similar explanations can plausibly be ported to other features of dynamic semantics that do not enjoy direct linguistic support. Further, I explain how conversational contexts which motivate the rejection of context fixity vindicate some of dynamic semantics' more puzzling predictions while still respecting our intuitions. Finally, I suggest that distinguishing semantic programs in terms of their idealizing assumptions allows us to understand putatively competing theories in considerably less antagonistic terms. This, in turn, allows us to more properly adjudicate differing predictions made by such frameworks.

---

[1]For instance, Excluded Middle is also violated, (Mandelkern, 2020, pp.386). As with violations of Non-Contradiction, intuitions on glosses that support rejection of the principle are incredibly difficult to come by. For simplicity, I restrict my attention to Non-Contradiction, and omit disjunction entirely. However, I expect many of the forthcoming considerations to apply to Excluded Middle as well.

1 Dynamic Semantics

1.1 Overview

The received view in natural language semantics takes the meaning of a sentence to be its truth-conditions, and semantic values are characterized in terms of propositions. Truth-conditional theories are mum with respect to the manner in which the meaning of a sentence uttered in context affects a state of information or conversation, as canonically, this is a job for pragmatics.[2] While the truth-conditional theorist has a litany of options for addressing this issue, adoption of truth-conditional semantics does not come with any particular commitments about how this process is carried out. Dynamic semantics for epistemic modals distinguishes itself from truth-conditional semantics by characterizing the compositional semantic values of sentences in terms of their discourse effects, rather than their truth-conditions. Instead, the meaning of a sentence is its capacity to change a state of information: a *context change potential* (CCP). Accordingly, the semantic values of the connectives are similarly defined in terms of CCPs. Such semantics can be called *compositionally dynamic*, which is to say that the discourse effects of sentences are determined by the compositional structure of the sentence itself, (Rothschild and Yalcin, 2016, pp.334). As each component of a sentence is processed, it updates the state of information to reflect this change. Since discourse effects are calculated compositionally, if a single sentence's compositional value is constituted by multiple sequential updates, the conversational context upon which later updates apply is not guaranteed to be the same as the initial context; context changes mid-sentence. This results in what can be called a *local context*, an intermediate state of information that reflects some, but not all, of the updates contained in the meaning of a given sentence, (Dever, 2013, pp.113). While a dynamic semantics is not required in order to appeal to local contexts, dynamic semantics is unique insofar as local contexts determine the compositional semantic values of sentences.[3]

1.2 Formal Details

I now consider the propositional fragment of the dynamic language proposed in Groenendijk and Stokhof (1991b).[4]

Let $\mathcal{L}_1$ be generated by the grammar:

$$\varphi ::= \mathrm{p} \mid \neg\varphi \mid (\varphi \land \varphi) \mid \Diamond\varphi$$

Dynamic semantics is intended to capture the manner in which sentences can update a state of information based on their compositional structure. We thus

---

[2]See Stalnaker (1978) for the canonical account.

[3]See Schlenker (2009) for a pragmatic account of local contexts.

[4]See Groenendijk and Stokhof (1991a), Veltman (1996), van der Does et al. (1997), Gillies (2004), von Fintel and Gillies (2007), Yalcin (2012b), Yalcin (2015), and Willer (2013) among many others, for developments and variations of the test semantics.

characterize a general framework for the way in which sentences can update states of information. This can be captured by an *update system*.

(4) **Update System**: $\langle \mathcal{L}, C, \cdot[\cdot] \rangle$ is an update system if and only if $\mathcal{L}$ is a set of sentences, C is a set of information states, and $\cdot[\cdot]$ is a function which maps sentences of $\mathcal{L}$ to operations on C.

An update system models the manner in which sentences of a given language can change a state of information. Let W be the set of all functions $\mathcal{A} \mapsto \{0, 1\}$. Our update system will be $\langle \mathcal{L}_1, \mathcal{P}(W), \cdot[\cdot] \rangle$. Some information state $s \in \mathcal{P}(W)$ is a set of possible worlds compatible with the information contained within the state. We characterize our semantics for $\mathcal{L}_1$ in terms of operations on $\mathcal{P}(W)$.[5]

(5) **Semantic Clauses**[6]

    1. $s[\mathrm{p}] = \{w \in s | w(\mathrm{p}) = 1\}$

    2. $s[\varphi \wedge \psi] = s[\varphi][\psi]$

    3. $s[\neg \varphi] = s - s[\varphi]$

    4. $s[\Diamond \varphi] = \{w \in s : s[\varphi] \neq \varnothing\}$

Updating a state with an atomic formula takes the intersection of the worlds in the state, and the set of worlds where that atomic is true. Conjunction is consecutive update, and negation takes the difference between the initial state, and what would result by updating the initial state with the non-negated formula.[7] Epistemic possibility (epistemic "might") performs a test, where the initial state is returned if the state is compatible with the prejacent of the modal. Otherwise, the test is 'failed' and the absurd state $\varnothing$ is returned.

    I now define what it is for a state to support some information.

(6) **Support**: $s \vDash \varphi \Leftrightarrow s[\varphi] = s$

A state supports some information if and only if updating the state with that information does not change the state. Next, I define consequence and consistency.

(7) **Dynamic Consequence**: $\psi_1, ..., \psi_n \Vdash \varphi \Leftrightarrow \forall s, s[\psi_1]...[\psi_n] \vDash \varphi$

(8) **Dynamic Consistency**: $\psi_1, ..., \psi_n$ is consistent $\Leftrightarrow \exists s, s[\psi_1]...[\psi_n] \neq \varnothing$

---

[5] According to Rothschild and Yalcin, a conversation system coupled with a semantics for the language determine what is called a *state system*. Whether, and in what capacity, a semantics is dynamic is determined by features of the state system. On this characterization, our present semantics is dynamic precisely because it fails to be idempotent and commutative. More on this presently.

[6] The present formulation only considers negation, conjunction, and epistemic possibility operators. Disjunction and quantification are suppressed, as the problem can be formulated in their absence.

[7] The failure of Non-Contradiction can be avoided by adopting the 'static' version of conjunction, which is set intersection. An example of this strategy can be found in the early pages of Veltman (1996). The problem, as Mandelkern (2020) observes, is that the static approach also prevents dynamic semantics from predicting many of the phenomena it is intended to predict, as the set theoretic operations are commutative.

(9) **Dynamic Inconsistency**: $\psi_1, ..., \psi_n$ is inconsistent $\Leftrightarrow \psi_1, ..., \psi_n$ is not consistent

Dynamic consequence is often called *update-to-test consequence*. A sequence of formula dynamically entails some formula φ if and only if every state updated with that sequence will support φ. We can define dynamic consistency in terms of update-to-test consequence. A sequence of updates is dynamically consistent if and only if there is a state that can be updated with that sequence without resulting in the absurd state. Note that this notion of entailment does not concern truth, and instead captures support preservation. Also note that the above notions are order sensitive. For instance, $\langle \Diamond p, \neg p \rangle$ is dynamically consistent while $\langle \neg p, \Diamond p \rangle$ is not.

Dynamic semantics is not limited to update-to-test consequence and also avails itself of other notions of entailment. Of interest here is an alternative often called *test-to-test consequence*. Its interdefinable notion of consistency is often called *coherence*.

(10) **Coherence**: $\psi_1, ..., \psi_n$ is coherent $\Leftrightarrow \exists s \neq \varnothing$, s.t. $s \vDash \psi_1, ..., s \vDash \psi_n$

(11) **Incoherence**: $\psi_1, ..., \psi_n$ is incoherent $\Leftrightarrow \psi_1, ..., \psi_n$ is not coherent.

Unlike dynamic consistency, coherence demands that a single state support each update individually, rather than in sequence. Coherence is not order sensitive, e.g. neither $\langle \Diamond p, \neg p \rangle$ nor $\langle \neg p, \Diamond p \rangle$ are coherent. Any update that is incoherent is not supported by any state other than $\varnothing$.

A few more definitions will be useful shortly.

(12) **Idempotence**: $[\varphi]$ is idempotent $\Leftrightarrow \forall s, s[\varphi] = s[\varphi][\varphi]$

An update is idempotent if and only if subsequent applications of the same update do not yield any change beyond the first. Importantly, a certain class of updates in this semantics are not idempotent.[8]

(13) **Distributivity**: $[\varphi]$ is distributive $\Leftrightarrow \forall s, s[\varphi] = \bigcup \{\{w\}[\varphi] | w \in s\}$

An update is distributive just in case updating a state with the formula is equivalent to the union of the individual updates of the singleton set of each world in that state, (van Benthem, 1989, pp.364). All contents that do not include the $\Diamond$ operator are distributive. Distributive updates can always be characterized by the union of an information state and a fixed set of worlds. This set of worlds will not change depending on the nature of the state it updates, and thus, updates that are distributive can be said to express propositions. Non-distributive updates do not express propositions, though they may entail them.

---

[8]Precisely speaking, updates, rather than formula, have the property of idempotence. For convenience, however, I will sometimes say that "φ is idempotent" to mean that "$[\varphi]$ is idempotent."

1.3 Selected Applications

The above semantics has been used to explain various behaviors of natural language, often involving epistemic modals and conjunction, that putatively fail to conform to patterns of classical logic. Here, I rehearse a select few that will be relevant for discussion.[9] Veltman (1996) observes a discrepancy in the following two discourses:

(14) ?? Adam might be at the door ...Adam isn't at the door.

(15) # Adam isn't at the door... Adam might be at the door.

According to Veltman, the sentences in (14) are assertible, *in that order*, at certain contexts, namely, those where one expects a particular visitor but is met with another. Flipping the order, as in (15), yields catastrophic infelicity. According to dynamic semantics, the discourse in (14) has the structure $\langle \Diamond p, \neg p \rangle$ while (15) is $\langle \neg p, \Diamond p \rangle$. This discrepancy can be explained by the fact that $\langle \Diamond p, \neg p \rangle$ is consistent, but $\langle \neg p, \Diamond p \rangle$ is inconsistent. Should we further hold, as dynamic semanticists do, that conjunction is consecutive update, then we can get the same result for predicting the related discrepancy in felicity between so-called *epistemic contradictions*.[10]

(16) ?? Adam might be at the door and he isn't.

(17) # Adam isn't at the door and he might be.

(18) $\Diamond p \wedge \neg p$

(19) $\neg p \wedge \Diamond p$

Since conjunction is merely consecutive update, (18) is consistent, while (19) is not. This captures the discrepancy in infelicity between(16) and (17).[11] Of course, the only difference between (16) and (17) is the order of the conjuncts. Dynamic semantics can explain this discrepancy due, in part, to the failure of commutativity of conjunction. Cases such as these also motivate failures of monotonicity. For instance, an agent undecided about whether p can assert or believe $\Diamond p$. However, should an agent learn that p is false and come to believe $\neg p$, they can no longer maintain $\Diamond p$. This suggests, contra the classical rule of monotonicity, that updating one's state does not guarantee that all of the

---

[9]Dynamic semantics has been commonly used to solve modal disagreements, as well as predicting the infelicity of epistemic contradictions. See Willer (2013) for a dynamic solution to both. See MacFarlane (2011) and MacFarlane (2014) for discussions of modal disagreement and see Yalcin (2007) for discussion of epistemic contradictions.

[10]See Gillies (2004) for an argument concerning why consecutive assertion can be treated as dynamic conjunction.

[11]This discrepancy in infelicity is captured in data gathered Knobe and Yalcin (2014). Sentences of the form $\Diamond p \wedge \neg p$ were considered infelicitous, but substantially less so than sentences of the form $\neg p \wedge \Diamond p$. This dovetails with Veltman's intuitions that one ordering is assertible, while the other is not. In addition, dynamic semantics can predict the infelicity of (16) in virtue of the fact that it is incoherent, despite the fact that it is consistent. This is to say that a state can sustain update with (18) without crashing, but no single state can support each update in (18). This explains why (16) sounds bad, but not as bad as (17).

formula supported by the original state will survive the update. This also helps explain the discrepancy in infelicity observed between (16) and (17). Importantly, the fact that (16) is only marginally assertible and does not seem to be felicitously embeddable can be explained by the fact that it is incoherent.

While (16) may be marginally assertible, observe that it can never be asserted twice in succession.

(20) # Adam might be at the door, and he isn't at the door, and he might be, and he isn't.

(20) is of the form $(\Diamond p \wedge \neg p) \wedge (\Diamond p \wedge \neg p)$, so according to dynamic semantics, its meaning is captured by the following sequence of updates: $[\Diamond p][\neg p][\Diamond p][\neg p]$. No state, however, can sustain sequential update with the middle two updates, $[\neg p][\Diamond p]$, without resulting in the absurd state. Thus, (20) is inconsistent and thereby infelicitous. In addition to the infelicity of (20), dynamic semantics also explains why the first instance of (16) can have a distinct discourse effect from the second. $[\Diamond p \wedge \neg p]$ is not idempotent, meaning that subsequent updates can have effects over and above the first. Thus, cases like (20) vindicate the presence of non-idempotent update in dynamic semantics. So-called static frameworks are unable to predict the infelicity of such cases semantically and will have to appeal to pragmatics.

Dynamic semantics also predicts related behaviors including modal disagreements and epistemic contradictions.[12] Further, the dynamic clause for conjunction allows for the semantics to be implemented in other frameworks designed to capture other dynamic phenomena, including presupposition projection and anaphora resolution. The important point, however, is that the motivations for the rejection of Commutativity and Monotonicity are clear, explicit, and motivated by concrete examples involving belief or assertion. The literature, however, does not contain any motivating examples for violations of Non-Contradiction.

2   Violations of Non-Contradiction

Dynamic semantics violates Non-Contradiction. This means that there is some state, s and some formula φ such that $s[\varphi \wedge \neg\varphi] \neq \varnothing$.[13] As demonstrated in Mandelkern (2020), Non-Contradiction fails when $[\varphi]$ is not idempotent. Consider two examples of non-idempotence:

(21) $\Diamond p \wedge \neg p$[14]

---

[12]See Willer (2013), Willer (2015), Lennertz (2019), and Skeels (2023) for examples and discussion.

[13]However, no single state can support both φ and $\neg\varphi$, thus, Non-Contradiction holds for test-to-test consequence.

[14]See Mandelkern (2020) for proof that $[\Diamond p \wedge \neg p]$ is not idempotent.

(22) $\Diamond p \wedge q$[15]

If we use either as a substitution instance for $\varphi \wedge \neg \varphi$, we get:

(23) $(\Diamond p \wedge \neg p) \wedge \neg(\Diamond p \wedge \neg p)$

(24) $(\Diamond p \wedge q) \wedge \neg(\Diamond p \wedge q)$

Despite being of the form $\varphi \wedge \neg \varphi$, (23) and (24) are consistent. For each, there exists a state that can sustain update with the formula without resulting in the absurd state. Proof of the consistency of (23) can be found in Mandelkern (2020). The proof for the consistency of (24) is as follows.

Fact: $(\Diamond p \wedge q) \wedge \neg(\Diamond p \wedge q)$ is dynamically consistent.

Proof: Let w be world s.t. $w(p) = 1$ and $w(q) = 0$ and let w' be a world s.t. $w'(p) = 0$ and $w'(q) = 1$. Let s be a state s.t. $w \in s$ and $w' \in s$.

$$
\begin{aligned}
s[(\Diamond p \wedge q) \wedge \neg(\Diamond p \wedge q)] &= s[(\Diamond p \wedge q)][\neg(\Diamond p \wedge q)] \\
&= s[\Diamond p][q][\neg(\Diamond p \wedge q)] \\
&= s[q][\neg(\Diamond p \wedge q)] \\
&= \{w'\}[\neg(\Diamond p \wedge q)] \\
&= \{w'\} - \{w'\}[(\Diamond p \wedge q)] \\
&= \{w'\} - \{w'\}[\Diamond p][q] \\
&= \{w'\} - \varnothing[q] \\
&= \{w'\} - \varnothing \\
&= \{w'\}
\end{aligned}
$$

$\{w'\} \neq \varnothing$ so $(\Diamond p \wedge q) \wedge \neg(\Diamond p \wedge q)$ is dynamically consistent.

■

Therefore, it is not the case that all substitution instances of $\varphi \wedge \neg \varphi$ are contradictions and Non-Contradiction does not hold.

Importantly, not all instances of $\varphi \wedge \neg \varphi$ are consistent, and it fails in a very particular set of circumstances. While Mandelkern (2020) proves that $[\varphi]$ is not idempotent only if $\varphi \wedge \neg \varphi$ is consistent, the converse also holds.

Fact: If $\varphi \wedge \neg \varphi$ is consistent, then $[\varphi]$ is not idempotent.

Proof: Suppose, for conditional proof, that $\varphi \wedge \neg \varphi$ is consistent. It follows from the definition of consistency that there exists a state s' s.t. $s'[\varphi \wedge \neg \varphi] \neq \varnothing$. It

---

[15]To show that $[\Diamond p \wedge q]$ is not idempotent, let s be an information state s.t. $s = \{w, w'\}$, $w(p) = w'(q) = 1$, and $w'(p) = w(q) = 0$. From the compositional values of the connectives, we can see that $s[\Diamond p \wedge q] = \{w'\}$, while $s[\Diamond p \wedge q][\Diamond p \wedge q] = \varnothing$. Therefore, $s[\Diamond p \wedge q] \neq s[\Diamond p \wedge q][\Diamond p \wedge q]$ and $[\Diamond p \wedge q]$ is not idempotent. Thanks to an anonymous referee for pointing out this instance of non-idempotence.

follows from the clause for conjunction that $s'[\varphi][\neg\varphi] \neq \varnothing$. Now suppose, for reductio, that $[\varphi]$ is idempotent. It follows from the definition of idempotence that $\varphi \Vdash \varphi$. By the definition of entailment, it follows that for all $s, s[\varphi] \vDash \varphi$. Therefore, $s'[\varphi] \vDash \varphi$. However, for any $s \vDash \varphi$, $s[\neg\varphi] = \varnothing$. Thus, $s'[\varphi][\neg\varphi] = \varnothing$, but this contradicts $s'[\varphi][\neg\varphi] \neq \varnothing$ established above. Therefore, $[\varphi]$ is not idempotent.

∎

Thus, $\varphi \wedge \neg\varphi$ is consistent if and only if $[\varphi]$ is not idempotent. This is to say that failures of Non-Contradiction not only include substitution instances where $[\varphi]$ is not idempotent, but are exhausted by them. This means that for idempotent updates, Non-Contradiction holds exactly as it does in the classical case. This includes all distributive (propositional) formulae, meaning that for any $[\varphi]$ that is distributive, $\varphi \wedge \neg\varphi$ is inconsistent. Thus, for any substitution instances that are truth-evaluable, Non-Contradiction holds. $p \wedge \neg p$, for example, is inconsistent. This explains why many instances of $\varphi \wedge \neg\varphi$ are straightforwardly infelicitous.

(25)   # San Francisco is in California and San Francisco is not in California.

Thus, dynamic semantics does not lose its capacity to predict that direct contradictions of truth-evaluable contents are infelicitous. In addition, this also ensures that dynamic semantics does not recognize anything resembling a true contradiction. Dialethism is avoided, since dynamic semantics admits of truth-value gaps, but not gluts. Formulae whose associated updates are not distributive do not express propositions and are not truth-evaluable. While there are instances of $\varphi \wedge \neg\varphi$ that are consistent, no instance that is consistent will express a proposition. In this case, recognizing the failure of Non-Contradiction does not commit one to true contradictions, but the substantially more conservative claim that certain states of information can be updated by certain formulae of the form $\varphi \wedge \neg\varphi$ without crashing.

Formulae of the form $\varphi \wedge \neg\varphi$ that are dynamically consistent must contain instances of φ s.t. $[\varphi]$ is not idempotent. Since (21) and (22) are not idempotent, (23) and (24) are consistent. Their natural language correlates are as follows.

(26)   ?? (Adam might be in his office, and Adam is not in his office) and it is not the case that (Adam might be in his office, and Adam is not in his office.)

(27)   ?? (It might be raining and it's windy) and it's not the case that (it might be raining and it's windy.)

The presence of glosses like (26) and (27) yields a two-part challenge to dynamic semantics. The first is dialectic. As we saw in §1.3, failures of Commutativity and Monotonicity were motivated by intuitions concerning reasonably straightforward utterances and inferences. To defend failures of Non-Contradiction in similar fashion would require uncontroversial cases where (26) or (27) were assertible. Opponents of dynamic semantics may be unwilling to grant that (26)

and (27) are in any way assertible. Given their palpable infelicity, any strategy that takes the assertibility, marginal or otherwise, of (26) and (27) as a motivating premise for the rejection of Non-Contradiction seems doomed to fail, since the motivating premise will be flatly rejected, as in Mandelkern (2020). This leaves us at a standoff since the standard dynamic strategy of motivating non-classical semantics with examples of utterances and discourses in context will likely be controversial and unconvincing, even if it is correct. Thus, positive motivation for the rejection of Non-Contradiction that may be convincing to those not antecedently sympathetic to the dynamic program will have to come via some other means. The first challenge that the dynamic semanticist faces is to provide convincing motivation for the rejection of Non-Contradiction in the absence of uncontroversial motivating glosses.

A second challenge arises from the apparent infelicity of (26) and (27). While alternative motivation is required for a convincing case against Non-Contradiction, we are not absolved from defending the controversial predictions above. This is to say that while dynamic semanticists may be able to convince their interlocutors that they should, in principle, reject Non-Contradiction, intuitions about particular sentence glosses remain. The standard dynamic strategy has been to treat sentences that are consistent, yet incoherent, as marginally assertible. Thus, (26) and (27) should be marginally marginally as well. This prediction is decidedly controversial, and it is far from obvious that it is correct. The proponent of dynamic semantics is thereby obligated to provide at least some examples of marginally assertible violations of Non-Contradiction.

In brief, (26) and (27) are simply too controversial to convincingly motivate the rejection of Non-Contradiction by themselves. We'll have to look elsewhere for further support. In addition, if standard dynamic practice is preserved, arguments must be made in favor of the marginal assertibility of (26) and (27). A successful defense of dynamic semantics' rejection of Non-Contradiction will need to answer both challenges. I begin with the first.

3   Context and Idealization

It will be instructive to consider the different ways in which truth-conditional and dynamic semantic theories interact with context, as well as the insights that motivate these choices.

3.1   Context, Canonically

The standard approach regarding semantic interactions with context is Kaplanian in spirit, if not in letter. Kaplan, motivated primarily by demonstratives and indexicals, argues that the denotation of a sentence is relativized to a context and an index. For Kaplan, a context is a four-tuple of an agent, time, location, and possible world. An index is an n-tuple with arbitrarily many parameters that model any other information that a sentence might be sensitive to. It is standardly taken, however, that the information in a Kaplanian context

should be enough to fully determine any information for an index, so the index is often suppressed.[16] Thus, it is pedestrian to relativize the denotation of a sentence to a context c which supplies all of the relevant contextually supplied information like so.

(28) $[\![\varphi]\!]_c$

On this approach, sentences are evaluated at single contexts, and any components of φ with contextually supplied values will have those values supplied by c. Kaplan further posits the notion of *truth at a context* and defines validity as preservation thereof. Thus, not merely sentences, but entire arguments are evaluated relative to a single, fixed, context. This strategy can be captured by the following principle.

(29)   **Context Fixity**: All components of a sentence (or argument) are evaluated relative to the same context.

According to Context Fixity, contexts do not shift mid-sentence. Each subformula of a sentence that requires any contextually supplied values will receive these values from the same context. The same goes for different premises in an argument, all of which are evaluated at a single context.[17] It is easy, however, to imagine cases of language use where context does change mid-sentence, or mid-argument. Kaplan recognizes this possibility, and is careful to distinguish between an *occurrence* and an *utterance*. An occurrence of a sentence at a context is simply a sentence context pair. We can evaluate an occurrence of a sentence at a context, and determine whether it is true, regardless of whether the sentence is uttered in that context (or whether the language of the sentence even exists at that context, etc.). Utterances, on the other hand, are speech acts, about which Kaplan says the following:

> Utterances take time, and are produced one at a time; this will not do for an analysis of validity. By the time an agent finished uttering a very, very long true premise, the premise may have gone false. Thus even the most trivial inferences, P therefore P, may appear invalid. Also, there are sentences which express a truth in certain contexts, but not if uttered. For example "I say nothing." Logic and semantics are concerned not with the vagaries of actions, but the verities of meanings, (Kaplan, 1989, pp.584).

Kaplan believes that occurrences, rather than utterances, are the proper subject matter of semantics. Even though no two spoken sentences can be uttered in the

---

[16]Since a context should fully determine an index, Lewis (1980) famously calls Kaplan's distinction between context and index a "distinction without a difference."

[17]The strategy of defining entailment in terms of truth at a context is ubiquitous in post-Kaplan approaches to context sensitivity as well. See Predelli (2005) for an example. Predelli, contra Lewis, maintains Kaplan's distinction between context and index, and thus treats entailment as preservation of truth at a context and an index. This is relevant because, while I articulate the Lewisian strategy for simplicity, nothing hangs on it.

exact same context, occurrences of these sentences can be evaluated at single, fixed contexts, and we can model whether they preserve truth at these contexts. In so doing, we can theorize entirely in terms of occurrences of sentences evaluated at the same context, and thus, validity that concerns preservation of truth at a context is an account of occurrence validity. According to Kaplan, if we do not do this, and theorize in such a way that context can change mid-sentence, we intuit immediate failures of inference patterns like repetition. To avoid dealing with situations like these, Kaplan thinks that we ought to characterize validity at the occurrence level, rather than the utterance level. Kaplan is extremely pessimistic about the prospects of the latter.

> I [Kaplan] am unclear even as to what arguments *ought* to come out as utterance-valid (as opposed to occurrence-valid). There are different notions of utterance-validity corresponding to different assumptions and idealizations. With no idealizations, the rules of repetition and double negation become invalid. This seems hopeless, (Kaplan, 1989, pp.585)

It is important to appreciate that Kaplan is not arguing that we adopt context fixity because it is true of natural language consequence. Indeed, it is almost certainly false, as any account of entailment which satisfies context fixity will fail to capture even the simplest cases of intercontextual entailment, e.g. If "It will rain tomorrow" is true at a context today, then "It will rain today." must be true at a context tomorrow.[18] This is not an objection to Kaplan's claims, but merely the observation that Kaplan's proposal is based on the perceived intractability of the alternative. This is to say that Kaplan is proposing that we make an *idealizing assumption*. A characteristic feature of such assumptions is that they are not made based on their truth, real or believed. Instead, they are made for practical reasons like simplicity, understandability, and tractability, (Potochnik, 2017, pp.43-44). Well known examples of such assumptions include frictionless planes and completely rational agents. Despite the fact that neither exist, each has proven indispensable in our theorizing. The same may plausibly be said of context fixity, as no two sentences are ever uttered at a single context. It has been, nonetheless, extremely profitable and informative to use this idealization in understanding indexicals and demonstratives. It has also provided the possibility of a generalized framework for theorizing about context sensitive meanings. Thus, according to the Kaplanian position, we needn't deny that contexts can change mid-sentence. Rather, given the intractability of theorizing in terms of utterance validity, we can idealize by holding context fixed. We can then cleanly discuss entailments of occurrences.

As far as Kaplan is concerned, Context Fixity is an idealizing assumption that is foisted upon us given the "hopelessness" of theorizing in its absence. Should we reject it, Kaplan believes that we must suffer the following:

---

[18]Such examples are famously discussed in Frege (1918). See Zardini (2014) for a logic intended to capture these entailments. Not surprisingly, the system presented in Zardini (2014) rejects context fixity.

1. We are forced to theorize in terms of utterances, rather than occurrences.

2. We have no clear intuitions about utterance validity.

3. We lose patterns like Repetition.

Fortunately, there is good reason to believe that Kaplan's worries are overstated. 1. is false, and there is room in the logical space for a theory which characterizes validity at the level of occurrences, but rejects context fixity. We might choose to evaluate particular occurrences at a context but allow subsequent sentences to modify the context. We may then characterize a notion of validity which accounts for information preservation through changing contexts. More on this in a moment, but the crucial point is that rejecting context fixity does not require that we theorize in terms of utterances. This also immediately absolves us of any responsibility to address 2. as we needn't discuss utterance validity at all. 3. remains, and patterns like Repetition would be lost, but this is far from hopeless. Rather, this is exactly what we would expect from a notion of entailment which rejects context fixity. Somewhat roughly, if different instances of the same sentence can have different outputs at distinct contexts, and we intend to capture a notion of validity that recognizes information preservation through context change, then we should expect failures of rules like Repetition. Put slightly differently, we would expect this account of entailment to reject the structural rule of *Reflexivity*.

(30) **Reflexivity**: $\varphi \Rightarrow \varphi$

As contexts change, the truth or acceptability of a sentence may also change. A satisfactory logic of context change would need to predict and explain precisely when this would happen.

Thus, rejecting context fixity is by no means hopeless and does not require that we brave the "vagaries of action" by theorizing in terms of utterances. We will, however, see some non-classical behaviors including failures of Reflexivity.[19] Importantly, by reasoning very similar to Kaplan's, we should also expect other non-classical behaviors, including failures of Non-Contradicton. Consider natural language sentences of the following form:

(31) φ_1 and it's not the case that φ_2.

Each instance of the sentence φ is given a metalinguistic marker in (31) above. It is easy to see that if φ_1 and φ_2 are evaluated at the same context, we understand (31) to be inconsistent. If the contents expressed by φ_1 and φ_2 are evaluated at the same context, then should have both the same output at that context. Our intuitive understanding of negation and conjunction also tells us that something

---

[19]I mention that not all rejections of context fixity result in failures of reflexivity. An example would include the propositional fragment of the $\mathcal{L}_1$ discussed above. Rather, we expect failures of reflexivity when we allow for what we might call *non-trivial context change*. Context change is non-trivial just in case the difference in contexts is substantial enough such that the same formula can have distinct meanings at each context. When I say a semantics rejects context fixity, I properly mean that it allows for non-trivial context change.

in conjunction with the negation of that very thing is contradictory, and thereby inconsistent. However, once we reject context fixity, we may evaluate φ_1 at a different context from φ_2, meaning that they are not guaranteed to have the same output. In this case, the very same intuitions about negation and conjunction yield a different conclusion, namely, that if the context has changed enough, (31) is consistent. This reasoning is not importantly different from that which motivates Kaplan's claims concerning the failure of repetition. We may thus conclude that rejecting context fixity opens the door for failures of Non-Contradiction in the same way it did for failures of Reflexivity. Importantly, this conclusion seems to arise from the very same intuitions that motivate the canonical approach. Thus, by Kaplan's lights, as well as our own, we can expect failures of Reflexivity and Non-Contradiction in accounts that reject context fixity.

3.2 Dynamic Contexts

The previous section established that there is room for theorizing which recognizes context change while still focusing exclusively on occurrences. I have also argued that, by Kaplanian lights, we should expect such a system to violate Reflexivity and Non-Contradiction. Of course, this region of the theoretical space is not uncharted wilderness and has been inhabited by dynamic semantics for quite some time. Dynamic semantics processes updates relative to a context (an information state) and outputs a new, updated, context. Subsequent updates, including updates in the same sentence or argument, are updated relative to the output of the previous update, which will often be different than the initial context. Such intermediate contexts are called local contexts, and it is often taken that the semantic significance of local contexts is characteristic of dynamic semantics. Dynamic semantics still trades in occurrences, and an occurrence of a sentence at a context will update that context. However, the context at which a sentence occurs can be changed by update with the sentence itself, and not just by pragmatic factors. Dynamically speaking, context change is not merely an artifact of the temporal extension of speech acts. Within a dynamic framework, the meanings of occurrences themselves can change at least one contextual parameter: the information state. It should thus be clear that dynamic semantics directly rejects context fixity while still trading in terms of occurrences. Viewed from the Kaplanian perspective, the dynamic semantics we presently consider recognizes a privileged contextual parameter, the information state, and allows that parameter to shift as it is updated. Occurrences of sentences are still evaluated at a single input context, but the output of that update (and the input of subsequent updates) can be different. Every other contextual parameter may well still be handled in Kaplanian fashion. This difference, however, is enough to reject context fixity.

Recognition that context fixity is an idealizing assumption, coupled with the observation that dynamic semantics rejects it, does not itself constitute a reason to reject the idealization. However, many of the paradigmatic motivations for dynamic semantics are tacit arguments against context fixity. It is worth

making these explicit. It is agreed by just about all parties in the literature that epistemic modals express possibility relative to some body of information. Disputes typically concern precisely how to model this shared intuition. Kratzerians think that the information state with respect to which the modal is relativized comes from the contextually supplied modal base.[20] Dynamic theorists think that it depends on the input state which the modal updates.

Assuming that the information state that the modal is sensitive to is contextually supplied, recall that (16) is marginally assertible in some circumstances. However, should one add a second modal, felicity drops dramatically.

(32)   # Adam might be at the door, and he isn't, and he might be.

(32) strikes as straightforwardly contradictory. Unlike (16), one cannot imagine even complex scenarios where (32) would be assertible. On the assumption that the information state that the modal is sensitive to is contextually supplied, and further assuming context fixity, the right instance of the modal should receive the same contextually supplied value as the left. If this were the case, then no new information would be added by the second modalized claim, and we would expect(32) to be consistent. We would similarly expect (16) and (32) to be equivalently assertible.[21] This is not what we observe, and the second modal in (32) seems to convey new information over the first which results in inconsistency and decreased felicity. This suggests that it is evaluated relative to a different context from the first, despite inhabiting the same sentence. If this assessment is correct, then it seems to be a straightforward case where context-fixity fails. Thus, when reasoning about epistemic modals, dynamic semanticist think that we should reject context fixity. This is made all the more convincing by the fact that once we do we can explain this discrepancy easily. Similar variants of this argument can be made that consider the difference in felicity between the ordering of conjuncts of epistemic contradictions. Similarly, should we extend our consideration beyond this fragment of dynamic predicate logic, we can make similar arguments involving pronominal anaphora. It should thus be little surprise that dynamic semantics is attractive in scenarios where context fixity falters, since the rejection of context fixity allows us to make predictions unavailable in its presence.

Thus, dynamic semantics not only rejects context fixity but is motivated in doing so. Since we should expect semantics which reject context fixity to violate Non-Contradiction and Reflexivity, we should expect dynamic semantics to violate both of these laws. This is exactly what we observe as neither hold in dynamic semantics. It may be obvious that a semantics which trades in context change will reject context fixity. What may be less obvious is that when context fixity is viewed as an idealizing assumption, rather than an inviolable pillar of theorizing, it vastly changes what rejecting the principle amounts to. Dynamic semantics is often seen by supporters and critics alike as exotic, esoteric, and

---

[20]See Kratzer (1977) and Kratzer (2012).

[21]It is not implausible that the repetition of the modal may yield pragmatically generated infelicity, but this would not be enough to account for the intuition that (32) is catastrophically infelicitous.

thoroughly punk rock, given its deviance from classical logic. It may be tempting to view the program this way when its characteristic feature, the rejection of context fixity, is fundamental to the canonical position. As much as I might enjoy this characterization, it is important to appreciate that the theoretical choice to reject context fixity seems far more temperate and benign once context fixity is recognized as not only an idealizing assumption, but a false one. This should not only change the way that we view and market dynamic semantics; it should also inform our understanding of the non-classical nature of dynamic entailment. By Kaplan's own lights, holding different idealizing assumptions will yield different notions of validity. Accordingly, recognizing context fixity as an idealizing assumption not only makes the dynamic approach to entailment seem less fringe and extreme, but it also makes the classical approach significantly less sacrosanct.

4   Context and Contradiction

I have argued that we should expect failures of Non-Contradiction and Reflexivity in our dynamic framework. This does not, however, vindicate the specific failures of Non-Contradiction that dynamic semantics predicts. There is a profound disanalogy between between the spatial and temporal changes in context which concerned Kaplan, and the changes in information states operative in dynamic semantics. Dynamic semantics concerns occurrences of sentences and thus, sentences will be evaluated (as updates) at discrete contexts. It does not recognize the spatial and temporal shifting that made utterance validity so difficult to handle. Instead, within dynamic semantics, the only way to change an information state is to update it with a formula. This places a significant constraint on when we can expect failures of Non-Contradiction. Consider formulae of $\mathcal{L}_1$ of the following form:

(33) $\varphi_1 \wedge \neg\varphi_2$

The kinds of failures of Non-Contradiction that concerned Kaplan would involve some temporal or spatial change that occurred between when φ_1 was uttered and the time that φ_2 was uttered. Kaplan's intuitions, however, can be generalized, and needn't involve time or space per se, and we should expect failures of (33) in cases where context, in whatever capacity, changes enough between update with φ_1 and update with φ_2 such that φ_1 at its context does not entail φ_2 at its context. Spatial and temporal changes to context are not recognized by our semantics, and thus, the kind of change that must be brought about between φ_1 and φ_2 that would result in a violation of Non-Contradiction can only come about by update. The clause for conjunction tells us that to update a state s with (33) is to update s with $[\varphi_1][\neg\varphi_2]$. What this means is that the only way for the initial context s to change before update with $\neg\varphi_2$ is by update with φ_1. Thus, without the help of changes in time and space, $[\varphi_1]$, itself, must change s enough so that $s[\varphi_1]$ can sustain update with $[\neg\varphi_2]$ without crashing. Any state s.t. $s \vDash \varphi$ cannot be updated with $\neg\varphi$ without crashing. Thus $[\varphi_1]$

must change the state enough s.t. $s[\varphi_1] \nVDash \varphi$. Slightly less formally, update with φ has to change the state enough that the output state does not support φ. This means that we should only expect failures of Non-Contradiction when substitution instances of φ fail to satisfy Reflexivity, i.e., $\varphi \nVDash \varphi$. This is to say that we would only expect (33) to be consistent in cases where φ does not entail itself.

Based on the intuitions about context change above, we should only expect the consistency of $\varphi \wedge \neg\varphi$ when φ violates reflexivity. It was also established, in §2, that $\varphi \wedge \neg\varphi$ is consistent if and only if $[\varphi]$ is not idempotent. These cases exactly coincide since it follows trivially from the definitions of dynamic entailment and idempotence that $\varphi \nVDash \varphi$ if and only if $[\varphi]$ is not idempotent. Thus, the same intuitions about context change which motivate the canonical view vindicate dynamic semantics' predictions about failures of Non-Contradiction. Namely, we should expect the consistency of $\varphi \wedge \neg\varphi$ precisely when φ does not entail itself. Such cases, fairly intuitively, coincide with cases where repeated update with φ can bring about context changes over and above those brought about by the first. This explains why we should expect failures of Non-Contradiction to be directly associated with non-idempotent update. Such failures match, not at all coincidentally, with the failures predicted by dynamic semantics.

The intuitions that motivate dynamic semantics, as well as the formal predictions it makes, dovetail perfectly with the Kaplanian intuitions about context change once we are willing to reject context fixity. The difference between the dynamic and the Kaplanian approach is that in dynamic frameworks, meanings themselves, and not only non-linguistic forces, can alter the contexts in which subsequent formula are evaluated. This allows for a separate avenue, unappreciated by Kaplan, for contexts to be changed by semantic forces alone. Kaplan's claims about meanings through changing contexts, however, remain pertinent. This is to say that the very failures of classical laws that Kaplan feared, are the very same violations of classical laws that dynamic semantics predicts and embraces. These insights provide a response to the first challenge that dynamic semantics faced. Even in the absence of motivating sentences glosses, we observe theoretical motivation for the rejection of Non-Contradiction in dynamic semantics. The second challenge remains.

5 Felicity Judgments

The previous sections explained when, and why, we should expect failures of Non-Contradiction from a theoretical perspective. While these claims may be sound, it remains less than clear how they can be fruitfully applied to natural language. The following, as well as similar violations of Non-Contradiction, demand substantial explanation.

(26)   ?? (Adam might be in his office, and Adam is not in his office) and it is not the case that (Adam might be in his office, and Adam is not in his office.)

(27) ?? (It might be raining and it's windy) and it's not the case that (it might be raining and it's windy).

(23) $(\Diamond p \wedge \neg p) \wedge \neg(\Diamond p \wedge \neg p)$

(24) $(\Diamond p \wedge q) \wedge \neg(\Diamond p \wedge q)$

Mandelkern says the following about sentences like (26) and (27):

> [Their] complexity notwithstanding, [(26) and (27)] seem to be perfectly inconsistent, like any sentence of the form $\varphi \wedge \neg\varphi$, (Mandelkern, 2020, pp.388).[22]

Mandelkern claims that sentences like (26) and (27) are contradictory in virtue of their form. Mandelkern's intuitions strike as legitimate to the extent that context fixity holds. However, as I have argued, in the absence of context fixity, we need not, and indeed should not, assume that sentences of the form $\varphi \wedge \neg\varphi$ are contradictory. If I am correct, then dynamic semantics need not be altered based on Mandelkern's stated criticism. However, while Mandelkern does not explicitly mention felicity judgments concerning sentences like these, one may nonetheless consider (26) and (27) to be infelicitous. On this latter point, it is difficult to earnestly disagree, and one may worry that dynamic semantics fails to appropriately capture our linguistic intuitions. I take it that dynamic semanticists remain obligated to predict the infelicity of (26) and (27). However, they are not obligated to make this prediction in virtue of the inconsistency of these sentences. This allows greater room for maneuver than Mandelkern suggests.

To explain this infelicity, I pursue a strategy considered, but ultimately rejected, in Mandelkern (2020). This strategy faces two challenges, one articulated in Mandelkern (2020) and one novel, each of which I attempt to diffuse. The strategy is to appeal to the incoherence of sentences like (26) and (27) to predict their infelicity.[23] While some sentences of the form $\varphi \wedge \neg\varphi$ are consistent, all such sentences are incoherent, (Mandelkern, 2020, pp.388). Thus, (26) and (27) are incoherent, and we can appeal to this to capture their infelicity. More generally, dynamic semantics avails itself of at least two notions of entailment, as well as consistency, each of which captures what van Benthem calls different "styles of inference." A proper solution to the issues at hand requires appeal to

---

[22]I modify Mandelkern's quote to apply to my own examples. Mandelkern's chosen examples involve disjunction, which I do not presently consider. However, since his observations are taken to apply to any sentence of the form $\varphi \wedge \neg\varphi$, he is committed these judgments here as well.

[23]Mandelkern (2020) not only concerns Non-Contradiction, but also Excluded Middle and Object-Language Non-Contradiction ($\Rightarrow \neg(\varphi \wedge \neg\varphi)$). What I have called Non-Contradiction, Mandelkern labels Meta-Language Non-Contradiction. Mandelkern's primary objection to the appeal to coherence is that, while it helps avoid certain challenges faced by Meta-Language Non-Contradiction, it does nothing to address the problems associated with Object-Language Non-Contradiction and Excluded Middle, each of which concern logical truth, (Mandelkern, 2020, pp.389). The good news is that my present proposal avoids this problem, as I only discuss Meta-Language Non-Contradiction. The bad news is that separate arguments will be required to vindicate the other two principles.

both. As we move through my proposal, I hope to shed at least some light on the role they each play.

5.1 Unembedded Violations of Non-Contradiction

Explaining the infelicity of (26) and (27) by appealing to their incoherence leads directly to the first problem. Dynamic semantics is partly motivated by the difference in felicity between (16) and (17). The dynamic strategy is to say that (16) is marginally assertible, while (17) is not. This explanation appeals to the fact that (18) is consistent but not coherent, while (19) is neither consistent nor coherent. This is supposed to capture the fact that while (16) is infelicitous, it is not catastrophically so, as it remains marginally assertible. (17), however, is catastrophically infelicitous. The idea is that while (16) is incoherent, one can nonetheless update a conversational state or a belief state with this sentence, and not incur a crash. It is in this sense, and *only* this sense, that (16) is 'less bad' than (17). Appealing to this explanation also demands that (26) and (27) can, in some cases, be marginally assertible since, like (16), they are consistent, but incoherent. This claim is controversial; so much so that using it as motivation to reject Non-Contradiction would have been unconvincing to all but the most devout disciples of dynamic semantics. I nonetheless take this to be the correct judgment. Bolstered with the machinery of dynamic semantics as well as my previous arguments, I believe a compelling case can be made. To do this, we briefly consider the dynamic account of belief.

A popular account of belief in dynamic semantics follows Heim (1992). According to this theory, belief is the fixed point of update.

(34) **Belief**: an agent a believes some content $\varphi \Leftrightarrow I_a[\varphi] = I_a$

An agent believes some content φ just in case the agent's information state does not change upon update with φ, i.e. their state supports φ. We can see that any consistent yet incoherent update can be applied to certain states without crashing, but no, non-absurd, state (this includes the belief state of the speaker and the conversational state of the collective interlocutors) will support the information contained in that update. If we take it that one's assertions should represent one's beliefs, then, even in the absence of any knowledge of the speaker's state, we understand that a speaker who has asserted an incoherent sentence cannot rationally believe what they've asserted. Given roughly Stalnakerian assumptions about the relationship between belief and assertion, such utterances should strike as infelicitous since they are rationally unbelievable, and any healthy conversational state will not support them. Consider Veltman's example of this phenomenon in (14). It is first asserted that Adam might be at the door, and it is subsequently asserted that he is not. We do not thereby conclude that the speaker believes both that Adam might be at the door, and that he is not. Rather, once the discourse is complete, we understand, on the assumption of speaker honesty, that they believe that Adam is not at the door, and the conversational state should reflect this. This point can be illustrated with the following Veltman inspired scenario. Suppose Bernard and Carolyn are

expecting Adam. The doorbell rings and and Bernard rushes excitedly to the door, which he opens mid-utterance, only to find a package delivered. Adam is nowhere in sight.

(35)   a. Bernard: ?? Adam might be at the door... and...*door opens*... he's not.
        b. Carolyn: # You think Adam both might be at the door and isn't?
        c. Bernard: No, Adam isn't at the door.

Notice first that Veltman's intuitions about these kinds of scenarios fundamentally rely on non-ideal contexts. (14), (16), and (35) all rely on contexts where the speaker's information changes during the utterance.[24] This is to say that the motivating cases are those where context fixity does not hold. Next, observe that Carolyn's response in (35b), as well as any variations thereof, either suggest a failure to understand (35a) or strike as facetious and uncooperative. Bernard cannot, and does not, believe the content of (35a), and it makes no sense to ask whether he does, hence the infelicity of the question. Rather, after Bernard's utterance of (35a), we expect Carolyn to understand that Bernard believes that Adam is not at the door, despite the previous claim that he might be. Bernard's belief state changed, and Bernard's subsequent contribution to the discourse reflects this change. This is further supported by the "No" at the beginning of (35c) in response to Carolyn's question. This is exactly the prediction that dynamic semantics makes as any state updated with (35a) will support that Adam is not at the door and will not support the claim that he might be. (35a) and its merely incoherent compatriots sound infelicitous, at least in part, because the speaker can never, on pain of contradiction, believe the entire contents of their asserted discourse. The only state that does support incoherent sentences is the absurd state. Thus, uttering an incoherent sentence strikes as a violation of some epistemic or doxastic norm, even if the sentence uttered is consistent. The infelicity of utterances of incoherent sentences can thus be explained by the violation of this norm.

This also gives us insight into what is less bad about assertions of merely incoherent sentences. In cases of merely incoherent assertion, the conversational state, as well as the belief state of the speaker, can each be determined compositionally, without any process of accommodation or further inquiry. No further claims need to be made, and no further questions need be answered in order to ascertain the speaker's beliefs. Thus, the beliefs of an agent who makes a merely incoherent utterance are not *captured* by that utterance, but they are *communicated* by it. After the utterance of (35a), we understand that Bernard believes that Adam is not at the door, and we understand this based on the sentence meaning alone. Contrast this with (15), which is both inconsistent and incoherent. No state can survive update with (15) without resulting in the absurd state, and the speaker's interlocutors will need to carry out some further investigation to determine the speaker's beliefs and to get the conversation back

---

[24]Temporal changes occur as well, but the semantically relevant change is in Bernard's information.

on track. Thus, despite the fact that the speaker's merely incoherent assertions will never be supported by the output state or their information state, a perfectly good conversational state is still produced, and the beliefs of the speaker can be ascertained. This does not mean that utterances of merely incoherent sentences are felicitous. Rather, it means that they are not catastrophically infelicitous, insofar as the beliefs of the speaker, and the conversational state which those beliefs partly determine, are perfectly recoverable. This, I argue, is the characteristic feature of the heretofore nebulous notion of marginal assertibility, at least to the extent it is employed in dynamic semantics.

With this explanation in hand, dynamic semanticists can explain the discrepancy in felicity judgments reported in Knobe and Yalcin (2014). Data on sentence glosses make clear that while merely incoherent as well as inconsistent assertions were considered to be infelicitous, merely incoherent assertions were judged to be infelicitous substantially less often. Dynamic semanticists can neatly explain why merely incoherent updates are bad: they do not express the speaker's beliefs, while also explaining why they are less bad than inconsistent ones: the speaker's beliefs can still be recovered from the update.

Something like the above explanation of the infelicity of incoherent sentence utterances is operating in many dynamic frameworks. More controversially, however, we should expect this explanation to apply to (26) and (27). If they are merely incoherent, and thereby marginally assertible, then we should expect the beliefs of those who utter them to be recoverable. Indeed, they are. There are two arguments in favor of this claim: one direct, and one indirect.

5.1.1 The Direct Argument

Consider the following variation of (35).

(36)   a. Bernard: ?? Adam might be at the door... and...*door opens*... he's not.
       b. Carolyn: # You think Adam both might be at the door and isn't?
       c. Bernard: No, it's not the case that he might be at the door and isn't.

This scenario is identical to (35) except for Bernard's response in (36c). Here, Bernard denies that it is the case that Adam both might be at the door, and isn't. Observe that Bernard's response is in no way infelicitous. It directly answers Carolyn's question, and expresses the negation of something that Bernard does not, and no one could rationally, believe. Notice, however, that Bernard can do this despite the fact that (36c) is the negation of what he has previously uttered. Bernard's discourse involves an assertion of $\Diamond p \wedge \neg p$ followed by an assertion of $\neg(\Diamond p \wedge \neg p)$. Nonetheless, Bernard does not contradict himself, since his information state changed during the discourse, and he is not committed to the contents of his initial utterance of (36a). Despite the marginal assertibility of (36a), his beliefs are perfectly recoverable, and, as before, it is understood that Bernard believes that Adam is not at the door. From this, it

follows trivially that Bernard believes that it is not the case that Adam might be at the door, and is not. Much like (35c) in the previous scenario, (36c) is merely clarificatory, and the claim that it is not the case that Adam might be at the door and isn't is dynamically entailed by Bernard's belief that Adam is not at the door. This can be antecedently understood by anyone who understood (36a).

Importantly, Carolyn contributes no new information to the discourse, in the sense that her question does not eliminate any relevant worlds from the context set. This is to say that there is no relevant change in context between Bernard's utterance of (36a) and (36c).[25] Thus, there is nothing stopping us from eliminating Carolyn's contribution to the conversation and combining Bernard's discourse into a single consistent conjunction. This point can be illustrated with a final variant where Carolyn does not interject.

(37)   Bernard: ?? Adam might be at the door... and...*door opens*... he's not... and, just to be clear, it's not the case that he might be at the door and isn't.

(37) features stops and starts, as well contextual cues that indicate that Bernard's information changes during the utterance. Nonetheless, it has the same logical form as (26). Of course, the sentence sounds terrible; it is complex, incoherent, and the second half is redundant. But complexity, incoherence, redundancy, and even infelicity notwithstanding, it remains perfectly consistent. It is possible to understand after utterance of (37), that Bernard believes that Adam is not at the door. It is similarly possible to understand that Bernard believes that it is not the case that Adam might be at the door and isn't. The availability of these inferences is explained by the fact that $\Diamond p \wedge \neg p$ dynamically entails both $\neg p$ and $\neg(\Diamond p \wedge \neg p)$. There should be no inconsistency when Bernard subsequently asserts an entailment of his previous claim, even if it sounds ridiculous. While Bernard cannot believe its content, it remains marginally assertible, as Bernard's beliefs remain recoverable, both formally and intuitively, exactly as is in (36). While (26) lacks the contextual cues of (36), it has the same logical form, and should consistent as well. Despite their consistency, (36) and (26) remain incoherent and are infelicitous as a result.

5.1.2   The Indirect Argument

Recall that $\Diamond p \wedge \neg p$ can *never* be felicitously uttered twice in succession, as in (20). Rather, after the first update with $\Diamond p \wedge \neg p$, a second is doomed to be inconsistent with the resultant local context. The reason, once again, is that any context updated with $\Diamond p \wedge \neg p$ will support both $\neg p$ and $\neg(\Diamond p \wedge \neg p)$. Our Veltman inspired intuitions similarly inform us that hearers will understand, upon update with the left conjunct of (20), that the speaker believes both that Adam is not at the door and that it is not the case that he might be and isn't. Thus the story that explains why Bernard's beliefs are recoverable in (26) is the

---

[25]Carolyn's question may alter Bernard's belief state with respect to what he thinks Carolyn believes he believes, but these will not affect the consistency of Bernard's claims.

very same story dynamic semanticists already appeal to in order to explain the contradictory nature of (20). This provides independent motivation as well as further support for the claim that a speaker's beliefs are perfectly recoverable after utterances of (26) or (37). They remain marginally assertible, despite their infelicity and should, thereby, be treated as consistent within the dynamic framework.

The same arguments can be made in favor of (27), with an added twist. While (24) is consistent, the only states that can successfully update with it are those that include no worlds where it is both windy and rainy. Since hearers evaluate updates relative to their own information state, any hearer who has not ruled out the possibility that is both windy and rainy will not be able to update with (27).[26] However, if the hearer's information state excludes worlds where it is both windy and rainy, they can understand that the first conjunct expresses that it is windy and not rainy, much like the previous cases. They will then update with the second conjunct which negates the claim that it both might be rainy and is windy. In such cases hearers will once again be able to recover the speaker's beliefs, rendering the utterance marginally assertible.

Thus, while (26), (27), and other violations of Non-contradiction are complex, infelicitous, and flout a maxim or two, the speaker's beliefs remain, in principle, recoverable. They should thereby be predicted to be consistent. In general, we may suspect that the perceived felicity of a discourse may well be related to the ease with which one can recover beliefs from it. Stops, starts, perceptible changes in context, and interjections may make this process more manageable. This would explain why Bernard's beliefs in (35), (36), and (37) were perhaps more intuitively recoverable than they would be if he uttered (26). This may similarly explain why sentences like (26) and (27) may be perceived to have increased infelicity over the previously mentioned discourses. Nonetheless, the beliefs of the speaker are, in principle, recoverable in each case.

The scenarios that support the consistency of (26) and (27) require precisely the same context shifts that Kaplan aimed to avoid and Veltman embraced. Indeed, it is unclear, at least to me, how one could coherently maintain Veltman's intuitions about (14) and (15) while denying the above, as my reasoning is a direct extension of the reasoning involved in those scenarios. It should thus be clear that allowing for the consistency of (26) and (27) brings about no internal tension in the dynamic program, and does not prevent dynamic semantics from predicting the infelicity of such utterances. For those who do not share Veltman's intuitions this may be a bridge too far. However, the bridge is no further than those dynamic semanticists have already invited them to cross.

5.2 Embedded Violations of Non-Contradiction

A second issue arises when sentences like (26) and (27) are embedded under epistemic modal operators. This objection is raised in Mandelkern (2020), where it is observed that if a sentence is consistent, but incoherent, it can be "boot-

---

[26]This can explain any perceived increase in infelicity from (26).

strapped" to coherence with the addition of a wide scoping modal operator, resulting in a sentence that is both consistent and coherent, (Mandelkern, 2020, pp.389).

(38) # It might be the case that [(Adam might be in his office, and Adam is not in his office) and it is not the case that (Adam might be in his office, and Adam is not in his office.)]

(39) $\Diamond((\Diamond p \land \neg p) \land \neg(\Diamond p \land \neg p))$

The objection holds that (39) should be inconsistent, since it expresses an epistemic possibility towards a direct contradiction. As before, this claim is only as strong as the claim that Non-Contradiction holds, which, as I have argued, is only applicable when context fixity holds. Since dynamic semantics rejects this idealizing assumption, it appears to incur no obligation to treat (38) as inconsistent.

Some may nonetheless consider (38) to be infelicitous. The same can be said for other consistent, but incoherent, updates embedded under modals.

(40) # It might be the case that (Adam might be in his office, and Adam is not in his office)

(41) $\Diamond(\Diamond p \land \neg p)$

(41) is both consistent and coherent, so we will not be able to appeal to either directly to predict the infelicity of (40). Instead, we may adopt a principle which stipulates that epistemically modalized formula with incoherent prejacents should be infelicitous. Such a principle predicts the infelicity of (38), (40), and similar constructions while remaining independently motivated. Such a principle would be very closely aligned with extant principles already adopted by dynamic semanticists that treat incoherent formulas under certain embeddings as infelicitous. Willer (2013) for instance, adopts such a principle concerning suppositions, (Willer, 2013, pp.68-69).[27] Thus, despite the fact that dynamic semantics rejects Non-Contradiction, it can nonetheless predict the infelicity of embedded violations thereof in a principled fashion.

Thus, dynamic semantics is able to make sense of its controversial predictions concerning violations Non-Contradiction. Violations of Non-Contradiction are predicted to be infelicitous, but remain marginally assertible. These claims differ substatially from canonical positions, but are well-motivated and compatible with the broader dynamic program.

---

[27]Such principles are perhaps unsatisfying since they are not properly compositional. One can provide a more thorough response by modifying the the test clause for epistemic possibility. The clause presently considered tests for consistency with the state updated, but one may instead adopt an alternative clause that tests for coherence with the state updated. If such a clause is adopted, (39) and (41) are treated as inconsistent, and the infelicity of their natural language correlates can be predicted compositionally. There is further independent motivation to adopt this clause that goes beyond the scope of the present paper, as it fundamentally changes the semantics. I leave this for future research.

6   What it Might Mean and What it Doesn't Mean

I have attempted to demonstrate that failures of Non-Contradiction follow from the dynamic rejection of context fixity coupled with broadly Kaplanian considerations about context sensitive meaning. The result is that formula of the form $\varphi \wedge \neg\varphi$ are dynamically consistent precisely when $[\varphi]$ is not idempotent. I have further argued that these predictions are perfectly compatible with, and motivated by, the natural language intuitions expressed by Veltman involving changing information in non-ideal contexts. Thus, allies and critics alike should neither be confused nor surprised by these results.

I doubt such arguments will convince classical semanticists to jump ship. I do, however, expect these arguments to explain why dynamic semanticists remain similarly unmoved in light of criticism that dynamic semantics "fails" to validate Non-Contradiction. Brute insistence that dynamic semantics is obligated to treat $\varphi \wedge \neg\varphi$ as inconsistent are based on the tacit assumption of context fixity, a principle dynamic semantics rejects. I consider the ideas expressed in the preceding sections to have been operative, at some level, in the development of dynamic semantics over the previous decades. My hope is that, in making them explicit, I have shed some light on why dynamic semanticists have the commitments that they do, and why certain objections to the program seem to have fallen on deaf ears.

While I hope my arguments about context fixity have been convincing, I suspect my claims about marginal assertibility to be somewhat less well received. The utterances in question and the contexts they inhabit are complicated, the sentences are difficult to evaluate, and we have a tendency to retreat to theory in complex cases. I've made an attempt to explain why I have the intuitions I have, but in cases like these, you got 'em or you don't. I expect something of a standoff here, and I recomend a "spoils to the victor" approach to the issue.[28] This is to say that fighting over intuitions about increasingly exotic and esoteric sentences should perhaps take a backseat to discussion involving which idealizing assumptions we should adopt and when we should adopt them. This will allow us to more profitably and convincingly adjudicate disputes between competing theories.

Kaplan recognized that different notions of entailment arise from different idealizing assumptions. As such, our intuitions about entailment and consistency are similarly beholden to these assumptions. This is all to say that individuating semantic theories based on their idealizing assumptions should inform, at least in part, what standards they should be held to. Debates concerning idealizing assumptions commonly occur in engineering and the sciences. Semantics should follow suit. I expect such discussions will be fruitful. I also suspect that relating theories in terms of their idealizing assumptions might orient them in a way that is less antagonistic than previously thought. Distinct aims often demand equally distinct idealizations, and it may turn out that we can share in

---

[28]See Lewis (1994).

the spoils.

References

van Benthem, J. (1989). "Semantic Parallels in Natural Language and Computation". *Studies in Logic and the Foundations of Mathematics.* vol.129: pp337–375. doi: 10.1016/S0049-237X(08)70133-2

van Benthem, J. (1996). "Exploring Logical Dynamcis". CSLI Publications.

Dever, J. (2013). "The Revenge of The Semantics-Pragmatics Distinction". *Philosophical Perspectives* vol.27: pp.104-144. doi: 10.1111/phpe.12032

van der Does, J., Groeneveld, W.Veltman, F. (1997). "An Update on "Might". *Journal of Logic, Language, and Information* vol.6: pp.361–380. doi: 10.1023/a:1008219821036

von Fintel , K. Gillies, A. 2007. "An Opinionated Guide to Epistemic Modality". in T. Szabó Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 2). Oxford University Press. doi: 10.1111/j.0029-4624.2004.00485.x

Frege, G. (1918). "Der Gedanke". Beiträge zur Philosophie des deutschen Idealismus, vol.1: pp.58–77.

Gillies, A. (2004). "Epistemic Conditionals and Conditional Epistemics". *Nous* vol.38: pp.585-616. doi: 10.1111/j.0029-4624.2004.00485.x

Groenendijk, J. and M. Stokhof (1991a). "Two Theories of Dynamic Semantics". In *In Logics in AI, Volume 478 of Lectures Notes in Computer Science.* Springer: pp.55-64. doi: 10.1007/BFb0018433

Groenendijk, J. and M. Stokhof (1991). "Dynamic Predicate Logic". *Linguistics and Philosophy.* vol.14: pp.39-100. doi: 10.1007/BF00628304

Heim, I. (1992). "Presupposition Projection and the Semantics of Attitude verbs". *Journal of Semantics*, vol.9: pp.183-221. doi: 10.1093/jos/9.3.183

Lewis, D. (1980). "Index, Context, and Content". in Kanger, S., Ōhman, S. (Eds.) *Philosophy and Grammar.* Springer. doi: 10.1007/978-94-009-9012-8$_6$

Lewis, D. (1994). "Humean Supervenience Debugged". *Mind.* vol. 103: pp.473-490 doi: 10.1093/mind/103.412.473

Kaplan, D. (1989). "Demonstratives: an Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". In Almog, J. Perry, J. Wettstein, H. (Eds.) *Themes from Kaplan.* Oxford University Press: pp.481-563.

Knobe, J. Yalcin, S. (2014). "Epistemic Modals in Context: Experimental Data". *Semantics and Pragmatics* vol.7, Article 10: pp.1-21. doi: 10.3765/sp.7.10

Kratzer, A. (1977). "What 'Must' and 'Can' Must and Can Mean". *Linguistics and Philosophy* vol.1: pp.337–355. doi: 10.1007/bf00353453

Kratzer, A. (2012). "The Notional Category of Modality". in Kratzer, A. *Modals and Conditionals*: pp.27-69. Oxford University Press. doi: 10.1093/acprof:oso/9780199234684.001.0001

Lennertz, B. (2019). "Might-Beliefs and Asymmetric Disagreement". *Synthese* vol.196: 4775–4805. doi: 10.1007/s11229-018-1688-9

MacFarlane, J. (2011). "Epistemic Modals are Assessment Sensitive". In Egan, A. and Weatherson, B. (Eds.) *Epistemic Modality* Oxford University Press, 144-178. doi: 10.1093/acprof:oso/9780199591596.003.0006

MacFarlane, J. (2014). *Assessment Sensitivity*. Oxford University Press. doi: 10.1093/acprof:oso/9780199682751.003.0003

Mandelkern, M. (2020). "Dynamic Non-Classicality". In *Australian Journal of Philosophy* vol.98: pp.382-392. doi: 10.1080/00048402.2019.1624376

Potochnik, A. (2017). "Idealization and the Aims of Science". The University of Chicago press. doi: 10.7208/chicago/9780226507194.001.0001

Predelli, S. (2005). *Contexts: Meaning, Truth, and the Use of Language.* Oxford University Press. doi: 10.1093/0199281734.001.0001

Rothschild, J. Yalcin, S. (2016). "Three Notions of Dynamicness in Language". *Linguistics and Philosophy* vol.39: pp.333-355. doi: 10.1007/s10988-016-9188-1

Rothschild, J. Yalcin, S. (2017). "On the Dynamics of Conversation". *Nous* vol.52: pp.24-48. doi: 10.1111/nous.12121

Schlenker, P. (2009). "Local Contexts". *Syntax and Semantics* vol.3: pp.1-78. doi: 10.3765/sp.2.3

Skeels, P. (2023). "Dynamic 'Might' and Correct Belief". *Inquiry*: pp.1-46. 10.1080/0020174X.2023.2249946

Stalnaker, R. (1978). "Assertion". in Portner, P. and Partee, B. (Eds.) *Formal Semantics: Essential Readings*: pp.147-161. Blackwell Publishing. doi: 10.1002/9780470758335.ch5

Veltman, F. (1996). "Defaults in Update Semantics". *Journal of Philosophical Logic* vol.25: pp.221-261. doi: 10.1007/BF00248150

Willer, M. (2013). "New Dynamics for Epistemic Modality". *Philosophical Review* vol.122: pp.45-92. doi: 10.1215/00318108-2315324

Willer, M. (2015). "An Update on Epistemic Modals". *Journal of Philosophical Logic* vol.44: pp.835-849. doi: 10.1007/s10992-015-9364-8

Yalcin, S. (2007). "Epistemic Modals". *Mind* vol.116: pp.983-1026. doi: 10.1093/mind/fzm983

Yalcin, S. (2012b) "Bayesian Expressivism". In *Proceedings of the Aristotlean Society* vol.112 (2), 123–160.

Yalcin, S. (2015) "Epistemic Modality *de re*". *Ergo* vol.2: pp.475-527. doi: 10.1111/j.1467-9264.2012.00329.x

Zardini, E. (2014). "Context and consequence. An Intercontextual Substructural Logic". *Synthese* vol.191: pp.3473–3500. doi: 10.1007/s11229-014-0490-6