Proof and Context

Oscar Henry Sannen

A thesis submitted to fulfil the requirements of the degree of
Doctor of Philosophy.

2025

THE UNIVERSITY OF
SYDNEY

Faculty of Arts and Social Sciences,
Discipline of Philosophy

This research was supported by an Australian Government Research Training Program (RTP) Stipend Scholarship.

No content produced by generative artificial intelligence tools has been used in the preparation of this thesis.

I certify that the intellectual content of this thesis is the product of my own work and that all assistance received in preparing this thesis and all sources have been acknowledged.

Oscar Sannen.

Acknowledgements

Foremost thanks are owed to my primary supervisor Nicholas J.J. Smith for many years of keen supervision and support. Not once have I left Nick's office without a renewed appetite for philosophy. I also wish to express much gratitude to my auxiliary supervisor Mark Colyvan, as well as Michael Nielsen, Melissa Fusco, Rasmus Pedersen, Brigitte Claire Everett, Wendy Xin and everyone else who offered patient ears and helpful suggestions during the writing of this thesis.

For their love, I would like to acknowledge my parents, Iain and Fiona, my brothers, Toby and Sam and my partner, Chloe. To my grandmother, Janice, who is no longer with us: thanks for being the first to show me what words could do.

The truly correct proof is one
that strikes a harmonious balance
between strength and flexibility.

—*Yōko Ogawa*

Contents

Introduction

This thesis concerns some overlooked aspects of conversational contexts. It seeks to give an account of how assertions find their place against a background network of presuppositions. The core idea of the project is that these assertions are both harder and easier to perform than prevailing theories of context and presupposition make it seem.

On the one hand, this network of presuppositions is extremely flexible. Some philosophers have posited rules to the effect that contexts will make the changes necessary to accommodate what a speaker says *by default*. However, little has been done to explain how these changes will be facilitated when an assertion is vague, probabilistic or incompatible with other presuppositions in play. A full explication of these rules of accommodation must also be able to explain what occurs in these not uncommon cases. On the other hand, we should be skeptical of the tendency to treat these rules of accommodation as universal laws. It is simply not the case that speakers *must* take for granted all claims put forward during a conversation. When you suspect that I am lying about P, you might reasonably ask for me to dispel your doubts before you take my word that P for granted. We expect that speakers are in a position to verify—at least to our satisfaction—the things which they say. The resulting picture is that contexts are quite flexible, but not unlimitedly so.

The first half of this thesis concerns the flexibility of contexts. It provides a novel model of presupposition which allows a diverse range of assertion types to be accommodated. A central feature of this model is that it allows presuppositions to come in degrees. This way, the extent to which speakers take an assertion for granted is proportional to either (i) the strength with which the assertion is made or (ii) the confidence which the speakers have in that assertion. I begin by considering two received models of context from Lewis and Stalnaker and show that, by combining these models with concepts from belief revision theory, they can be made receptive to assertions which contradict or undermine presuppositions which are already in force. These models are then extended once more to include facts about each speaker's partial beliefs. These partial beliefs are used to construct a definition of partial presupposition where a speaker presupposes that P to degree n just in case they believe there is a common n-strength partial belief that P. The final section in this first half generalises the theory of context revision developed earlier in order to include the revision of these partial presuppositions. It is argued that updating the probabilities which underlie these partial beliefs ought to be mediated via imaging, as opposed to Bayesian conditioning. This is essentially because Bayesian conditioning is a probabilistic analogue of the naïve theories of context insofar as conditioning isn't capable of updating a probability measure with incompatible information.

The second half of this thesis concerns the recalcitrance of contexts. It provides a theory of *proof*, considered as a kind of speech-act, which has the effect of raising an audience's confidence in one's assertions. This notion of proof, which also comes in degrees, is used to associate different contexts with different levels of tolerance for accommodation. These levels of tolerance are, in turn, understood to be a measure of how difficult it is to introduce a presupposition in a given context. This measure is also applied in a consideration of legal standards of proof, where it is suggested that legal claims are only accommodated when sufficiently supported with non-probabilistic sorts of evidence. In this half, I also discuss how concerns about logical closure encroach on an adequate theory of proofs. These concerns are addressed by imbuing the model of contexts with tools necessary to represent speakers who have diverging presuppositions about laws of logic, which leads me to a notion of proof which isn't closed under entailment. Finally, I also address a series of Gricean constraints on the appropriateness of a proof (amongst other speech acts) in any given context. Whether a proof is appropriate depends on multiple factors, including (i) whether the proof is sufficiently strong, (ii) whether it is relevant for the purposes of the conversation, (iii) whether the proof is sufficiently (or overly) informative.

The theories of proof and context put forward in this thesis come together to paint the following picture: virtually any kind of assertion is one which can be accommodated in any kind of context, so long as its assertor is prepared to take on the relevant burden of proof.

Chapter 1

Partial Presupposition

Whether a speech-act is acceptable, successful, or even interpreted correctly in any given conversation depends on what is being presupposed. Unless it is presupposed that I have a bike, it is unacceptable for me to say that my bike was stolen. Unless it is presupposed that I have the means to repay a debt, my request for a loan will be unsuccessful. Unless it is presupposed that I am referring to my friend, and not a cow in the field named Mary, my assertion that 'Mary is ill' will be misinterpreted.[1] The upshot here is that proofs, considered as a category of speech-act, are similarly dependent on presuppositions for their success. If I attempt to prove John's innocence by showing that he was at work from 9am to 5pm, this demonstration will not work unless it is also presupposed that the murder took place outside of John's office during that timeframe.

An adequate theory of presupposition is also needed to account for how proofs change the context they are performed within. In many instances, the goal of a proof is immediate—defending the introduction of a new presupposition as a basis for further inquiry. In some cases, a group may decide to split the burden of proof amongst its members, so that each person performs their own proof in order to reach some broader conclusion. Habgood-Coote (2022, p. 1111) gives the example of a team of detectives trying to answer the question: *Who committed the murder?* Instead of every member of the team answering that broad question, they might each be given smaller 'subquestions' which together bear on the main goal of their inquiry: e.g., *What was the murder weapon? Was Mr. X at the scene of the crime? What does forensics say about the fingerprints?* and so on. When an answer is given to one of these subquestions it can then be used as a presupposition when working on other subquestions. If forensics shows the fingerprints were Mr. X's, then agents working on the question as to whether Mr. X was at the scene of the crime may use this assumption in their reasoning. Because of this, presuppositions are important for understanding not only when a proof is successful, but also what the outcome of that success *is*. The goal of this chapter will be to give an adequate characterisation of presupposition used for the analysis of proof offered in Chapter 2.

---

[1] This example is from Goddard and Routley (1973).

1.1 Some Existing Models of Context

In this section, we will cover two different models of contexts with respect to presupposition: Lewis' (1979) notion of conversational score and Stalnaker's (1999) notion of context sets. It's useful to examine both models because one is broadly syntactic in nature while the other uses a non-syntactic possible worlds approach. A common problem for both theories is their failure to address what might be called 'partial presupposition', which occurs when presuppositions are taken on tentatively due to unreliable or hedged assertions.

1.1.1 Lewis' Conversational Score

Lewis' (1979) 'Scorekeeping in a Language Game' proposes that all assertions in a conversation are evaluated for their acceptability relative to an imagined conversational scoreboard, much like how appropriate behaviour is determined by actual scoreboards in rule-governed games such as baseball. For instance, in baseball, it is only acceptable for a batter to walk to first base once four 'balls' have been thrown outside of the strike zone. The scoreboard keeps track of how many balls have been thrown to assess at any given time whether a batter's walking to first is appropriate. The linguistic equivalent, Lewis suggests, keeps track of items which are relevant to the assessment of assertions made in a conversation. Instead of recording runs, balls, strikes and so on, the conversational scoreboard contains abstract entities such as lists of presupposed sentences, rankings of objects in terms of their salience, a function which assigns referents to various names, a ranking of possible worlds in terms of their similarity to the actual world, etc. These items play the role of things like the 'offside line' in soccer or the 'strike count' in baseball in that they determine whether a conversationalist's acts are rule-abiding. For instance, if I say 'John couldn't make it' while one of the presuppositions on our current scoreboard says that John is here with us, then my assertion is unacceptable. Similarly, if I make the assertion 'That dog is happy!' in a context where there are no salient dogs, my assertion will be unacceptable (and incomprehensible), due to the scoreboard containing no salient dogs.[2]

---

[2] This second case shows how an assertion's unacceptability doesn't necessarily come from its being false. Instead, the assertion is unacceptable because there's not enough contextual information to determine a referent for the singular term 'that dog'.

The similarities between scoreboards used in games and those used in conversations are two-fold: the conversational scoreboard determines the meaning and acceptability of various assertions, but it is also updated by assertions in the process of a conversation in the same way that the offside line moves with the position of the soccer ball or the second-last opponent. Focusing on presupposition, Lewis' (1979, p. 340) suggestion is that presuppositions are added to the score according to the following rule of accommodation:

ACCOMMODATION If at time t something is said that requires presupposition P to be acceptable, and if P is not presupposed just before t, then—*ceteris paribus* and within certain limits—presupposition P comes into existence at t.

For example, if I say to you that my bike was stolen and it was not presupposed before I said this that I had a bike, then at the time of my assertion it is suddenly presupposed in order to make my assertion acceptable. For all assertions S, the content of S is included as one of the presuppositions required for S's acceptability. So, when I assert that "My bike was stolen", and this assertion is accommodated, the proposition that my bike was stolen is also added to the conversational score.

As well as requiring that presuppositions be *added* to the score, some assertions require that presuppositions be *removed* from the conversational score. If the presupposition that John has never smoked is on the scoreboard, then fully accommodating my assertion that "John has quit smoking" requires removing the presupposition that John has never smoked. This kind of presuppositional displacement doesn't follow from ACCOMMODATION (which only gives a condition for when presuppositions are added), but it might be achieved by proposing the following kinematic rule for conversational scoreboards:

DISPLACEMENT If at time t the presupposition P is successfully introduced to the score, then any presupposition which is incompatible with P is removed from the score at t.[3]

---

[3] While Lewis doesn't actually give DISPLACEMENT as one of the rules governing kinematics for conversational score, it follows the same general structure as a subsequent rule he gives for updating the comparative salience of objects (Lewis 1979, p. 349).

Here, a presupposition P' is incompatible with P just in case it's impossible for both presuppositions to be true simultaneously—i.e., they're mutually exclusive. This may be because P' is the negation of P, or because of some other kind of logical inconsistency between the two, but it might also be because of a 'metaphysical' inconsistency—e.g., if P is the presupposition that my bowtie is green and P' is the presupposition that my bowtie is red.

However, this simple account of displacement fails to account for the fact that sets of presuppositions, while maintaining pairwise compatibility, are sometimes jointly incompatible. For instance, if a group is committed to the list of presuppositions $\mathfrak{P} : \{P, P \to Q\}$ just before t, and someone makes an assertion which requires the presupposition that $\neg Q$ at t, then DISPLACEMENT doesn't rule out the addition of this presupposition so that at t the list of presuppositions is $\mathfrak{P}_t : \{P, P \to Q, \neg Q\}$. Each pair of propositions in $\mathfrak{P}_t$ is compatible, but the set as a whole is inconsistent, since there is no way of making all three true. Since the list of presuppositions is intuitively understood as a body of claims which speakers take to be true, a minimal requirement is that they can all be true simultaneously. This motivates a variation on DISPLACEMENT which ensures the compatibility of sets of presuppositions rather than pairs of presuppositions:

SET-WISE DISPLACEMENT Where $\mathfrak{P}$ is the set of presuppositions which is on the score just before time t, if an assertion is accommodated at time t which requires the presupposition that P, then at t the set of presuppositions becomes a consistent subset $\mathfrak{P}_t \subseteq \mathfrak{P} \cup \{P\}$

Notice that this definition is incomplete. It doesn't specify *which* consistent subset $\mathfrak{P}_t$ should become when it is presupposed that P. Whenever $\mathfrak{P} \cup \{P\}$ is a consistent set of propositions, then clearly it should be required that $\mathfrak{P}_t = \mathfrak{P} \cup \{P\}$, being a straightforward instance of ACCOMMODATION. However, when $\mathfrak{P} \cup \{P\}$ is inconsistent but has multiple consistent subsets, it's unclear in principle which presuppositions speakers will retain. Take the previous example where an assertion requiring the presupposition that $\neg Q$ is made against the background of a conversational score including $\mathfrak{P} : \{P, P \to Q\}$ as the list of presuppositions. $\{P, \neg Q\}$, $\{P \to Q, \neg Q\}$, $\{P\}$, $\{P \to Q\}$ and $\{\neg Q\}$ are all consistent subsets of $\mathfrak{P} \cup \{\neg Q\}$ but speaker presupposition alone doesn't provide enough information to determine which of these subsets a conversation will end up adopting for $\mathfrak{P}_t$.

This lack of information suggests that a full account of conversational kinematics requires an additional ‘abstract entity’ (besides lists of presuppositions) which facilitates the selection of consistent subsets of presuppositions. I suggest that the notion of an *epistemic entrenchment ordering* from AGM style belief revision theory (Alchourrón, Gärdenfors and Makinson 1985; Gärdenfors 1988; Lin 2019) is an excellent candidate for this entity. This ordering is some relation $\geq$ defined over the powerset of the set of all possible beliefs which intuitively ranks belief sets in terms of their indispensability.[4] Where B is a set of propositions and $X, Y \subseteq B$, $X \geq Y$ is interpreted as the claim that X is at least as epistemically entrenched as Y. While this original notion of entrenchment was devised to represent comparative commitment to various belief states, we will use it to represent a group of speakers' comparative commitment to various sets of presuppositions.[5] Thus, we will interpret $X \geq Y$ as saying that, for some group of speakers, the presuppositions in X are at least as important (for the purposes of their conversation) as those in Y. In the interest of keeping the conversation going, speakers will prefer to retain presuppositions which are more entrenched when forced to revise their conversational scoreboard. For example, if we are discussing the details of a play, then presuppositions about what occurred according to the play's narrative might be more indispensable than presuppositions about what occurred in the theatre hall; I might accommodate the presupposition that no-one was harmed, but not at the cost of abandoning the presupposition that Macbeth was slain by Macduff in a conversation about the Scottish play's themes of fate and greed.

Unlike belief revision theory, the presuppositional version of an epistemic entrenchment relation is defined as holding between sets of presuppositions. The main difference between belief revision theory's “belief sets” and sets of presuppositions is that belief sets are closed under logical consequence, which

---

[4] In belief revision theory, there are actually two distinct types of epistemic entrenchment relation. The type utilised here is a relation between sets of propositions following Gärdenfors' (1988, p. 80) gloss, whereas the other type is a relation between individual sentences in the language $\mathcal{L}$ (See Gärdenfors, 1988. p. 90; Fermé and Hansson, 2018, p. 31).

[5] Gärdenfors (1988, p. 93) claims that, outside of the metatheoretic cost-benefit analysis which goes on in science, one must look to conversational contexts to read off facts about epistemic entrenchment. So, this novel combination of belief revision and context revision is not entirely unprecedented.

means that for all belief sets B, if $B \vdash P$ then $P \in B$. In this respect, sets of presuppositions are closer to"belief bases", which are (potentially finite) sets of propositions B where an agent with belief base B believes P just in case $B \vdash P$ (Fermé and Hansson, 2018, p. 49). However, even belief bases may not be the perfect analogue because it's not clear whether Lewis (1979) held that a group of speakers presuppose all consequences of presuppositions listed explicitly on their conversational score. Regardless, belief sets and sets of presuppositions are both sets of sentences from some propositional language $\mathcal{L}$, so in either case $\geq$ is a relation over $\wp(\mathcal{L})$—it holds between sets of sentences from the language $\mathcal{L}$. We assume that $\geq$ is a total order on $\wp(\mathcal{L})$, such that it is transitive, antisymmetric, reflexive and connected. While this might be an unrealistically strong requirement for all conversational contexts, it is a minimum requirement for conversational scores which are rich enough to decide which subset of $\mathfrak{P} \cup \{P\}$ speakers end up adopting, which is the target phenomena here.

Given a set of presuppositions $\mathfrak{P}$ and a presupposition which is incompatible with P, the members of $\wp(\mathcal{L})$ which are of concern are the largest possible subsets of $\mathfrak{P}$ which are consistent with P. This captures the intuitive idea that the conversational score ought to retain as much information as possible when being updated. This intuition is central to AGM style belief revision theory, where it is called the principle of *informational economy* (Gärdenfors 1988, p. 49; Fermé and Hansson 2018, p. 13). Following Fermé and Hansson (2018, p. 21) call $\mathfrak{P}'$ a maximal subset of $\mathfrak{P}$ which fails to imply P just in case:

(i) $\mathfrak{P}' \subseteq \mathfrak{P}$

(ii) $\mathfrak{P}' \nvdash P$

(iii) There is no $\mathfrak{P}''$ such that $\mathfrak{P}' \subset \mathfrak{P}'' \subseteq \mathfrak{P}$ and $\mathfrak{P}'' \nvdash P$

The set of all maximal subsets of $\mathfrak{P}$ which fail to imply P is written $\mathfrak{P} \perp P$ in Gärdenfors' (1988) notation. Condition (iii) ensures the relevant sense of maximality by making it such that any proposition which is left out of $\mathfrak{P}'$ would make the set imply P once added. Once this set of maximal non-P-entailing subsets is established, the epistemic entrenchment relation ($\geq$) is used to define a function which selects all most-entrenched candidates in $\mathfrak{P} \perp P$ for the new set of presuppositions (Gärdenfors 1988; Lin 2019).[6] Wherever $\mathfrak{P} \perp P \neq \varnothing$ let this selection function γ be such that,

$$\gamma(\mathfrak{P} \perp P) = \{X \in \mathfrak{P} \perp P : X \geq Y \text{ for all } Y \in \mathfrak{P} \perp P\}$$

and $\gamma(\mathfrak{P} \perp P) = \{\mathfrak{P}\}$ otherwise. This function is the key to a full definition of DISPLACEMENT, which leverages *partial meet contraction* from belief revision theory (Gärdenfors 1988; Lin 2019).

PARTIAL MEET DISPLACEMENT Where $\mathfrak{P}$ is the set of presuppositions which is on the score just before time t, if an assertion is accommodated at time t which requires the presupposition that P, then at t the set of presuppositions becomes $\mathfrak{P}_t = \bigcap(\gamma(\mathfrak{P} \perp \neg P)) \cup \{P\}$.

Notice that ACCOMMODATION can be thought of as a special case of displacement where $\gamma(\mathfrak{P} \perp \neg P) = \{\mathfrak{P}\}$, either because $\neg P$ is a tautology (such that $\mathfrak{P} \perp \neg P = \varnothing$) or because $\mathfrak{P}$ is already consistent with P (such that $\mathfrak{P} \perp \neg P = \{\mathfrak{P}\}$. So, PARTIAL MEET DISPLACEMENT is general enough to underlie all kinds of presuppositional update which Lewis' theory countenances. To see how PARTIAL MEET DISPLACEMENT facilitates genuine displacement, consider the previous example where $\mathfrak{P} : \{P, P \to Q\}$ and someone asserts that $\neg Q$. Suppose the presupposition P is more indispensable than any other presupposition to the speakers. This can be represented with an entrenchment relation over $\mathfrak{P} \perp \neg\neg Q \subseteq \wp(\mathcal{L})$ where $\{P\} \geq \{P \to Q\}$, $\{P\} \geq \{P\}$ and $\{P \to Q\} \geq \{P \to Q\}$. Relative to this entrenchment ordering $\gamma(\mathfrak{P} \perp \neg\neg Q) = \{\{P\}\}$, so when $\neg Q$ is asserted at t the list of presuppositions on the conversational score is $\mathfrak{P}_t = \bigcap(\gamma(\mathfrak{P} \perp \neg\neg Q)) \cup \{\neg Q\} = \{P\} \cup \{\neg Q\} = \{P, \neg Q\}$.

In total, there are two distinct possibilities when an assertion S is made that requires the presupposition that P under Lewis' (1979) theory:

(i) Accommodation – P is added to the conversational score, accommodating S. Furthermore, by PARTIAL MEET DISPLACEMENT, propositions which are jointly incompatible with P (if there are any) are removed from

---

[6] Without the requirement that $\geq$ be a total order, it would be possible for these functions to return an empty set even when $\mathfrak{P} \perp P$ is non-empty. See Lin (2019, p. 363).

the conversational score.

(ii) Rejection – P is not added to the conversational score, either because someone explicitly objects to S, or because S is asserted by someone with poor epistemic standing with respect to S, such that S is not made acceptable merely by assertion.

There are two further sub-possibilities to consider in the case of Rejection:

(a) Simple rejection: S is rejected in such a way as to leave it open whether S, i.e., neither S nor not-S is presupposed. (Example: Someone says, "I have a winning lottery ticket!" and someone else responds, "There's no way you could possibly know that.")

(b) Negative rejection: S is rejected because someone makes an unchallenged assertion of S's negation, or there is a pre-existing presupposition of S's negation which is not successfully challenged by the assertion of S. (Example: Someone asserts, "My bike was stolen" and someone else replies, "Your bike wasn't stolen – I just saw you lock it up outside.")

In both sub-possibilities, the presupposition P required by S is not added by S's assertion alone, but in the case of rejection via negation P may still end up being presupposed because the negation of S also requires P for its acceptability. In the bike example, for instance, "My bike was stolen" and "Your bike wasn't stolen" both require the presupposition that you had (at least at one point) possession of a bike for their acceptability.

1.1.2 Stalnaker's Context Sets

Stalnaker's (1999) theory of presupposition uses sets of possible worlds in place of a 'conversational scoreboard'. This model is motivated by the assumption that conversations are a cooperative game of inquiry, where speakers take turns eliminating possible worlds with their assertions in order to single out some particular world as actual (Stalnaker, 1984).

While Stalnaker's theory is also about group presupposition, it begins with an account of individual speaker attitudes towards propositions. Each speaker in the conversation is said to presuppose the proposition P if and only if they believe that P, believe everyone in the conversation believes that P and believe that all these beliefs are taking place (Stalnaker, 1999, p. 49). Essentially, an individual speaker presupposes P just in case they take P to be a common belief of parties to the conversation (See Lewis 1969, Fagin et al. 1995).[7] This captures the idea that presuppositions encode information which is available (or at least seems to be available) to all conversational participants – that which is taken for granted as common background. This feature is what distinguishes *pre*supposition from regular supposition or assumption. When the content of a presupposition only *seems* to some speaker to represent a fact which everyone takes for granted, then that speaker presupposes something which others don't. That is, agents can presuppose P where it is not a common belief that P. This happens, for instance, whenever an uninformed traveller drives on the wrong side of the road in a different country. The traveller believes that people drive on the right side of the road in the country they are in and also believes that everyone believes this. However, it may actually be a convention to drive on the left side of the road, in which case the traveller presupposes P where it is actually a common belief that $\neg P$.[8]

Relative to each agent, there is a context set $\boldsymbol{C}$ which is the set of all possible worlds which are compatible with that agent's presuppositions. Stalnaker (1999, p. 66) doesn't define exactly what these possible worlds are but argues they are necessary for any theory of conversation insofar as conversations utilise a notion of 'live alternatives' that speakers distinguish between with their assertions. A proposition P is compatible with one of these live alternatives just in case P isn't false in that alternative. A group of conversationalists operate under the presupposition that P if and only if P is true in every world of each conversationalist's context set.

It is possible for two speakers in a conversation to have different context sets, where one of them takes for granted certain possibilities that the other doesn't. This follows from the fact that agents can take P to be a common belief, thereby presupposing that P, even in instances where P is not commonly believed. When at least two members of a conversation have different context sets, then the context is said to be defective (Stalnaker 1999, p. 85). These defective contexts are undesirable for the conversationalists insofar as

---

[7] At this point, only a rough understanding of 'common belief' is necessary. In §1.3.1 below, we return to the concept and make it precise.

[8] Or rather, it *was* a common belief that $\neg P$ before the traveller entered this community of drivers, thereby making it neither a common belief that P nor $\neg P$.

they will cause assertions to have unexpected consequences. If Bob presupposes that tomorrow there will be a federal election in a conversation where others do not, assertions such as "We'll meet at the school when we go to vote" conveys different information to Bob then it does to others—since Bob will take this information to entail a meeting taking place tomorrow, whereas others will take this information to entail a meeting on some other day (or perhaps no particular day at all). Stalnaker (1999, p. 85) proposes that, because the enterprise of conversation is based on communication, conversationalists will work to ensure that they share the same context set and avoid miscommunication. When there is a single context set shared by all members of a conversation (i.e., they share the exact same set of presuppositions), the context is said to be nondefective, and we can talk about *the* context set as opposed to each individual's. Going forward, we will assume contexts are nondefective unless otherwise specified.

Given this picture of presupposition, how are presuppositions added to or removed from the context set? As is the case with Lewis' conversational score, the background set of presuppositions is generally changed in one of two ways. Either (i) some conspicuous event happens at the scene of a conversation, such that all possibilities where this event did not happen are removed from the context set or (ii) someone asserts that P and this assertion is not challenged, such that all possibilities incompatible with P are removed from the context set (Stalnaker, 1999, pp. 86—87).

Note that these both describe circumstances for adding presuppositions to the context set, but they do not on their own explain how presuppositions are to be removed. Stalnaker (1973, p. 455) does suggest that presuppositions can be removed from the context set—for instance, when they are "denied, challenged, retracted or forgotten" but doesn't explain how the context set is to be updated once a presupposition has been removed. Indeed, subsequent literature has classified Stalnaker's theory of conversation as 'propositional', meaning that "conversational update is always just a matter of adding some proposition to the [common background]" (Rothschild and Yalcin, 2017, p. 27). Evidently, if Stalnaker's theory *is* propositional in this sense, then there is no way of updating a context set which *removes* propositions from the status of being presupposed.

According to Stalnaker (1973, p. 452), assertions are speech acts which agents use to divide the context set into two, mutually exclusive and exhaustive sets: the set of worlds in which the proposition asserted is true and the set of worlds in which the proposition asserted is false. If someone makes an assertion that P, call P the set of all worlds in which P is true and $\overline{P}$ the set of all worlds in which P is false. This double-use of capital italic letters to represent propositions *and* sets of possible worlds is appropriate in the context of Stalnaker's (1984, 1999) theory because he defines propositions as functions from possible worlds to truth values. Thus, each proposition P can also be interpreted as a characteristic function which describes a set of worlds (also denoted P) at which the function P returns the value true. Strictly speaking, propositions are *not* sets of possible worlds according to this view, but the mapping between them allows this perspicuous notational shorthand.

According to Stalnaker's rule (ii), every unchallenged assertion of P made relative to some context set $\boldsymbol{C}$ will produce a new context set $\boldsymbol{C}'$, which is the intersection of P and $\boldsymbol{C}$. There are two types of assertion to consider which both remove the presupposition that P:

1. Elimination via negation: Someone asserts the negation of P, or else something which entails P's negation.
2. Simple elimination: Someone challenges the presupposition that P, but they do so without asserting P's negation. This could also be because the conversationalists simply forget that they are presupposing P, or because the speaker who asserted P retracts their assertion.

First, consider what happens when a context set is updated by the first kind of elimination. We can't simply say that the new context set will also be the intersection of $\boldsymbol{C}$ and $\overline{P}$ (i.e., the set of worlds in which the negation of P is true), because seeing as P is already presupposed according to $\boldsymbol{C}$, it follows that $\boldsymbol{C} \subseteq P$, which makes $\boldsymbol{C} \cap \overline{P} = \varnothing$. Having the context set be empty is undesirable for two reasons. Firstly, it removes the goal from inquiry-based conversation, which is to eventually single out some possible world as actual. If *no worlds* have been singled out as possible, then there is no way for the conversation to proceed. Secondly, given Stalnaker's model, whenever the context set is empty it follows that every proposition is presupposed, since the empty set is a subset of all sets. Speakers should be able to take back a presupposition without stopping the conversation or presupposing everything whatsoever, so a separate rule for eliminating presuppositions via negation must be introduced.

ELIMINATION VIA NEGATION Call P_t the set of all propositions presupposed or successfully asserted up until time t in the conversation. At any time t, the context set $\boldsymbol{C}$ is $\bigcap P_t$. If any proposition $Q \subseteq \overline{P}$ is successfully asserted at some time t' after t, where $P \in P_t$, then P is removed from the list of successfully asserted propositions, such that $P_{t'}$ includes Q but not P. At t' the context set changes accordingly to become the intersection of all propositions in $P_{t'}$.

Note that Q in the above definition is any proposition which entails the negation of P, which is why the rule is called elimination via negation. Like the DISPLACEMENT rule suggested for Lewis' conversational score, Q might just be the negation of P, but it could also be anything else which is logically or 'metaphysically' incompatible with P, where metaphysical incompatibility is understood as mutual exclusivity in terms of truth at possible worlds. Here's an example which utilises ELIMINATION VIA NEGATION.

NUMBER GAME

A: "I'm thinking of a number between 1 and 6" $\boldsymbol{C}_1 = \{1, 2, 3, 4, 5, 6\}$

B: "Is it an even number?"

A: "Yes." $\boldsymbol{C}_2 = \{2, 4, 6\}$

B: "Is it less than 3?"

A: "Yes." $\boldsymbol{C}_3 = \{2\}$

B: "So it's 2?"

A: "No – I take back what I said earlier, it's not even." $\boldsymbol{C}_4 = \{1\}$

B: "Ah, so it's 1?"

Every answer given by A in this example expresses some proposition which changes the context set, depicted on the right. First, A sets up the initial context set by claiming they are thinking of a number between 1 and

6. If each number represents the world where A is thinking of that number, then this initial context set is $\boldsymbol{C}_1 = \{1, 2, 3, 4, 5, 6\}$. Then A claims they are thinking of an even number. Call this proposition E. The set of worlds in which E is true is $E : \{0, 2, 4, 6, ...\}$. When A asserts E against the background of $\boldsymbol{C}_1$, the context set is updated (by intersection) to become $\boldsymbol{C}_1 \cap E$ which is $\boldsymbol{C}_2 = \{2, 4, 6\}$. Then, in this second context, A asserts the proposition $L = \{1, 2\}$. When intersected with $\boldsymbol{C}_2$, L delivers the context set $\boldsymbol{C}_3 = \{2\}$. Once the conversation has progressed to $\boldsymbol{C}_3$ at time t, P_t includes the propositions $\boldsymbol{C}_1$, E, and L. When A asserts the negation of E in context $\boldsymbol{C}_3$, the elimination via negation rule is followed so that E is removed from $P_{t'}$ and replaced with $\overline{E} = \{1, 3, 5, ...\}$. At t', the context set changes to become the intersection of all propositions in $P_{t'} = \{\overline{E}, \boldsymbol{C}_1, L\}$, which is $\boldsymbol{C}_4 = \{1\}$. This elimination via negation rule allows propositions to be removed from the common background without removing other information that is still pertinent to the conversation. In *NUMBER GAME*, agent B can continue presupposing that A is thinking of a number less than 3 even when A rescinds their answer about the number's parity.

Despite this elimination rule working in simple cases like *NUMBER GAME*, it breaks down when we consider assertions which are incompatible with a given context without entailing the negation of any proposition which has been accepted during the course of a conversation. Essentially, this elimination rule has the same problem as the naïve DISPLACEMENT rule given for Lewis' conversational score in §1.1.1. Consider a case where A has flipped two coins, so the space of possible worlds is $\{HH, HT, TH, TT\}$ where XY means coin 1 landed on side X and coin 2 landed on side Y.

COIN GAME

A: "I've flipped two coins. You will have to guess which is which."
$\boldsymbol{C}_1 = \{HH, HT, TH, TT\}$

B: "Is coin 1 heads?"

A: "Yes."
$\boldsymbol{C}_2 = \{HH, HT\}$

B: "Is coin 2 heads?"

A: "Yes."
$\boldsymbol{C}_3 = \{HH\}$

B: "So they're both heads?"

A: "No, sorry – at least one of the coins landed on tails."
$\boldsymbol{C}_4 = \varnothing$

Here, A's final remark is clearly incompatible with the context set at stage three (i.e., $\boldsymbol{C}_3$) but it fails to outright contradict any of A's assertions. Where t is a time occurring just before A makes their last assertion, $P_t = \{\{HH, HT, TH, TT\}, \{HH, HT\}, \{HH, TH\}\}$. Just after t, when A asserts $\{HT, TH, TT\}$, ELIMINATION VIA NEGATION rules that $P_{t'}$ become $P_t \cap \{HT, TH, TT\} = \varnothing$, since for all $P \in P_t$, $\{HT, TH, TT\} \nsubseteq \overline{P}$. To allow for A and B to rescind the presupposition that both coins are heads without abandoning their context set altogether, ELIMINATION VIA NEGATION must be amended.

Our goal, then, is to derive a procedure which takes speakers from one context set to another when a proposition is asserted which is incompatible with the prior context set. Due to a paper from Grove (1988), this sort of process has already been described for the purposes of presenting a possible-worlds formulation of AGM style belief revision theory. This presentation from Grove (1988, p. 159) is based on Lewis' (1973) systems of spheres. Call $\mathbb{S}$ a system of spheres centred on some set of possible worlds $\boldsymbol{C}$ when:

(S1) $\subseteq$ forms a total order on $\mathbb{S}$; for all $S, S' \in \mathbb{S}$ either $S \subseteq S'$ or $S' \subseteq S$.

(S2) $\boldsymbol{C} \in \mathbb{S}$ and $\boldsymbol{C} \subseteq S$ for all $S \in \mathbb{S}$.

(S3) Where $\boldsymbol{W}$ is the set of all possible worlds, $\boldsymbol{W} \in \mathbb{S}$ and $S \subseteq \boldsymbol{W}$ for all $S \in \mathbb{S}$.

(S4) If $P \cap S \neq \varnothing$ for some $S \in \mathbb{S}$, then there is a smallest sphere in $\mathbb{S}$ — denoted S_P — such that $P \cap S_P \neq \varnothing$.

Intuitively speaking, this subset relation over spheres in $\mathbb{S}$ plays the same role as the epistemic entrenchment relation $\geq$. $S \subseteq S'$ can be read as saying that S is at least as viable as S' when it comes to selecting a new context set (relative to $\boldsymbol{C}$). Note that this is different from the interpretation which Lewis (1973, p. 14) provides for the inclusion relation used in evaluating counterfactual conditionals, which reads $S \subseteq S'$ as saying that all worlds in S are at least as similar to those in $\boldsymbol{C}$ as are any from S'. An important difference between the two is that our notion isn't beholden to any pretence of objectivity: whether an alternative context set is viable for some group of speakers depends on facts about their preferences. Whether two worlds are similar 'metaphysically speaking' isn't interest-dependent in the same way. Speakers can't make a counterfactual conditional true simply by favouring antecedent worlds where the consequent is true, but speakers can make it true that they will adopt context $\boldsymbol{C}_1$ over $\boldsymbol{C}_2$ by favouring the former.[9]

Given a context set $\boldsymbol{C}$ and a system of spheres $\mathbb{S}$ centred on it, here is a new rule for ELIMINATION:

PARTIAL MEET ELIMINATION If at time t the context set is $\boldsymbol{C}_t$ and a speaker successfully asserts P at some time t' just after t, then at t' the context set becomes $\boldsymbol{C}_{t'} = S_P \cap P$.

This elimination rule involves two steps. First, presuppositions which are incompatible with P are removed from $\boldsymbol{C}_t$ by expanding to the smallest sphere which contains some worlds where P is true. Then, to accommodate the presupposition that P is true, all worlds where P is false are removed from this expanded set via intersection. In the special case where P is already compatible with $\boldsymbol{C}_t$ (such that $S_P = \boldsymbol{C}_t$), PARTIAL MEET ELIMINATION reduces to addition of the presupposition that P via intersection. This dual-step procedure was also characteristic of PARTIAL MEET DISPLACEMENT, which removes presuppositions which are incompatible with P before adding in P as a presupposition.

Returning to the *COIN GAME* problem, we assume that A and B adopt a system of spheres which represents the entrenchment of various possibilities

---

[9] The point is worth stressing here because it bears on a separate issue discussed in §1.4.3 below. There, we see how Fusco (2023, p. 461) has been hesitant to endorse a particular restriction on a process called imaging because it relies on a subjectively constrained function which selects 'close' possibilities. If closeness is understood metaphysically speaking, then this subjective constraint is problematic. However, there is no need for this hesitation when closeness is understood in an entirely subjective manner.

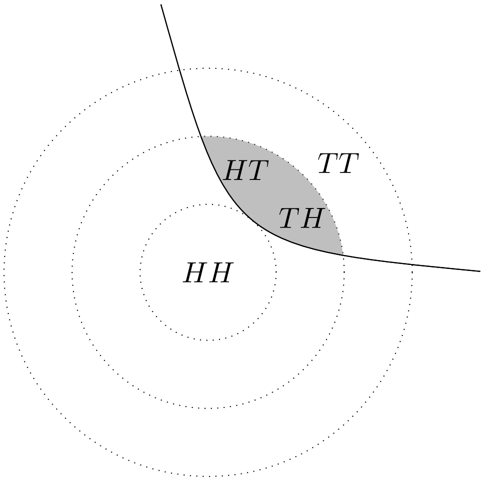

Figure 1.1: A sphere system centred on $\boldsymbol{C}_3$ representing *COIN GAME.*

according to their conversation, which is what will facilitate their elimination of some but not all of the presuppositions held at the time of A's incompatible assertion. This is a system of spheres $\mathbb{S}$ centred on $\boldsymbol{C}_3 = \{HH\}$, and suppose that worlds where one of the coins landed heads (i.e., where at least one of the other assertions A has made is true) are closer to HH than all others. This system can be depicted graphically as in Figure 1.1, where HT and TH are equally close to HH but strictly closer than TT. The curved line circumscribes a set of worlds in which A's final assertion is true—the claim that at least one of the coins landed tails. The shaded area indicates the new context set once A has asserted $P = \{HT, TH, TT\}$. Since $S_P = \{HH, HT, TH\}$, when A makes their final assertion the context set becomes $\boldsymbol{C}_4 = \{HH, HT, TH\} \cap \{HT, TH, TT\} = \{HT, TH\}$.

The second kind of elimination, simple elimination, takes place when a speaker removes commitment to a certain presupposition without asserting anything which contradicts the removed presupposition. This elimination corresponds to the first step in PARTIAL MEET ELIMINATION's two-step process – the widening of $\boldsymbol{C}$ to S_P. This might occur when speakers forget a particular presupposition they had made earlier, or when a speaker successfully calls for the retraction of an earlier presupposition without the addition of any new information. In these cases, the contextual change can be accounted for with the following rule:

SIMPLE ELIMINATION If at time t the context set is $\boldsymbol{C}$ and the system of spheres centred on it is $\mathbb{S}$ and the presupposition that P is successfully challenged, retracted or forgotten at time t' just after t – then at time t' the context set becomes $S_{\neg P}$, which is the smallest sphere around $\boldsymbol{C}$ where P is not presupposed.

It's worth briefly considering one more type of elimination, which might be called the 'nuclear option'. Sometimes a conversation will proceed in a way that is unacceptable to the conversationalists. Perhaps they have reached an absurd conclusion which they didn't anticipate, and this leads them to question all the assumptions they have taken for granted. Perhaps they are having trouble coordinating their context sets so as not to be defective. In these cases, it makes sense for the conversationalists to 'wipe the slate clean' and eliminate *all* presuppositions. This can be achieved by renouncing previous propositions and establishing some new context set $\boldsymbol{C}$.

1.2 Qualified Assertion and Partial Accommodation

Missing from these modified accounts of conversational kinematics is a strategy for situations where an assertion is made without full support from its speaker and situations where a speaker is only partially informed on the matter of their assertion. In this section, I spell out some ways these two types of assertion may happen and demonstrate why Lewis (1969) and Stalnaker's (1999) theories of context fail to represent them. The conversational moves discussed in this section suggest that one feature which context sets and lists of presuppositions both lack is a structure which allows for intermediate degrees of presupposition.

1.2.1 Qualified Assertion

Besides outright asserting some proposition P, it is possible for speakers to make various qualifications about P. Consider the following:

(1) It will probably be windy.

(2) There is a 75% probability of it being windy.

(3) It's fairly windy today.

(4) It's ... windy [said with hesitation].

Fully accommodating any of these four assertions requires that contexts neither add the presupposition that it will be windy nor remove it. Instead, accommodating these assertions requires that listeners appropriately adjust their confidence in its being windy. Assertions (1) and (2) are examples containing probability operators which are a focus of the literature on epistemic modals (Yalcin 2012; Swanson 2016; Moss 2018; Rudin 2018). These operators allow speakers to express claims about the likelihood of propositions. Accommodating an assertion which contains one of these probability operators requires redistributing a probability measure defined over sets of worlds in a context set (or alternatively, over propositions in a list of presuppositions), but often it needn't require the removal or addition of any of these sets of worlds (or propositions).[10] For instance, if someone asserts (1), then an accommodating response would be to increase the amount of probability bestowed on the claim that it will be windy, provided it wasn't sufficiently high already. If, instead of raising probability slightly, listeners add the presupposition that it will be windy to the context, they will have 'over-accommodated' (1) by adding strictly stronger information to the context. This over-accommodation eschews good faith communication since it opens (1)'s asserter to unfair scrutiny. It would be unreasonable to call them out for saying something misleading when the day turns out to be calm and windless. Furthermore, this over-accommodation will inevitably give rise to defective contexts. Since it is acceptable for A to assert (1) even when they are not certain it will be windy, whenever B over-accommodates by ruling out all possibilities where it's windy A and B's context sets will diverge. Alternatively, not adding the presupposition that it will be windy (and doing nothing else) amounts to ignoring the assertion of (1) entirely, which is clearly not a case of accommodating (1). The two rules for accommodation given so far – PARTIAL MEET DISPLACEMENT and PARTIAL MEET ELIMINATION – only describe methods for adding or removing presuppositions in their respective models of context. As such, they do not facilitate the accommodation of these qualified assertions, which require something intermediate. The goal

---

[10] See also Greco (2015, p. 185).

of §1.3 below is to provide an extension of these models which does so.

The difference between (1)-(2) and (3)-(4) is that the former make qualifications about the probability of its being windy whereas the latter make qualifications about the truth of its being windy. Since 'windy' is a vague term, it's possible not just to qualify the chance of wind but the intensity of wind as well. Note that if P contained no vague terms, it would still be possible to make probabilistic qualifications about P but intensive qualifications would no longer be appropriate. One can say "The next candidate is probably at least 195cm tall" but not "The next candidate is fairly at least 195cm tall". Much like how there is a spectrum of probabilistic operators along the lines of (2), Smith (2008) suggests there is a spectrum of utterances which express different degrees of confidence in the proposition being asserted. The degree of confidence behind these utterances is determined by hedging phrases (such as 'fairly' or 'sort of') and behavioural markers (such as hesitation before asserting, like in example (4)). If accommodating an unqualified assertion that S requires "taking S to be true", then accommodating an assertion like (3) or (4) requires "taking S to be true to a certain degree". Again, this kind of accommodation isn't tracked by either of the rules discussed in §1.1 since they only prescribe methods for taking propositions to be true or false, or dropping a commitment either way.[11]

Probabilistic qualification and intensive qualification are different kinds of qualification. As such, we should expect that they make different changes to the context in which they are performed. Contexts which accommodate (1) shouldn't automatically accommodate (3) and vice versa. This point is highlighted by the fact that both kinds of qualification can be combined to produce a unique effect on the context. Consider the following combination of (1) and (3):

(5) It will probably be fairly windy.

This kind of assertion can be understood as a proposal to furnish the common ground with a high attribution of likelihood to the event where the proposition that it's windy is fairly true. This furnishing is a composition of two operations on the context. First, worlds where it is fairly windy are added

---

[11] SIMPLE ELIMINATION is the relevant rule which accounts for when a proposition which was once taken as true is now neither taken to be true nor false.

to the context set (if necessary). Then, the probability of this set of worlds is increased relative to its complement (again, if necessary). A formal model for these kinds of complex changes in context will be given in §1.4.5 below, but the thumbnail given here demonstrates how the two kinds of qualification can be combined to perform distinct operations on the context.

1.2.2 Partial Accommodation

Another feature of conversations that is missing from both models of context is a response for situations where an assertion is made by a partially informed speaker. Suppose we are playing a game of pub trivia and a question is asked—in which city was the seventies rock band *Soyol Erdene* formed? One teammate, John, emphatically says without hesitation that the answer is Kathmandu. Another teammate, Chloe, challenges John and claims they are from Ulaanbaatar. In cases like these, it wouldn't be unusual to write down both answers while deciding which to go with. To ultimately facilitate a decision here, it would be prudent for us as a trivia team to consider the respective reliability of John and Chloe's assertions when it comes to musical matters. At this stage, before the team decides on an answer, it makes sense to understand the team as *partially accommodating* both assertions. They aren't taking either for granted, but perhaps they take it for granted that one of them is correct and their confidence needn't be equally distributed across both.

To extend Lewis' baseball metaphor, the scoreboard (of runs, balls, walks, etc.) is not the only thing which measures various players' contributions in a game – alongside this score is a running log of player statistics. One of these statistics is a player's 'on-base percentage', which is the proportion of plate appearances where the player reaches at least one of the non-home base plates. If a player's on-base percentage is poor enough, they won't be sent up to bat. Similarly, if a trivia player's answers are often wrong, the team won't automatically make the presuppositional adjustments needed to accommodate their answers. Instances like these occupy the 'certain limits' of ACCOMMODATION, where full accommodation isn't automatically granted to any speaker regardless of their conversational or epistemic standing.[12]

---

[12] Lewis (1979, p. 347) claims that conversations are accommodating by default, and that this is what makes them different from other species of rule-governed games. If a batter walks to first base after three balls have been thrown, then they've made an error and their walking to first doesn't suddenly make it the case that four balls were

A full account of partial presupposition must also include a theory of how partial accommodation of unqualified assertions leads to partial presupposition.

1.3 Degrees of Presupposition

Given the need for an account of partial accommodation and full accommodation of hedged assertions, we need a theory of how presuppositions can come in degrees of strength. The easiest way to do this is by changing a single detail of Stalnaker's account, while keeping the remaining details the same: simply change the requirement that presuppositions are what an agent takes to be common belief with the requirement that presuppositions are what an agent takes to be common *degree of* belief. In this section, we define degrees of belief, and then return in §1.3.1 and §1.3.2 to see how they can be used in Stalnaker's presuppositional pragmatics.

Following Smith (2010, 2014) and Jeffrey (1986), we define an agent's degree of belief that P as an expectation of P's truth-value. This means that beliefs are treated as a kind of confidence measure in various circumstances weighted by a proposition's degree of truth in each of these circumstances. The benefit of this approach, as demonstrated by Smith (2010), is its ability to clearly distinguish two kinds of psychological state which would be run together by a simpler account where degrees of belief are merely represented as probability functions over propositions: intermediate confidence in a proposition due to vagueness and intermediate confidence in a proposition due to uncertainty. Consider the proposition T expressed by 'Tom Cruise is bald.' If Tom Cruise is subjected to an experiment along the lines of Schiffer (2000), where a team of suspect experimental scientists pluck hairs from his head one-by-one, then at a certain point spectators will end up with a partial belief that T. Similarly, if an agent is told that Tom Cruise's hairdresser uses

---

thrown. This is unlike conversations, which tend to accommodate most assertions in order to make them acceptable. However, overstating this asymmetry obscures the fact that sometimes speakers' assertions are not so easily accommodated, as in cases of rejection or partial accommodation. Even in baseball, if a ball is thrown outside the strikezone but an umpire calls it a strike, then the scoreboard is changed to accommodate the umpire's assertion and a strike is added. While it may be true, on average, that conversations are more accommodating than other rule-governed activities, sometimes speakers don't accommodate and sports scoreboards do.

a coinflip to decide whether to give Tom a normal haircut or shave off all his hair, then after Tom Cruise's next visit to the barber they will have a partial belief that T. While the result is the same in both cases, the origin of these partial beliefs is different. It's also possible that agents have a partial belief stemming from a mixture of vagueness *and* uncertainty; for example, if a coin-toss is used to determine whether a significant amount or all of Tom Cruise's hairs are plucked. Defining degree of belief as expected degree of truth allows one to model all three scenarios differently, even when the particular degree of belief is the same in all cases.

This definition has two components: (i) a function which represents the degree of likelihood an agent places on various possible scenarios and (ii) a function which gives the truth-values of propositions in these scenarios. We take the first function to be a probability measure Pr which is a function from subsets of a set of possible worlds $\boldsymbol{W}$ to $[0,1]$ that satisfies the following three Kolmogorov Axioms for all $A, B \subseteq \boldsymbol{W}$:[13]

(K1) ADDITIVITY: If $A \cap B = \varnothing$, then $Pr(A \cup B) = Pr(A) + Pr(B)$

(K2) POSITIVITY: $Pr(A) \geq 0$

(K3) NORMALIZATION: $Pr(\boldsymbol{W}) = 1$

Intuitively speaking, $Pr(A)$ represents how likely the agent takes it to be that the actual world, @, is a member of A, where 1 means an agent is certain $@ \in \boldsymbol{W}$ and 0 means the agent is certain that $@ \notin \boldsymbol{W}$.

Here, I only consider agents whose probability measure is defined over sets $\boldsymbol{W}$ which are finite. There are a few reasons for this restriction. First of all, it's important that each possibility in $\boldsymbol{W}$ represents a circumstance which one could — at least in principle — prove to be the case. An infinite $\boldsymbol{W}$ simply contains more possibilities than there are opportunities for this kind of proof. Since proofs are a kind of act, and acts require resources which are inexorably finite (e.g. time, memory, chalk, etc.) performable proofs are outnumbered by possibilities from an infinite $\boldsymbol{W}$. This echoes

---

[13] Taking this likelihood function to be a probability measure follows Smith(2010), as opposed to Jeffrey (1986) who (like de Finetti (2017[1970], p. 12) takes *expectation* to be primitive, and then defines probability measures in terms of these expectations.

Ramsey's suspicion (1990[1926], p. 79) that when agents set themselves a question which has infinitely many distinct answers, in order to "consider the answers [they] must lump them into a finite number of groups." This holds even for extremely broad questions such as, 'What is the world like in every respect?' If there are finitely many ways of answering this question, then there are only finitely many possibilities which agents can consider during an inquiry. Note that it might not be the case that this broad question *has* finitely many answers. If worldly properties like length, temperature and mass are continuous then there will be uncountably many answers to such a question (e.g., at least one for every real-valued temperature between 0 and 100 degrees which a cup of coffee might possess). Instead, what is being suggested is that when it comes to the process of *giving* an answer there is a finite cap on the number of answers one could give. This line of argument is similar to Smith (2022, p. 5) who suggests that $\boldsymbol{W}$ must be finite because each possibility must be an outcome which agents could uniquely desire or support with their evidence to some degree.[14] If $\boldsymbol{W}$ were infinite, then agents would require implausibly rich powers of preferential and evidential discrimination, such as being able to say for any arbitrarily similar pair of temperatures which they'd prefer their coffee to be.

For those who are more confident in our powers of distinction, who think that agents could distinguish any two members in an infinite set of worlds, much of the theory developed below still holds, albeit with certain caveats. Firstly, the axiom of additivity (K1) might then be replaced with (K1)*:

(K1)* COUNTABLE ADDITIVITY: For any countably large collection of $\boldsymbol{W}$'s subsets A_i for i: 1, 2, 3 ..., if $A_i \cap A_j = \varnothing$ for all i and j, then $Pr(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} Pr(A_i)$

This axiom gives shows how the probabilities of certain sets relate to the probability of (countably) infinite constructions made out of them. Besides this modification, §1.3.1 below advocates an analysis of subjective possibility in terms of subjective probability, where the rough idea is that an agent takes some world w to be possible just in case they assign the event $\{w\}$ some non-zero likelihood. Levi (1989) shows that this reduction is inadequate

---

[14] Although Smith's (2022) argument is made in relation to non-probabilistic Dempster-Shafer belief functions, the same sentiment carries over to the probabilistic case—measures of plausibility or likelihood which are used for practical deliberation can't be spread infinitely thin.

when the probability measure in question is defined over an infinite sample space. First, suppose the sample space is uncountably large. Perhaps the players of NUMBER GAME expand their number-theoretic horizons and player A asks B to guess which number from $\mathbb{R}$ he is thinking of. For B to have a probability measure which respects the fact that A could've picked *any* number in $\mathbb{R}$, this measure would have to assign nonzero probability to $\{x\}$ for each $x \in \mathbb{R}$. However, doing so would lead to a function which doesn't respect the probability axioms, since (by (K1) *or* (K1)*) the probability of $\mathbb{R}$ would end up being greater than 1—violating (K3). Instead, suppose A and B play a modest infinite-numbered guessing game where A asks B to guess which rational number from $\mathbb{Q}$ he is thinking of. Unlike the uncountable case, here it is possible for B to adopt a probability measure which does assign positive probability to $\{x\}$ for each $x \in \mathbb{Q}$, but somewhat counterintuitively it will not be possible for this measure to be a uniform distribution. There will have to be a pair $x, y \in \mathbb{Q}$ such that $Pr(\{x\}) \neq Pr(\{y\})$, otherwise (by (K1) or (K1)*) the probability of $\mathbb{Q}$ would be greater than 1. This leads Levi (1989, p. 371), following de Finetti, to suggest that even in the countably infinite case, one should adopt a probability measure which assigns 0 to some events which are still taken to be serious possibilities. Namely, B should adopt a measure Pr such that $Pr(\{x\}) = 0$ for all $x \in \mathbb{Q}$ to reflect the fact that each number is equally likely. Of course, doing so requires abandoning countable additivity (Levi 1989, p. 380), since otherwise $Pr(\mathbb{Q}) = 0$. So, the force of this second part of Levi's argument will depend on how unimpeachable (K1)* is taken to be.

Levi (1989) and more recently Easwaran (2014, p. 4) ultimately suggest that a full representation of an agent's epistemic state should include *both* their probability measure and some specification of which worlds they consider possible. So, instead of simply being the probability measure Pr, B's epistemic state might be the pair $\langle Pr, \mathbb{R} \rangle$ where $Pr(\{x\}) = 0$ for all $x \in \mathbb{R}$ but the inclusion of each x in the set $\mathbb{R}$ signifies B's reckoning that A could be thinking of any real number. If you are convinced that agents can entertain infinitely many possibilities at once, then a full description of what presuppositions agents bring to the context will require use of these epistemic state pairs.

The other component of expected truth-values (besides a probability measure) is a valuation function, ν, which takes world-proposition pairs as input and delivers truth-values as output. That is, where $\mathcal{L}$ is a set of propositions in some logical language, and $\mathfrak{T}$ is some set of truth-values, $\nu : \boldsymbol{W} \times \mathcal{L} \rightarrow \mathfrak{T}$. Take $\mathcal{L}$ to be the language of propositional logic: that is, the set of well-formed formulae defined in the standard way using basic propositions (denoted with capital letters in italics), connectives ($\neg, \wedge, \vee, \rightarrow, \leftrightarrow$) and parentheses. As in Smith (2014, p. 1032) the logical details of ν, such as what values it assigns compound formulae, are left open for the moment. However, there is one restriction we do place on the set $\mathfrak{T}$ at this point. In order to represent situations where a partial belief results from vagueness without uncertainty, the set $\mathfrak{T}$ must include degrees of truth. To satisfy this condition, we will make $\mathfrak{T} = [0, 1]$, the continuum of truth-values from fuzzy logic.

Given these two functions (a valuation function ν and a probability measure Pr), an agent's degree of belief that P is defined as the expected truth-value of P. That is:

EXPECTED TRUTH-VALUE: $E(P) = Pr(\{w_1\}) \cdot \nu(\langle w_1, P\rangle) + ... + Pr(\{w_n\}) \cdot \nu(\langle w_n, P\rangle)$

Where $w_1, ..., w_n$ is an exhaustive list of all possible worlds in $\boldsymbol{W}$.[15] We define an agent's degree of belief in P as $E(P)$ and denote it as such. Notice that, when an agent is certain the world is a particular way—i.e., there is a w such that $Pr(\{w\}) = 1$—then $E(P)$ is identical to P's degree of truth at w. When this is the case, any partial belief is due to vagueness alone, and not uncertainty. Conversely, whenever $\nu(\langle w, P\rangle)$ is 1 or 0 for each $w \in \boldsymbol{W}$, then $E(P)$ is the same as whatever value the probability measure assigns to the set of worlds in which P is true. When this is the case, any partial belief is due to uncertainty alone, and not vagueness. Smith (2010, p. 500) calls these circumstances uncertainty-free and vagueness-free beliefs respectively.

1.3.1 Common Belief in Stalnaker's Theory

Before using expected truth-value as a substitute for belief in Stalnaker's model, it's important to give a precise definition of a concept which was roughly outlined earlier: common belief. The notion of a 'common' propositional attitude comes from Lewis (1969, p. 56), where he claims it is common

---

[15] Although we are dealing with a finite sample space $\boldsymbol{W}$, the expected truth-value of some propositions can still be given when the sample space is infinite. See Smith (2010, p. 499) for details.

knowledge that _____ amongst the members of some group G just in case there is a state of affairs A and,

> (1) Everyone in G has reason to believe that A holds.
>
> (2) A indicates to everyone in G that everyone in G has reason to believe that A holds.
>
> (3) A indicates to everyone in G that _____

While Lewis does call this attitude 'common knowledge', later on he suggested it might more accurately be called "overt belief" since, unlike knowledge, it might be common knowledge that P where P isn't true (See Lewis, 1978, p. 44). Regardless, the 'commonality' of Lewis' belief/knowledge attitude comes from the fact that, whenever it is common knowledge/belief that P in some group G, everyone in G believes/knows that P, believes/knows that everyone in G believes/knows that P, and so on up to an arbitrarily high order. Here, we will disambiguate by using common knowledge to mean this kind of commonality applied to knowledge and common belief to mean this kind of commonality applied to belief. Because of this, the objects of common belief are good candidates for presupposition; they form the corpus of propositions which agents in G have some mutually understood agreement upon.

In the model of presupposition developed by Stalnaker, these 'common' propositional attitudes can be represented by adding a tool from epistemic logic—in particular, what is needed is an accessibility relation for each agent in the group G. This is a binary relation R_i, which is a subset of $\boldsymbol{W} \times \boldsymbol{W}$ where $\langle x, y\rangle \in R_i$ if and only if the agent i thinks that the world y is possible relative to x. For instance, if w is a possible world where i turns off their desk lamp at 6:00pm, then R_i will most likely include the pair $\langle w, v\rangle$ where v is a possible world in which i turns their desk lamp on just after 6:00pm. In this case, R_i will not include $\langle w, z\rangle$, where z is a possible world in which i turns the desk lamp off just after 6:00pm because the desk lamp is already switched off. This relation is used to determine the agent i's various epistemic attitudes such as knowledge and belief, so it's an important component of any modal epistemic logic. However, to avoid multiplying theoretical entities beyond necessity, we can and will reduce this notion of an accessibility relation R_i to the previously introduced notion of an agent i's probability measure, Pr. At every world w, i will be in some kind of epistemic state dictated by their probability measure in w, which we will denote as Pr^i_w. I propose that the pair $\langle x, y \rangle$ is in R_i if and only if i's epistemic state in x is such that $Pr^i_x(\{y\}) \neq 0$.[16] This equates subjective possibility with subjective non-zero probability.

Note that the relation R_i itself is not world-relative. Whether the pair $\langle x, y \rangle$ is in the *global* relation R_i depends on *local* facts about the probability measure Pr^i_x held by agent i at world x. At each world x, the agent is in a doxastic state which describes (for them) which worlds are accessible from x and the relation R_i compiles all these local facts about accessibility into a relation over $\boldsymbol{W}$. It doesn't make sense, within this framework, to let facts like whether y is accessible from x to depend on what i believes in some completely unrelated world z and so the relation R_i isn't world-relative.

Earlier when we introduced the probability measure Pr for analysing degree of belief as expected truth-value, it wasn't indexed to a particular possible world. The received picture there was something like this: agents are presented with a set of possibilities $\boldsymbol{W}$ and they consider how likely they take each subset of $\boldsymbol{W}$ to be. This weighing up of likelihoods occurs *outside* of any particular world, so to speak. The only modification to this picture required to make sense of indexed probabilities involves dropping the assumption that agents can do this weighing up of possibilities outside of modal space; agents only distribute their confidence amongst worlds in $\boldsymbol{W}$ *at* a particular world in $\boldsymbol{W}$. Thus, an agent's probability measure varies depending on which world in $\boldsymbol{W}$ is actual. Adding this indexation allows us to represent agents who are uncertain about their current epistemic state. For instance, an agent might be uncertain as to whether they have high or low credence in some proposition P. By indexing an agent's probability measure to particular worlds, this scenario is represented by introducing two worlds (w and v): one in which the agent's expectation of P's degree of truth is high and the other in which it is low. A scenario where an agent is uncertain about their credence in P is simply a possible world where they assign a non-zero probability to both w and v. To account for these cases, instead of talking about an agent's probability measure in general, we talk about the agent i's probability measure Pr^i at some particular possible world w, and index Pr^i as Pr^i_w accordingly (likewise for $E^i_w(P)$). In contexts where the agent's probability measure is held fixed across all possible worlds, these

---

[16] A similar sort of framework with indexed probability measures for multiagent epistemic logic is briefly proposed in Shoham and Leyton-Brown (2008, p. 424).

world-relative indices will be dropped for reasons of perspicuity.

Admittedly, the value of such a framework depends on whether it is possible for an agent to be uncertain about their degree of belief that P, which in turn depends on whether an agent's degree of belief is a dispositional or introspective quality. If, following Ramsey (1990[1926], p. 65), you take an agent's degree of belief that P to be the strength of their disposition to behave *as if* P, then it is relatively straightforward to conceive of examples where an agent fails to know their credence in P. For instance, a blindly faithful sports fan may think they have a full strength belief that their team will win an upcoming match, but nonetheless fail to act in the ways which take this fact for granted – e.g., accepting any bet which is profitable contingent on them winning, not checking the TV to see if they actually win, etc. On the other hand, if degree of belief is a kind of introspectively accessible distribution of confidence in various propositions, then it is hard to see how an agent could be uncertain about their degree of belief that P. If one adopts this latter conception of partial belief, they can help themselves to the indexing of probabilities but they must do so with an added requirement that ensures agents assign 0 probability to all possible worlds where their probability measure is different.

INTROSPECTION: If $Pr_w(P) = x$ and $Pr_w(\{v\}) \neq 0$, then $Pr_v(P) = x$.

This restricted type of indexing with INTROSPECTION is still useful when modelling groups of agents who are uncertain about *each other's* beliefs, which is important for delineating the presuppositions a group is operating under. However, despite these indexed probabilities being amenable to INTROSPECTION, Ramsey(1990[1926], pp. 66—67) points out the implausibility of this constraint by noting that the beliefs which we hold most strongly are often accompanied by no introspectively accessible feelings which reliably inform that strength. Consider your belief that the air which surrounds you is breathable, which presumably you believe to a high (if not, the highest) degree. Beliefs like this are rarely cohabited by an apparent feeling of conviction—we *know* that we believe them so strongly only because we observe ourselves behaving as such. Going forward, the unrestricted indexing of probabilities is favoured for this reason.[17]

---

[17] Williamson's (2000, pp. 93—113) argument against the existence of mental states with a "luminosity" condition can also be construed as supporting the idea that agents can sometimes have a credence in P without knowing their credence in P. If there are no

When establishing an epistemic logic using the accessibility relation R_i, the following three rules are sometimes imposed to render R_i an equivalence relation:

REFLEXIVITY For all $x \in \boldsymbol{W}$, $\langle x, x \rangle \in R_i$

TRANSITIVITY For all $x, y, z \in \boldsymbol{W}$, if $\langle x, y \rangle \in R_i$ and $\langle y, z \rangle \in R_i$ then $\langle x, z \rangle \in R_i$

SYMMETRY For all $x, y \in \boldsymbol{W}$, if $\langle x, y \rangle \in R_i$ then $\langle y, x \rangle \in R_i$

This restriction is thought to "capture the intuition that agent i considers t possible in world s if in both s and t agent i has the same information about the worlds, that is, the two worlds are indistinguishable to the agent" (Fagin et al., 1995, p. 18). The sentiment here is that agents have some body of true or veridical information, and every world which is inconsistent with that information is ruled out of the space of possibilities. In this way, the accessibility relation has no structure beyond forming sets of worlds which are consistent with the body of information an agent has. However, for the purposes of modelling presupposition (which, unlike knowledge, doesn't rely on agent's body of true information) it is unnecessary to place these three restrictions on the accessibility relation R_i. Indeed, there are multiple scenarios we will want to represent where an agent's epistemic state doesn't conform to at least one of the three conditions above. Here is an example for each:

(i) REFLEXIVITY For this condition to hold, an agent i's probability measure must always be such that $Pr^i_x(\{x\}) \neq 0$. That is, in every scenario they will always believe that the scenario they are in is possible – or equivalently, they always consider the actual world as possible. However, suppose i finds a coin on the street and forms the belief that the coin is weighted in such a way as to always land heads. i flips the coin and covers it with his hand before he can see it. Call h the possible world where the coin is weighted and lands on

---

mental states such that an agent always knows they are in that mental state, then it is possible for an agent to be in the mental state of believing P to degree x without knowing they believe P to degree x.

heads and t the possible world where the coin is regular and lands on tails. If the coin turns out to be regular and lands on tails, then i's probability measure will be $Pr^i_t(\{t\}) = 0$, violating reflexivity.

(ii) TRANSITIVITY For this condition to hold, an agent i's probability measures must always be such that if $Pr^i_x(\{y\}) \neq 0$ and $Pr^i_y(\{z\}) \neq 0$, then $Pr^i_x(\{z\}) \neq 0$. Suppose that i is playing a game where a ball is hidden under one of three cups. $\boldsymbol{W} = \{w, v, u\}$ where w is the world in which the ball is under the left-most cup, v the middle cup and u the right-most cup. In world w, the hider of the ball tells i that there's an equal chance of it being in the left-most cup or the middle one, such that $Pr^i_w(\{w\}) = Pr^i_w(\{v\}) = 0.5$. In world v, the hider of the ball says there is a similar probabilistic split between worlds v and u. So, the agent's probability measures are such that $Pr^i_w(\{v\}) = 0.5$ and $Pr^i_v(\{u\}) = 0.5$, but $Pr^i_w(\{u\}) = 0$, violating transitivity.

(iii) SYMMETRY For this condition to hold, an agent i's probability measures must always be such that if $Pr^i_x(\{y\}) \neq 0$, then $Pr^i_y(\{x\}) \neq 0$. Take the same example from (i), where $\boldsymbol{W} = \{h, t\}$. In world t, i thinks h is possible, since $Pr^i_t(\{h\}) = 1$. For symmetry to be satisfied, i would also need to think that t is possible in h. However, this is not the case because $Pr^i_h(\{t\}) = 0$.

Given this relation which represents an agent's subjective possibility, we can then go on to define epistemic attitudes such as belief and common belief. We say that an agent i believes that P (denoted $\mathbf{B}_iP$) at a world w if and only if $\nu(\langle x, P\rangle) = 1$ for all x such that $\langle w, x\rangle \in R_i$, or equivalently, for all x such that $Pr_w(x) \neq 0$. The intuitive idea here is that an agent believes P in w whenever it is true in every world they consider possible from w – i.e., they take it as given that P is true no matter what. This equates belief simpliciter with full degree of belief, as the only way $\mathbf{B}_iP$ can be true at w is if $E^i_w(P) = 1$. This analysis is contentious, since some have insisted there is no direct relationship between "plain belief" and degrees of belief. Their concerns will be addressed and put to rest in §1.3.3 below, but at this point we will just assume the analysis is correct.

In Fagin et al. (1995), the truth-condition which we have provided for ascriptions of outright belief is given for ascriptions of knowledge, where 'i knows that P' is translated as $\mathbf{K}_iP$. However, from the perspective of the current account, these conditions are inappropriate for knowledge. This is due to two departures from in the account Fagin et al. (1995). The first point of departure is that we do not assume R_i to be an equivalence relation. This constraint is important for knowledge ascriptions, since otherwise it would be possible for an agent to 'know' P in situations where P is false, violating the factivity of knowledge: $\mathbf{K}_i P \rightarrow P$. The second point of departure involves an interpretation of the relation R_i, which Fagin et al. (1995, p. 18) take to encode information which is accessible to the agent. According to them, $\langle w, v\rangle \in R_i$ just in case the agent i has some information at w which fails to rule out v. This information is taken to be veridical information, which is what allows for i's state to constitute knowledge that P when P is true across all worlds which i deems possible. In contrast, we have imposed no such constraint on R_i. Whether the pair $\langle w, v\rangle$ is a member of R_i instead depends only on the (potentially uninformed or misinformed) opinions and beliefs of the agent i at the world w. To sidestep debates about whether knowledge requires more than just belief and veridical information, it's appropriate for the concerns of this project to drop any attempt to give truth-conditions for sentences of the form $\mathbf{K}_i P$ and instead be content with the idea that knowledge will ultimately require belief and some additional constraints which we won't attempt to determine.[18] After all, this project is about how proofs interact with the contexts in which they are performed. It is not an attempt to solve the gargantuan task of sorting those proofs which deliver knowledge from those which don't.

Following Fagin et al. (1995), we say a group of agents, G, all believe that P (denoted $\mathbf{E}_G P$) at some world w if and only if $\mathbf{B}_i P$ is true at w for all $i \in G$. For higher orders of group belief, we use $\mathbf{E}_G^0 P$ as an abbreviation for P, and $\mathbf{E}_G^{k+1} P$ as an abbreviation for $\mathbf{E}_G \mathbf{E}_G^k P$. The highest possible order of group belief is *common belief*, denoted $\mathbf{C}_G P$, and has the following truth-conditions:

$$\nu(\langle w, \mathbf{C}_G P\rangle) = 1 \text{ if and only if } \nu(\langle w, \mathbf{E}_G^k P\rangle) = 1 \text{ for } k = 1, 2, 3...$$

Now that we have a formal account of common belief, we can see the role it plays in Stalnaker's theory of context. His theory began with an account of

---

[18] Even Fagin et al. (1995, p. 8) are quick to qualify that while their analysis offers a reasonable truth-condition for ascriptions of knowledge, it is not the definitive account by any means. Instead, they suggest that different notions of knowledge will require different truth-conditions and humbly submit that their offering covers one.

individual speaker presupposition, where a speaker presupposes that P when they take it to be a common belief that P amongst their group of speakers, G. These individual presuppositions were then used to map out a context set $\boldsymbol{C}_i$ for each individual i, which is the set of all worlds which are compatible with i's presuppositions. These context sets form a non-defective context just in case $\boldsymbol{C}_i = \boldsymbol{C}_j$ for all $i, j \in G$, and a defective context just in case $\boldsymbol{C}_i \neq \boldsymbol{C}_j$ for some $i, j \in G$.

In converting these ideas to the epistemic logic framework, we will start by defining context sets and work backwards towards presuppositions. At any possible world w, the agent i's context set is the set of worlds which i deems possible from w – that is, $\boldsymbol{C}_i$ at w is the set of all x such that $Pr^i_w(\{x\}) \neq 0$, or equivalently, the set of all worlds x such that $\langle w, x\rangle \in R_i$. At any world w, the group G's conversational context is non-defective just in case $\boldsymbol{C}_i = \boldsymbol{C}_j$ for all $i, j \in G$ and defective just in case $\boldsymbol{C}_i \neq \boldsymbol{C}_j$ for some $i, j \in G$. Following Stalnaker, we say that an agent i presupposes P just in case they believe it to be a common belief that P amongst G – that is, $\mathbf{B}_i\mathbf{C}_G P$. If the context set is non-defective, then i's belief is true and it is a common belief that P, and therefore P is presupposed by everyone in G. If the context set is defective, then $\mathbf{C}_G P$ is false, in which case certain members of G besides i do not presuppose that P.

To see how Stalnaker's model works in the epistemic logic framework, here's a concrete example where two agents' presuppositions don't align, making a defective context. Suppose Alice is a driving instructor teaching Bob how to drive. Bob has red-green colour-blindness and can't distinguish the two colours. Alice is unaware of Bob's inability to distinguish red from green and tells Bob to stop at an upcoming red stop-sign. Call R the proposition that the stop-sign is red, and G the proposition that the stop-sign is green. In this scenario, Alice believes it is a common belief that R, and therefore presupposes that R when she tells Bob to stop at the red sign. The situation can be represented graphically as it is in Figure 1.2.

In this modal diagram, w is the world in which the relevant scenario is taking place, where the stop-sign is red and Alice informs Bob to stop there – so, $\nu(\langle w, R\rangle) = 1$, while $\nu(\langle w, G\rangle) = 0$. Alternatively, u is a world just like w except the stop-sign is green – so $\nu(\langle u, R\rangle) = 0$, while $\nu(\langle u, G\rangle) = 1$. Due to his colour-blindness, Bob's probability measure in w assigns equal likelihood to two possibilities: $Pr^B_w(\{w\}) = Pr^B_w(\{u\}) = 0.5$. This means

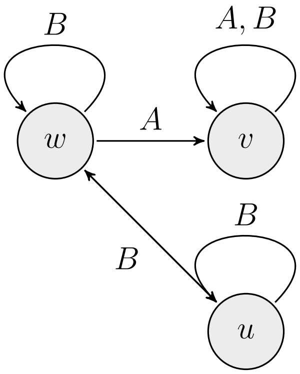

Figure 1.2: In this modal diagram R_{Alice} and R_{Bob} are represented as arrows, where there is an arrow from world x to world y labelled i just in case $\langle x, y\rangle \in R_i$.

that $\langle w, u\rangle$ and $\langle w, w\rangle$ are both members of R_{Bob}.[19]

On the other hand, v is a world just like w except that in v Bob is certain that the stop-sign is red, such that $Pr_v^B(\{v\}) = 1$. Perhaps this is because Bob isn't colour-blind at v, or he knows that all stop-signs are red. The only world which Alice thinks is possible, given her probability measure at w, is this world v, such that $Pr_w^A(\{v\}) = 1$. In v, it is a common belief amongst the group $G : \{Alice, Bob\}$ that R—such that $\nu(\langle v, \mathbf{C}_G R\rangle) = 1$. This follows from the fact that, in v, $\mathbf{B}_{Alice}R \wedge \mathbf{B}_{Bob}R$, $\mathbf{B}_{Alice}(\mathbf{B}_{Alice}R \wedge \mathbf{B}_{Bob}R) \wedge \mathbf{B}_{Bob}(\mathbf{B}_{Alice}R \wedge \mathbf{B}_{Bob}R)$, ... are all true. Because v is the only world accessible via R_{Alice} from w, Alice believes at w that it is a common belief that R (i.e., $\nu(\langle w, \mathbf{B}_{Alice}\mathbf{C}_G R\rangle) = 1$). As stated above, this is what it means for Alice to presuppose that R while instructing Bob. On the other hand, Bob does not presuppose that R, and this follows from the fact that $\mathbf{B}_{Bob}R$ is false at w, such that the lowest order requirement for common belief, $\mathbf{E}_G R$, is not satisfied at one of the worlds which Bob considers possible.

It's easy to see that the context is defective at w: Bob's context set $\boldsymbol{C}_{Bob}$ contains w and u, whereas $\boldsymbol{C}_{Alice}$ only contains v. This defective context here is what allows Bob and Alice's presuppositions to diverge regarding R. We would expect, as Stalnaker notes, that these defective contexts are prone

---

[19] Since Bob's colour-blindness goes both ways and he would have similar issues of discrimination if the stop-sig where green, $\langle u, u\rangle$ and $\langle u, w\rangle$ are also members of R_{Bob}, as shown Figure 1.2.

to failures of communication. Alice might tell Bob to stop at the red sign, and Bob will struggle to follow this instruction. However, the context set at v is non-defective, since there $\boldsymbol{C}_{Bob} = \boldsymbol{C}_{Alice} = \{v\}$. At v, we should expect that Bob has no problem at all when following Alice's instruction.

1.3.2 Common (degrees of) Belief

In the previous section, an epistemic logic was developed to model different kinds of speaker attitude towards propositions: attitudes like belief, knowledge and common belief. This allowed us to revisit Stalnaker's definition of presupposition in a modal framework. In this section, we give an account of partial presupposition by enriching this epistemic logic with an additional class of speaker attitudes – degrees of belief.

When an agent i believes that P – when $\mathbf{B}_i P$ is true – they are in a position where they have the highest possible degree of belief that P, and $E^i(P) = 1$. However, it's possible for an agent's expectation of P's truth to be less than 1, so it will be useful to introduce a new class of epistemic operators to describe these partial beliefs. When an agent i has an n-strength belief that P, we will express this as $\mathbf{B(}n\mathbf{)}_i P$, and give it the following truth-conditions:

$$\nu(\langle w, \mathbf{B(}n\mathbf{)}_i P\rangle) = 1 \text{ iff } E^i_w(P) = n$$

So, for example, in Figure 1.2 above, $\mathbf{B(}0.5\mathbf{)}_{Bob} R$ is true at w, since $E^{Bob}_w(R) = 0.5$. There can also be higher order beliefs about these degrees of belief. For instance, when rolling a die, some group G might have a common 0.5-strength belief in E, where E is the proposition that the die lands on 2, 4 or 6. This common degree of belief works just like common full belief, except the lowest order beliefs of each agent are partial. Where $\mathbf{C(}n\mathbf{)}_G P$ denotes that group G has a common n-strength belief that P,

$$\nu(\langle w, \mathbf{C(}n\mathbf{)}_G P\rangle) = 1 \text{ iff } \nu(\langle w, \mathbf{C}_G \mathbf{B(}n\mathbf{)}_i P\rangle) = 1 \text{ for each } i \in G$$

So, a group G has a common n-strength belief that P just in case they all believe P to degree n, fully believe that everyone in G has this n-strength belief that P, and fully believe that everyone fully believes that this is the case, and so on.

A natural way of adding partial presupposition to Stalnaker's account of presupposition developed in the previous section is by removing the tacit assumption that the propositions which agents take to be common beliefs are always believed to the highest possible degree.

PARTIAL PRESUPPOSITION An agent i presupposes P to degree n just in case they believe there is a common n-strength degree of belief that P amongst G. That is, i presupposes P to degree n if and only if $\mathbf{B}_i\mathbf{C}(n)_GP$ is true. A group G operates under an n-strength presupposition that P just in case it is true that everyone in G presupposes P to degree n.

In general, it will be possible for a group G to presuppose P to degree n even if some members of G are in different epistemic states with respect to P. This can be the case when some members of G have vagueness-free partial beliefs whereas others have uncertainty-free partial beliefs. Suppose, for instance, there are two agents A and B. A decides to trick B into playing an unfair coin-game using an old penny which has corroded to th extent that both sides simultaneously resemble heads and tails. Because of this strange corrosion, it will be true to degree .5 that the coin lands heads, no matter which of the two sides it lands on. Without seeing the coin, B believes it is fair, and in good enough condition to distinguish the two sides. This situation is depicted by the flow-chart in Figure 1.3.

Suppose w is the world where the penny is corroded, and H is the proposition that the coin lands heads. No matter which side the coin lands on, $\nu(\langle w, H\rangle) = 0.5$ because of its ambiguous faces. v and u are worlds where the coin is regular, and in v it lands heads whereas in u it lands tails, so $\nu(\langle v, H\rangle) = 1$ and $\nu(\langle u, H\rangle) = 0$. At w, A and B both presuppose H to degree 0.5, since their expectations of H's truth are both 0.5 and this holds in all worlds accessible from w by each of their respective possibility relations, so they both believe they have a common 0.5-strong belief that H. However, B mistakenly believes that this is because the coin is regular, and that A has joined B in assigning equal credence to the possibility of it landing on either side. This example shows that agents can have a different working model of the world even if they share the same degree of presupposition in the exact some propositions.

Instances like these look like a problem for the current account, since there

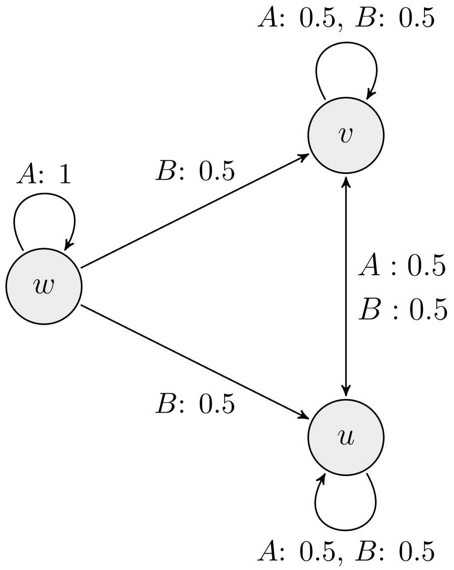

Figure 1.3: Corroded penny game

is a sense in which two agents can presuppose P to degree n for completely different reasons, thereby lacking shared or common ground with respect to P. However, this lack of shared ground can be explained by the fact that, while a group of agents might share the same degree of belief with respect to P and even take it for granted that everyone has this degree of belief, it is still possible for them to do so in a defective context, where the space of worlds which is taken for granted isn't shared. A truly nondefective context in the current framework requires that agents assign positive probability to all and only the same worlds *as well as* sharing the same degrees of presupposition in each proposition. At w, the context is defective since $\boldsymbol{C}_A = \{w\}$ and $\boldsymbol{C}_B = \{v, u\}$, so A and B lack shared ground concerning the proposition H despite both presupposing it to the same degree.

To summarise this section, I will conclude by giving a formal summary of *what* constitutes a context at this point. Contexts are 5-tuples $\mathfrak{C} = \langle G, \mathbb{P}, \boldsymbol{W}, @, \nu\rangle$ where

G is a set containing speakers engaged in the conversation.

$\mathbb{P}$ is a set containing probability measures Pr^i_w defined over $\boldsymbol{W}$, one for each pair consisting of a speaker $i \in G$ and a world $w \in \boldsymbol{W}$

$\boldsymbol{W}$ is a set of worlds

@ is a member of $\boldsymbol{W}$ called the *actual* world

ν is a valuation/interpretation function from $\mathbf{W} \times \mathcal{L} \rightarrow [0, 1]$

Furthermore, a context set $\boldsymbol{C}_i$ can be identified for each speaker in $i \in G$ relative to a given context $\mathfrak{C}$ which is simply the set $\{w : Pr^i_@(\{w\}) \neq 0\}$. In non-defective contexts, where $\boldsymbol{C}_i = \boldsymbol{C}_j$ for all $i, j \in G$, we can talk about *the* context set of $\mathfrak{C}$, denoted $\boldsymbol{C}_{\mathfrak{C}}$. We can change the context set(s) associated with a given context by selecting different members of $\boldsymbol{W}$ as actual.

1.3.3 Belief is Belief to the Highest Degree

In §1.3.1 I pushed for an analysis of belief where $\mathbf{B}_i P$ is true just in case $E^i(P) = 1$. In a slogan, this analysis equates belief with full degree of belief. Others have objected that the relationship between outright belief and degrees of belief is tenuous at best. Consider Spohn (2012, p. 47):

> [...] there is a deeply rooted tendency to project belief onto some scale of (un)certainty [...] On the other hand, the scale does not really fit; belief is clearly not maximal certainty, but also not any probabilistic degree below the maximum.

The intuition guiding this remark of Spohn's is that agents can supposedly believe P while doubting P, whereas assigning P the highest degree of probability constitutes being certain that P. Spohn suggests an agent might believe that there is milk in the fridge while still assigning some nonzero probability to the event that they forgot the milk carton's being empty, which by additivity and normalization means that Pr('There is milk in the fridge') $\neq 1$. Adopting a policy where belief is understood as degree of belief above a "threshold" just short of 1 is also marred by difficulty. Suppose you set the threshold as follows:

$$\nu(\langle x, \mathbf{B}_i P\rangle) = 1 \text{ iff } E^i_x(P) \geq 1 - \epsilon \text{ for some small } \epsilon > 0$$

Spohn (2012) uses Kyburg's (1983) presentation of the *lottery paradox* to show that any choice of ϵ for this threshold is too large. Consider a fair lottery which contains $n \geq \frac{1}{\epsilon}$ tickets from which one lucky winner will be drawn. Let A_i be the proposition that ticket number i wins. The fairness of the lottery guarantees that $Pr(\{i\}) = \frac{1}{n}$, where i is the world where ticket number i is drawn. This means that $Pr(\overline{\{i\}}) = \frac{n-1}{n}$ for each i, from which it follows that $E(\neg A_i) = \frac{n-1}{n}$. Since $\frac{n-1}{n} \geq 1 - \epsilon$ then according to any threshold analysis of belief, a rational agent a believes that ticket i will not win, for each ticket i – i.e., $\mathbf{B}_a \neg A_i$. However, by the same analysis, they also believe that at least one ticket will win because $E(A_1 \vee ... \vee A_n) = 1$. If belief is closed under conjunction, then this doxastic awkwardness is compounded by the fact that they will also believe the following contradiction: $(A_1 \vee ... \vee A_n) \wedge (\neg A_1 \wedge ... \wedge \neg A_n)$.[20] Thus, Spohn (2012, p. 44) concludes, a threshold analysis of belief leads to absurdity no matter what choice of $\epsilon > 0$ is made.

While I take the same moral as Spohn from the lottery paradox, I also think it's a mistake to go one step further and suggest that this kind of analysis is faulty even when $\epsilon = 0$. Someone who faces an arbitrarily large fair lottery doesn't believe of any individual ticket that it will lose, but this fact is best explained by an understanding of belief as full degree of belief. Contrary to Spohn's intuitions, the agent who assigns nonzero probability to the fridge being empty simply doesn't believe there is milk in their fridge. Instead, following a natural interpretation of believing P as "taking P to be true" which is even endorsed by Spohn (2012, p. 45), I suggest that believing P requires assigning no positive probability to worlds where P is anything short of true. Using a taxonomy from Jackson (2020), the view I am advocating, where $\mathbf{B}_i P$ is true if and only if $E^i(P) = 1$, can be categorised as a credence-first reduction of belief to credence 1.

This kind of response is aligned with Ramsey (1990[1926]) and Jeffrey (1970) who both defend a notion of partial belief where an agent's degree of belief is fully characterised by their behavioural dispositions. The strength of a belief that P directly corresponds to how strongly an agent takes P for granted in their actions. The strongest kind of belief possible with respect to P involves being disposed to behave as if P is true. This is marked by behaviour such as using P as a premise in one's reasoning, sincerely answering that P when asked whether P and so on. Intermediate degrees of belief are

---

[20] Kyburg's (1983, p. 233) solution to the paradox involves rejecting closure under conjunction for beliefs, but it's worth noting that there is something incoherent about agent a's beliefs even when they don't believe the explicit contradiction. Under a dispositional understanding of belief, the agent will be prone to contradictory behaviour like throwing out their ticket since they expect it to lose *and* holding onto it since it's a possible winner.

marked by more tentative dispositions such as a willingness to enter what are perceived as favourable gambles on P and sincere but partial assertions about P when asked whether P. One objection to the credence-first analysis of belief pushes on this distinction. They claim that sometimes we confidently assert (and thereby fully believe) propositions which we wouldn't be prepared to take arbitrarily risky gambles on (See e.g. Maher (1993, p. 133), Jackson (2020, p. 2)) However, I follow Jeffrey (1970, p. 161) in thinking that this sense of "full belief" or "certainty" is mysterious given our dispositional account of belief: if you hesitate to accept bets at arbitrarily short odds on P, then there is a sense in which you don't take P's truth for granted even if you go around claiming to believe P.[21] This view starts to make belief look like an increasingly rare attitude. If we believe only those propositions which we are prepared to accept arbitrarily risky bets on, then it appears as if we don't believe in much at all. Indeed, Jeffrey (1970) goes on to argue that theorists should get on by replacing attributions of belief that P with attributions of sufficiently high credence. However, unlike Jeffrey, one can maintain that this demanding requirement on full belief is compatible with its total ubiquity. People often enter gambles in an extended sense while being completely aware of risks which they voluntarily ignore.

A stronger objection comes from Maher (1993, p. 135) and Levi (2004, p. 9). They both attack the reduction by giving examples of agents who assign subjective probability 1 to an event which they don't fully believe. Suppose you are testing a compound for some continuous variable like mass, height or temperature. Relaxing our constraint on the finitude of $\boldsymbol{W}$, here your probability measure would be defined over some uncountable sample space $\mathbb{R}$ where each $x \in \mathbb{R}$ denotes the possible world where the compound takes value x. In order for your probability measure to be coherent, it must assign zero probability to each event $\{x\}$. This means that, where A_x is the proposition that the compound takes value x, $E(\neg A_x) = 1$—i.e., you must assign maximum credence to the claim that the compound does not take the value x. However, you may also believe that there is nothing in principle stopping the compound from having this value, so you fail to believe $\neg A_x$

---

[21] Dodd (2017) defends a credence-first reduction of belief against this objection by denying the link between one's doxastic profile and the bets they'd accept. When offered with a gamble on P, an agent who might've believed P would become aware of possibilities which previously they hadn't considered and in doing so would come to disbelieve P. This kind of view is called *sensitivism* in Greco(2015, p. 186).

despite assigning it maximum credence. Here, there are two things which can vindicate the credence-first reductionist against this objection. Firstly, as discussed in the outset of §1.3, one might be skeptical of agents who are able to entertain infinitely many possibilities in their attributions of likelihood to various events. Thus, an agent would not be 'forced' to assign zero probability to an event which they deemed possible (i.e., $\{x\}$), because they are not 'forced' to consider infinitely many possibilities in the first place. Even an experimental scientist who is interested in what real value some compound takes is required to group possibilities in accordance with the margin of error found in their scale, thermometer or ruler.

Regardless, even if this kind of epistemic state is possible, "full certainty" can no longer be identified with maximum probability. When the sample space is infinitely large an agent's epistemic state is to be modelled by a pair $\langle Pr, \boldsymbol{W} \rangle$ where Pr is their probability measure and $\boldsymbol{W}$ is a set of outcomes which the agent deems possible. Full certainty for these infinite contexts requires that a proposition be true across all worlds in $\boldsymbol{W}$. If you are performing an experiment to ascertain the real-valued height/weight/temperature of some compound then your epistemic state is the pair $\langle Pr, \mathbb{R} \rangle$, meaning that while you have credence 1 in $\neg A_x$ you are not fully certain of it since there is a world in $\mathbb{R}$, namely x, in which it is false that $\neg A_x$. So while you don't believe $\neg A_x$, you aren't fully certain of it either. Again, despite Spohn's remarks we can still interpret belief as maximum certainty, it's just that infinite sample spaces create an additional variable (above one's probability measure) which factors into the relevant scale of certainty.

Yet another problem for the advocated analysis is discussed by Roorda (1995, p. 8), Dodd (2017, p. 4614) and Jackson (2020, p. 2). There are some beliefs which we are more confident in than others. For instance, I am more confident that my friend was born in November than I am that he was born on the fourteenth of November. However, if I fully believe both of these claims—that is, my expectation of their truth value is 1—this difference in confidence can't be the result of different degrees of belief in the two propositions. Dodd's (2017, p. 4614) response to this problem involves distinguishing contexts where I am comparing the probability of the two propositions from contexts where I believe both. When I am comparing my confidence in both propositions, I become aware of and assign positive probability to possibilities which I might appropriately ignore when believing both – e.g. the possibility that I am misremembering the precise date but not the month of my friend's birth. This response seems sound in some cases, but there's an alternative explanation of the phenomena which allows for comparative confidence amongst full beliefs. Suppose Alice claims she is more confident that P than she is that Q, yet her probability measure is such that $E(P) = E(Q) = 1$. This difference in confidence might be due to the fact that, were her beliefs to change, Alice would more readily give up the belief that Q than she would her belief that P. Using Grove's (1988) model of belief revision theory presented for contexts in §1.2 above, this conviction of Alice's can be modelled by her adoption of a system of spheres where $S_{\neg Q} \subset S_{\neg P}$. That is, for Alice, the closest alternative state of belief where she doesn't believe Q is strictly contained within the closest alternative state of belief where she doesn't believe P. Similarly, I am more confident in the month of my friend's birth rather than the exact date because possibilities where I have remembered the month but forgotten the exact day are closer and more plentiful than possibilities where I have forgotten the month.

Say that an agent is at least as confident in P as they are in Q just in case $S_{\neg Q} \subseteq S_{\neg P}$ according to their system of spheres $\mathbb{S}$. This definition generates a "confidence order" on all propositions, but those who object to the credence-first reduction might suggest this ordering doesn't provide a rich enough model. They might suggest that, not only can we compare beliefs to see which is held with more conviction, we can also compare them to see the proportion between these convictions. For instance, Alice might have two beliefs but hold one of them with twice as much confidence as the other. This would require a model which can also countenance *degrees* of confidence amongst full beliefs. However, this model would have no natural interpretation inside Ramsey's (1990[1926]) view of credence where a belief's degree is the strength of one's disposition to act as if the belief's content is true. Once this strength is maxed out, at credence 1, there's no room for finer distinctions: fully believing a proposition requires simply acting as if it is true. If Alice fully believes two propositions, P and Q, then her degree of belief in the two is identical. Admitting degrees of confidence into our theory would require a separate interpretation of the degrees in question, perhaps something like "strength of one's disposition to abandon said belief". However, interpretations along these lines are awkward and comparatively unnatural. One cannot come to learn Alice's 'degree of confidence' in P or Q by finding her current betting behaviour, since abandoning a belief involves *changing* betting behaviour—i.e., Alice abandons her belief that P when she is no longer disposed to buy all bets which pay out under the condition that

P. Instead of attempting to provide such an interpretation, we stick with an ordinal understanding of comparative confidence, which has the following clear and natural interpretation: an agent is more confident in her belief that P than Q just in case, were the agent to stop believing $P \wedge Q$, she would stop believing Q before stopping her belief that P.

Another problem which looms large for this analysis of belief has to do with the *dynamics* of belief and probability. The problem is that if probability measures are updated via conditionalization, there is no series of updates to Pr which reduces A's probability when $Pr(A) = 1$. However, it also seems obvious that agents can come to stop believing things they previously took for granted. This is a concern aired by Jeffrey (1965, p. 171) and Spohn (2012, p. 45) uses it to bolster his argument against credence-first analyses of belief. So far, I have only attempted to present a static model of probabilistic presupposition. In the next section, we turn to the question of which update rules are appropriate for these new kinds of context, and in doing so we'll question the legitimacy of this problem for the credence-first analysis of belief. For now, it suffices to note that the force of this challenge derives its legitimacy from the assumption that conditionalization is the sole rule for updating probability functions but in the following sections we will consider updating rules where attributions of probability 1 can be rescinded.

1.4 Contextual Dynamics

The partial presupposition model developed in §1.3.2 is static. It places constraints on what an arbitrary context should look like, but doesn't explain how contexts should change once an assertion is accommodated. Now we seek to supplement this static model with some updating rules along the lines of PARTIAL MEET DISPLACEMENT and PARTIAL MEET ELIMINATION. The guiding idea for these new rules is that different types of qualified assertion operate on the context differently. This is how we can ensure that probabilistic and intensive qualification aren't automatically co-accommodated, which was flagged earlier (in §1.2.1) as a goal for our conversational kinematics. First, in §1.4.1, I consider updating rules for probabilistic assertion. There, I will survey some existing updating rules given in the literature on probabilistic assertion (Yalcin 2012, Rudin 2018). One issue with these updating rules is akin to the problem for Lewis' (1969) ACCOMMODATION rule discussed above: they present a way of "adding" information to a probabilistic context, but they don't prescribe any method of revising the context with incompatible information. After this, in §1.4.2, we consider updating rules for intensive qualified assertion given in Smith (2008), which are found to have the same shortfall. In §§1.4.3-4, I survey some different types of *imaging*, which are a family of update rules which do allow for the retraction of previously held probabilistic commitments. Then, in §1.4.5, I will use these imaging rules to present a system of contextual dynamics for (intensive and probabilistic) qualified assertions and partial accommodation.

All systems of probabilistic contextual dynamics considered in this section employ the notion of conditional probability. On a common (but not ubiquitous) interpretation, the conditional probability of P given Q (denoted $Pr(P|Q)$) is defined as the probability of their intersection divided by the probability of Q:

$$Pr(P|Q) = \frac{Pr(P \cap Q)}{Pr(Q)}, \text{ provided that } Pr(Q) \neq 0$$

Q in this formulation is called the *condition*. If $Pr_Y = Pr(X|Y)$ for all events X, then we call Pr_Y the *conditionalization* of Pr on Y. Gärdenfors (1988, p. 105) points out that the non-zero proviso immediately limits conditional probability's role in a Bayesian theory of belief revision. As argued in §1.3.3 above, agents believe a proposition just in case they assign probability 1 to a set of worlds in which it is true. Due to additivity and normalisation, this requires them to assign 0 probability to any set which includes worlds where that proposition is not true. Now, suppose the agent obtains information to the effect that this proposition *isn't* true. Because conditionalization is undefined when the condition has probability 0, they are unable to conditionalize on this incompatible information. In this way, conditionalization is analogous to Stalnaker's intersective rule for context sets: it merely provides a method for incorporating information which is already compatible with a given epistemic state.

Objectors to this analysis of conditional probability suggest that we should treat conditional probability as the primitive notion and derive unconditional probability from it.[22] These accounts prefer the following equation, inter-

---

[22] Hájek (2003, p. 315) attributes this kind of view to himself and an extensive list of forerunners in the theory of probability

preted as a constraint on conditional probabilities, which is well-defined even when the condition has probability 0:

$$Pr(P|Q)Pr(Q) = Pr(P \cap Q)$$

While this approach does allow for conditional probabilities to have some value even when the condition has probability 0, it offers little in the way of a procedure for the genuine revision of probability measures. Hájek (2003, p. 286) stipulates a few additional constraints on conditional probabilities (e.g., that the probability of a non-empty event P *given* P should always be 1), however these stipulated constraints are not rich enough to ensure that conditionalization will yield the least gratuitous redistribution of probabilities necessary to incorporate new information. Suppose Alice is fairly confident she will draw a red ball from the urn, but she also believes the urn contains no black balls, such that for Alice $Pr(R) = 0.9$ and $Pr(B) = 0$. Since $Pr(\cdot|B)$ is only minimally constrained according to this approach which takes conditional probabilities as primitive, it's possible that $Pr(R|B) = 0$, so long as $Pr(B|B) = 1$ and any other stipulative requirements are met. Of course, this might be a perfectly rational probability to adopt if Alice has evidence to the effect that the urn's containing black balls precludes its containing red ones. However, taking conditional probabilities to be minimally constrained permits Alice to perform these drastic conditionalizations without such evidence. Without a story about how one's body of evidence imposes constraints on conditional probabilities, the primitive conditional probabilities approach is underdeveloped.

1.4.1 Probabilistic Accommodation

Yalcin (2012) considers two different models of probabilistic contexts, and in doing so presents two different accommodation rules for probabilistic assertions. The first of these models suggests that probabilistic contexts are 'sharp' information states, which are pairs $\langle C, Pr \rangle$ where C is a subset of the possible worlds $\boldsymbol{W}$ and Pr is a probability measure defined over subsets of $\boldsymbol{W}$. Since Yalcin (2012, p. 12) restricts his attention to nondefective contexts, this section uses subscripts to indicate which sharp information state probability measures are contained in, as opposed to which agent they are held by. Unlike Yalcin, here I use italicised captial letters ambiguously to mean propositions or sets of worlds, as was done when handling Stalnaker's theory. This ambiguity will always be resolved by context—italic capital letters are flanked by set-theoretic operators when being used as an event and logical operators when being used as a proposition. Update functions $[\cdot]$ are functions of the form $\alpha \to (i \to i')$ where α is a proposition and i/i' are sharp information states. Here are two illustrative constraints which Yalcin (2012, p. 15) puts on these update functions, where $\triangle$ is an operator which means '...is probable':

$$i[P] = \langle C_i \cap P, Pr_i(\cdot|C_i \cap P)\rangle$$

$$i[\triangle P] = i \text{ iff } Pr_i(C_i \cap P) > .5, \text{ else } i[\triangle P] = \langle \varnothing, Pr_i(\cdot|\varnothing)\rangle$$

The first rule combines Stalnaker's intersective update rule with conditionalization. All worlds where P is not true are removed from the context set and then, the probability function in i is conditionalized on the new context set. The second rule is a 'test' which returns the information state i unchanged when P is probable according to i and returns the absurd information state otherwise. On Stalnaker's account, all context sets which are included in the set of worlds where P is true presuppose that P. On Yalcin's probabilistic extension, an information state is said to *accept* (aka presuppose) some proposition P iff $i[P] = i$—that is, if updating with P fails to change the information state.

One issue with sharp information states as a model of probabilistic contexts is the requirement that speakers have a common background which is rich enough to determine specific probabilities for each event. Since this kind of unanimous ultra-specific probabilistic agreement is rare, Yalcin's second model of probabilistic contexts uses *blunt* information states, which are defined as sets of sharp information states. Update functions for blunt information states are determined using sharp update functions. Where I is a blunt information state,

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

That is, to update a blunt information state with α, you simply remove all of its sharp information states which fail to accept α (if there are any). This second model has the added advantage of allowing probabilistic qualifications to be non-trivially informative. Yalcin (2012, pp. 18-19) points out that when a sharp context is successfully updated with a proposition like $\triangle P$, no change in context takes place. However, when a blunt context set is successfully updated with $\triangle P$, some of its sharp context sets may be removed—namely, those which don't attribute at least 0.5 probability to P.

Rudin (2018) develops Yalcin's context probabilism further by posing an additional desideratum for any model of probabilistic accommodation, called 'slippery intermediate credences'—which holds that contexts should be flexible when it comes to successive changes of a proposition's intermediate probability. For example, while watching a game of soccer, it seems permissible to revise the presupposed probability of a particular team winning as the game progresses. A spectator might say 'They'll probably win' when the team is up by two goals, 'They'll probably lose' when down by a goal in the last ten minutes, and then 'They'll probably win' again when they score two more in quick succession. Rudin suggests that contexts should be able to accommodate all three of the spectator's assertions in order of appearance. More generally, Rudin (2018, p. 372) suggests that "expressions of some range of uncertainty about ϕ must be followable without contradiction by expressions of some other range of uncertainty about ϕ—potentially a non-overlapping range." Rephrasing in terms of Yalcin's update functions, where $\%_{[x,y]}$ is an operator which means "...has a probability in the interval $[x, y]$", this requirement says that for all non-empty P, $i[\%_{[x,y]}P][\%_{[w,z]}P] \neq \langle \varnothing, \mathrm{Pr}_i(\cdot|\varnothing)\rangle$ (unless $[x, y]$ is $[1, 1]$ or $[0, 0]$).

Yalcin's sharp and blunt information states both fail to capture this notion of a slippery intermediate credence. To see why in the blunt case, consider an information state I which only contains the sharp information state i. Suppose $Pr_i(X) = 0.8$, then $I[\%_{[0.8,0.9]}P][\%_{[0.2,0.5]}P] = \varnothing$, since the first update returns $I = \{i\}$ but the second update removes i. For the sharp case, consider the same sequence of updates on the sharp information state i—this state will 'pass' the first test and 'fail' the second, yielding the absurd sharp information state $\langle \varnothing, \mathrm{Pr}_i(\cdot|\varnothing)\rangle$.

Rudin attempts to solve this inflexibility by replacing Yalcin's updates for blunt states with point-wise conditionalization. Rudin (2018, p. 388) takes information states I to be a set of probability functions defined over some finite set of worlds W. The notion of a context set for each probability measure can be recovered by taking C_{Pr} (what Rudin calls Pr's WORLD-IMAGE)[23]to be the intersection of all events which have probability 1—i.e.,

---

[23] This is equivalent to the recovery of context sets used in §1.3.1 above. Recall that

$$C_{Pr} = \bigcap \{X \subseteq \boldsymbol{W} : Pr(X) = 1\}$$

.

Given a set of probability measures I, one incorporates the information that P has a probability in the range $[x, y]$ by replacing I with I', which is the set of all probability measures Pr satisfying the following two requirements:

(i) $Pr(P) \in [x, y]$

(ii) There is some $Pr_i \in I$ and some *maximal* event $E \subseteq \boldsymbol{W}$ such that Pr is the conditionalization of Pr_i on E. Say that E is maximal just in case there is no event E' such that $E \subset E'$ and $Pr_i(P|E') \in [x, y]$.

Let us say that one of Rudin's information states I *accepts* an attribution of the probability interval $[x, y]$ to a proposition P just in case $I[\%_{[x,y]}P] = I$. Acceptance of a proposition outright is a special case where the interval of attribution is [1,1]. The maximality constraint guarantees that the new information state I' doesn't end up including unreasonably informative probability measures. Without it, updating a state I with the claim that P has probability 1 would lead to an I' which includes probability measures conditionalized on the event $P \cap Q$, for some arbitrary $Q \subset P$—e.g., updating with the information that it will certainly rain tomorrow will import probability measures which assign probability 1 to the event that it rains tomorrow *and* exactly two-thousand-and-one polar bears will go swimming. Without the maximality constraint, contexts would constantly shift from certainty about Q's probability to uncertainty about Q's probability even when the updates have nothing to do with Q.

The problem with Rudin's (2018) proposal relates to an issue which Jeffrey (1965) flagged for conditionalization in general: sometimes, the space of propositions is too impoverished for conditionalization to provide the desired flexibility. Consider a simple space of possibilities $\boldsymbol{W}$ which only contains three worlds: r, g and b, where each of these worlds is exactly the same except for an RGB diode which briefly emits red, green or blue light. Suppose a group of agents agree it is equally likely that any of the

---

there $\boldsymbol{C}_{Pr}$ was defined as $\{w \in \boldsymbol{W} : Pr(\{w\}) \neq 0\}$. Since $\boldsymbol{W}$ is finite, the smallest of its subsets which is assigned probability 1 is simply the set of all and only those worlds which have positive probability.

three colour flashes will occur, which can be represented by their adoption of an information state I which only contains the probability measure Pr s.t. $Pr(\{r\}) = Pr(\{g\}) = Pr(\{b\}) = \frac{1}{3}$. The diode flashes but due to poor visibility conditions (perhaps they are viewing it in broad daylight), the agents come to have 0.8 credence in its having flashed blue and 0.2 credence in its having flashed green. Using Rudin's update rule, there's no way for such a shift in information states to take place, simply because there's no event which can be conditionalized upon to produce the correct probability distribution. This leads to a violation of slippery intermediate credences. If G expresses the proposition that the light flashed green, the counterexample is $I[\%_{[0.\dot{3},0.\dot{3}]}G][\%_{[0.2,0.2]}G]$, which is empty relative to this set-up, since the sole probability function in I is removed by the second update.

Jeffrey (1965, p. 166) frames his objection to conditionalization in terms of linguistic resources, saying that it might be the case that an agent's subjective probabilities change because of some observation which is inexpressible in their language. Here, the problem is in terms of the limited class of events under consideration (i.e., subsets of $\boldsymbol{W}$), but the upshot is the same either way: instead of mediating probabilistic update through conditionalization on some event, probabilistic updates should directly shift probabilities in accordance with the update's content. Jeffrey (1965, p. 169) proposes this can be done with a generalisation of conditionalization which has since come to be known as Jeffrey conditioning. Suppose an observation (or assertion or piece of evidence) causes an agent to abandon their probability measure Pr in favour of a new measure Pr' which assigns probability x to the event A. So long as A's initial probability is neither 0 nor 1, the remaining details of this new probability measure are described as follows,

$$Pr_{(x,A)}(\cdot) = Pr(\cdot|A)x + Pr(\cdot|\overline{A})(1-x)$$

Note that regular conditionalization is the special case where $x = 1$. Of course, observations can cause an agent to directly revise their probabilities for more than one event—take the observation of the RGB diode from the previous discussion, for example. So in general, given a partition $\mathcal{X}$ on $\boldsymbol{W}$ with n cells and respective input probabilities x_i for each cell,

$$Pr_{(x_i,\mathcal{X})}(\cdot) = Pr(\cdot|A_1)x_1 + ... + Pr(\cdot|A_n)x_n$$

Replacing conditionalization with Jeffrey conditioning yields an updating rule which *does* satisfy the condition of "slippery intermediate credences" but it compounds another problem stated at the outset of this section. Instead of being undefined when the condition has probability 0, Jeffrey conditioning is undefined when *any* of the multiple conditions used has probability 0. This entrenches a problem of "grippy extreme credences", such that whenever an event has probability 0 or 1, it is impossible to revise that event's probability with any sequence of Jeffrey conditionings. This is why Jeffrey (1965, p. 171) ultimately imposes a constraint on credences called regularity:

REGULARITY $Pr(A) \in (0,1)$, unless A is empty or $\boldsymbol{W}$.[24]

However, given that the model of presupposition developed in §1.3.1 allows for full belief in propositions which aren't logical truths, we will not follow Jeffrey in imposing regularity. Instead, we seek a more general form of updating which has the potential to change the probabilities of events which are assigned probability 1 or 0.

Interestingly, not only is Rudin (2018, p. 372) unconcerned about the grippy-ness of extreme credences, but goes so far as to enshrine this feature as a desideratum for models of probabilistic contexts, saying that "the proposition φ must *not* be followable without contradiction by the proposition $\neg\varphi$". That is, information states which result from updates the form $I[\varphi]...[\neg\varphi]$ must always be the empty. However, while it's clear that contradictory assertions occasionally disrupt the common ground, it's not clear that contexts must proceed in this fashion, or that they're lawless when they don't. To the contrary, in sections §1.1.1 and §1.1.2 some laws were proposed which do govern the accommodation of assertions which are incompatible with the context of their utterance: PARTIAL MEET DISPLACEMENT and PARTIAL MEET ELIMINATION. Below, in §1.4.5, we give similar laws for assertions which are incompatible with probabilistic contexts of utterance.

1.4.2 Intensive Accommodation

Smith (2008, pp. 246—248) presents a model of context which tracks changes in relation to a proposition's presupposed degree of truth. This model consists in a set of probability measures (one for each conversationalist) where

---

[24] This regularity constraint is amenable to Jeffrey's (1970, p. 161) proposal for replacing all attributions of "full belief" with attributions of sufficiently high partial belief, discussed in §1.3.3 above.

context sets can be recovered from these probability measures using Rudin's notion of a WORLD-IMAGE discussed above.[25] Against this background, a speaker presupposes that P is true to degree n if and only if their probability measure assigns 1 to a set of worlds where P is true to degree n. A non-defective context is one where every speaker's probability measure has the same world-image, meaning that if any of them take a proposition to be true to some degree, they all take it to be true to that degree. Interestingly, this type of non-defective context still allows for epistemic disagreement, since speakers might assign different probabilities to the subsets of this shared world-image (Smith 2008, p. 247).

Recall from §1.2.1 that speakers are able to qualify assertions with hedging phrases, attributions of degrees of truth, hesitance and other behaviours which are characteristic of middling confidence. We take it that for every degree of truth n, there is a corresponding degree of strength or confidence with which an assertion can be made. When a speaker makes an n-strength assertion that P, this can be understood as a proposal to change the context set by removing all worlds where P isn't true to degree n. Smith's (2008, p. 248) proposed accommodation rule for this kind of intensive qualification is conditionalization. Speakers accommodate an n-strength assertion that P by conditionalizing their probability measure on the set of worlds where P is true to degree n. That is, where P_n denotes the set of worlds where P is true to degree n, each speaker will replace their Pr with Pr', obtained as follows:

$$Pr'(X) = \frac{Pr(X \cap P_n)}{Pr(P_n)}$$

Since this rule is undefined when the probability of P_n is zero, it does not allow speakers to *change* the degree of truth which P is presupposed to have. Instead, it only permits speakers to go from being undecided about P's degree of truth to having a presupposition on the matter. Smith's (2008, p. 248) justification for this restriction on conditionalization claims that the issues vanish when considering the illegitimacy of making an n-strength assertion that P when it is already presupposed that P isn't true to degree n. However, this kind of conversational move is just a many-valued generalisation of the

---

[25] Smith (2008, p. 247) actually proposes to take a context set relative to some measure Pr as the set C such that $Pr(C) = 1$. However, this method of extraction would make it such that every superset of C is also a context set associated with Pr. Assuming that there is a unique context set which can be associated which each probability measure, we choose to adopt Rudin's method of extraction.

same kind of move used in conversations like COIN GAME, so it is legitimate in at least some cases. This is yet another reason why, in the next section, we introduce an alternative method of revising probabilities which is still well-defined when the condition has probability zero.

1.4.3 Imaging Generally

Imaging is a procedure for revising probability measures which is well-defined even when the event being incorporated has probability 0. It was first given in Lewis (1976, p. 310) as a framework for attributing probabilities to Stalnaker's counterfactual conditionals. Following the contemporary presentation in Fusco (2023), a probability measure Pr^X is said to be the *image* of Pr on some event X when,

$$Pr^X(\{w\}) = \begin{cases} 0, & \text{if } \nu(\langle X, w\rangle) \neq 1 \\ Pr(\{w\}) + \displaystyle\sum_{w' \in \overline{X} \mid \sigma(w', X) = w} Pr(\{w'\}), & \text{if } \nu(\langle X, w\rangle) = 1 \end{cases}$$

Here, $\sigma(w, X)$ is a function $\boldsymbol{W} \times \wp(\boldsymbol{W}) \to \boldsymbol{W}$ which intuitively gives the 'closest' world to w where the proposition X is true (to degree 1). As for an intuitive gloss on imaging as a whole, the idea is that imaging on some event X involves shifting probability assigned to any world where X isn't true over to the nearest possible world where X is true.

One might have the same objections to this function σ that people have to Stalnaker's (1968) notion of a "selection function" used in the evaluation of counterfactual conditionals. Is it true, in all cases, that there will be a *unique* closest world to w in which X is true? Not if two or more worlds are tied for similarity in the relevant sense. Suppose X is the claim that Bizet and Verdi are compatriots, borrowing an example from Lewis (1973, p. 80). X-worlds where both composers are French are just as similar to the actual world as X-worlds where both composers are Italian, since they require identical levels of departure from reality viz. the revision of one person's nationality. In cases like these, it makes sense that probability assigned to a world where X is not true may be distributed amongst a *set* of closest worlds where X is true when imaging on X. This is achieved by changing the function σ to a function whose range includes sets of possible worlds, i.e. $\sigma(w, X) : \boldsymbol{W} \times \wp(\boldsymbol{W}) \to \wp(\boldsymbol{W})$. Following Fusco (2023, p. 448), we place the following two requirements on this set-valued function σ:

(Success) $\sigma(w, X) \subseteq X$

(Strong Centring) If $w \in X$, then $\sigma(w, X) = \{w\}$

There is an obvious similarity between this function σ so-constrained and the systems of spheres given in §1.1.2 above. $\sigma(w, X)$ can be interpreted as a function which returns $S_X \cap X$ in a system of spheres centred on $\{w\}$. That is, it returns the most viable alternative belief state to $\{w\}$ relative to the constraint that X holds.

Now that worlds are related to *sets of* closest X-worlds, we introduce a *transfer function* $T_{w,X} : \{v \in \sigma(w, X)\} \to [0, 1]$ which describes how the probability from a non-X world is to be divided amongst a set of closest X-worlds. Since all probability must be vacated from the non-X worlds, we require that $\sum_{v \in \sigma(w,X)} T_{w,X}(v) = 1$, ensuring that the full proportion of w's probability is sent somewhere X-friendly, even if this destination ends up being w itself. The addition of a transfer function lends itself to a new revision process called *general imaging*, defined as follows:

$$Pr^X(\{w\}) = \begin{cases} 0, & \text{if } \nu(\langle X, w\rangle) \neq 1 \\ Pr(\{w\}) + \displaystyle\sum_{w' \in \overline{X} | w \in \sigma(w', X)} Pr(\{w'\}) \cdot T_{w',X}(w) & \text{if } \nu(\langle X, w\rangle) = 1 \end{cases} \tag{1.1}$$

Interestingly, Fusco (2023, p. 455) shows that conditionalization can be construed as a special case of general imaging. To get conditionalization out of imaging, one simply has to put the following two restrictions on σ and $T_{w,X}$:

$\sigma(w, X) = \{v \in X : Pr(\{v\}) > 0\}$, for all $w \in \overline{X}$

$T_{w,X}(v) = \frac{Pr(\{v\})}{Pr(X)}$

This way, events which are subsets of X retain the proportion of probability they have relative to X's, just as they do under the ratio definition of conditionalization. Also note that it inherits the problem of being undefined when the probability of the condition is zero, here because it leads to undefined values in the transfer function. Fusco (2023, p. 461) does point out that this reduction of conditionalization requires a relativisation of σ to a particular probability measure, which is problematic if worldly similarity is understood 'metaphysically', in the sense that whether two possible worlds are similar depends solely on facts about the worlds themselves. Metaphysical similarity is important for things like evaluating counterfactual conditionals, where we're interested in whether such propositions are true. However, for our purpose, which is to provide laws for how agents update their beliefs during the course of a conversation, this metaphysical notion of similarity is clearly inadequate. Two possibilities may be metaphysically similar without being similar in the sense which is relevant for conversational dynamics. We might be engaged in a conversation where it is presupposed that the solar system must have an odd number of planets. Relative to this background, possible worlds where there are 11 planets in our solar system may be considered more similar to the actual world than possible worlds where there are 10, despite the later possibilities being more similar to actuality overall.

Now, recall that the guiding principle of belief revision is the concept of *informational economy*, which holds that a revised belief state ought to be as similar to the original belief state as consistency allows (Gärdenfors 1988, p. 49; Fermé and Hansson 2018, p. 13). As Lewis (1976, p. 311) observed, conditionalization and (unconstrained) general imaging both achieve informational economy in different senses. Conditionalization is conservative insofar as it preserves the proportional distribution of probabilities assigned to subsets of X, whereas imaging is conservative insofar as all probabilistic mass is transported a minimal distance indicated by σ. However, since conditionalization just is a special case of general imaging, using imaging to conditionalize adheres to this principle of informational economy in *both* senses. This leads us to a preliminary rule for probabilistic revision:

Conditional Revision If an agent's probability measure is Pr and furthermore $Pr(X) > 0$, then they should revise Pr with the information that X by conditionalizing Pr on X—that is, by replacing Pr with Pr_X.

Assuming that general imaging is the fundamental rule for probabilistic revision, conditional revision constitutes a two-part requirement for those wishing to form the belief that X. First, it requires that agents who assign non-zero probability to X must adopt a function σ which represents all (subjective) live possibilities in X as equally similar to any world in X's complement. Secondly, it requires that they adopt a transfer function T such that $T_{w,X}(v) = {}^{Pr(\{v\})}/_{Pr(X)}$ for all $v, w \in \boldsymbol{W}$. Together, these requirements guarantee that informational economy is achieved insofar as subsets of X retain their proportional probability. As for the first requirement, however, one might suggest that it's possible for strict subsets of $\{v \in X : Pr(\{v\}) > 0\}$ to be maximally similar to some world outside of X. If so, then CONDITIONAL REVISION would violate informational economy in the imagistic sense, leading agents to shift probability further than necessary for the accommodation of X. However, prioritising informational economy in the imagistic sense comes at a high cost, since it leads to a gratuitous distortion of those probabilities assigned to subsets of X.

Take this illustrative example. Suppose someone offers you a covered box which could possibly contain (i) nothing, (ii) a dollar, (iii) 10 dollars or (iv) 100 dollars. They tell you that they have flipped a fair coin to decide whether or not the box will be empty, and they also tell you that should the box contain money, the amount was decided by spinning a wheel which has equal area designated for each of the options (ii), (iii) and (iv). This leads you to a probability measure Pr which is such that $Pr(E) = 0.5$, where E is the event where the box is empty and $Pr(\{x\}) = 0.1\dot{6}$ where x is the world where the box contains \$$x$. Since the world where the box contains one dollar is more similar to the world where the box is empty than any other (by smallest deviation in terms of your wealth), you may adopt a function σ such that $\sigma(0, \overline{E}) = \{1\}$.[26] By the requirement that, for each $v \in \sigma(w, X)$, the values of $T_{w,X}(v)$ must sum to 1, this will only be possible if $T_{0,\overline{E}}(1) = 1$—i.e., if the world 0 wills all of its probability to the world where the box contains \$1. Now, if you are told that the box is not empty, your new probability measure after imaging on $\overline{E}$ will assign $0.\dot{6}$ probability to the outcome where the box has one dollar and $0.1\dot{6}$ probability to each of the two remaining outcomes. That is, you will now take it as probable that the box contains one dollar *despite* not being given any evidence to that effect. Indeed, because the

---

[26] It is par for the course in probability and decision theory to give toy examples where monetary value is used as a placeholder for some other value of significance. If you scoff at the idea that individual wealth could determine worldly similarity, you can either imagine the agent as someone who is concerned about nothing other than money or else imagine a suitable series of progressively life-altering box contents.

box-filler has informed you of the fair spinning wheel used to determine the box's contents you have contrary evidence to the effect that (ii), (iii) and (iv) should occur with equal probability. To avoid cases like this, where probabilistic evidence is destroyed for the sake of minimal revision in the imagistic sense, a rule like CONDITIONAL REVISION requires that agents conditionalize wherever possible. This way, the agent's revision can be seen as satisfying informational economy in both senses, since their conditionalization can be reconstructed as a special case of general imaging where the value of $\sigma(0, \overline{E})$ is the set $\{1, 10, 100\}$.

Of course, this doesn't mean that agents should *always* incorporate new information via conditionalization. If the new information X is assigned zero probability by Pr, then general imaging is the relevant rule for revision:

IMAGING REVISION If an agent has the probability measure Pr and $Pr(X) = 0$, then they should revise Pr with the information that X by general imaging Pr on X—that is, by replacing Pr with Pr^X.

This rule is favourable, in part, simply because general imaging *is* well-defined when the probability of X is zero whereas conditionalization isn't. Agents are better off settling for one sort of informational economy in cases where they can't adhere to both. However, IMAGING REVISION isn't merely a compromise. When X is assigned probability 0, then raising its probability to 1 already involves abandoning all probabilistic structure distributed across subsets of $\overline{X}$, since $Pr(\overline{X})$—and the probability of any subset of $\overline{X}$—is in turn reduced to zero. So, any distortion of probability assigned to subsets of X happens in the context of a shift which already distorts the probability measure as a whole.

As an added bonus, IMAGING REVISION also buttresses the analysis of outright belief as full *degree of* belief. Recall that Jeffrey (1965, p. 171) and Spohn (2012, p. 45) rejected this analysis on dynamic grounds. If an agent assigns probability 1 to some event E, i.e. $Pr(E) = 1$, then there is no sequence $Pr_X, Pr_{XY}, Pr_{XYZ}, \ldots$ of conditionalizations which reduces the probability of E thereafter. They contest that if belief is full degree of belief and conditionalization is the only medium of belief change, then we can't stop believing propositions which we have come to believe. Clearly, we can stop believing things, so this analysis of belief must be faulty. However, with IMAGING REVISION, we are in a position instead to reject the claim that conditionalization mediates all belief change. Agents can stop believing propositions which they believed earlier because, with imaging, they can revise their probability measure in a way that lowers the probability of E even when $Pr(E) = 1$. This will be the case whenever an agent wishes to incorporate the information that some event in $\overline{E}$ obtains, which will kick IMAGING REVISION into gear since $Pr(X) = 0$ for all $X \subseteq \overline{E}$.

1.4.4 Imaging Partially

It's easy to see how CONDITIONAL REVISION and IMAGING REVISION in tandem could feature in accommodation rules for unqualified assertions in probabilistic contexts. One would simply need to replace the idea that speakers accommodate through conditionalization with the principle that they do so through imaging *or* conditionalization depending on the presupposed probability of the assertion's content. However, doing so now would leave a gap in the accommodation rules, since it would fail to provide guidance when a qualified assertion is made or in cases where accommodation is only partially granted. So, before we integrate imaging into our theory of contextual dynamics, we will first seek an imaging-based analogue of Jeffrey conditioning which allows for probability measures to be revised with information which is itself more-or-less certain.

There is a growing debate on this very topic of partial imaging which stems from Eva and Hartmann's (2021) attempt to give a Bayesian model of various kinds of supposition.[27] There, they posit four categories of supposition: (i) full indicative supposition, (ii) full subjunctive supposition, (iii) partial indicative supposition and (iv) partial subjunctive supposition. Whether supposition is full or partial depends on whether the proposition being supposed is taken as certain or taken as having some intermediate probability. As for whether supposition is indicative or subjunctive, Eva and Hartmann (2021) distinguish using the following pair of similes:

(Indicative) When an agent supposes the truth of P in the indicative mood, they change their beliefs as if they were to learn the truth of P.

---

[27] See Eva and Hartmann (2021), Nielsen (2022), Fusco (2023) and Zhang (2024).

(Subjunctive) When an agent supposes the truth of P in the subjunctive mood, they change their beliefs as if they were to learn that P were made true by some 'local miracle' or 'divine intervention', with all else held as fixed as possible.[28]

Conditionalization models the belief change required by (i), imaging models the belief change required by (ii) and Jeffrey conditioning models the belief change required by (iii) and so the open question is, 'What models the belief change required by (iv)?' Eva and Hartmann (2021) suggest the most intuitive answer may be the least promising. They consider an extension of imaging which is analogous to Jeffrey conditioning, which is called Jeffrey imaging:

Jeffrey Imaging Given a partition $\mathcal{X}$ of $\boldsymbol{W}$ with n cells labelled $X_1, ..., X_n$ and an input probability $q(X_i)$ for each of the cells, the Jeffrey image of Pr on these input probabilities, denoted $Pr^{(q,\mathcal{X})}$, is defined as follows:

$$Pr^{(q,\mathcal{X})}(\cdot) = Pr^{X_1}(\cdot)q(X_1) + ... + Pr^{X_n}(\cdot)q(X_n)$$

That is, to find the probability of an event E after jeffrey imaging, one images on each of the cells in $\mathcal{X}$ and weighs the probability of E on these images by a proportion determined by the input values q. Eva and Hartmann's (2021) concern with Jeffrey imaging is that it fails to obey a highly intuitive rule for belief revision:

Monotonicity Condition Whenever an agent raises the probability of X and lowers the probability of $\overline{X}$ in changing their probability measure from Pr to Pr', they must ensure that $Pr(Y) \leq Pr'(Y)$ for all $X \subseteq Y$.

With the idea behind this condition being that raising the probability of an event X shouldn't decrease the probability of an event Y which is guaranteed

---

[28] These two moods also outline a distinction in approaches to decision theory, as shown by Joyce (1999, p. 181). Given a general framework of decision theory where agents ought to select an available action which, *supposing* they have performed it, maximises expected utility, Evidential Decision Theory (EDT) interprets this supposition in the indicative mood, whereas Causal Decision Theory (CDT) interprets this supposition in the subjunctive mood.

by the occurrence of X. For example, one shouldn't increase the subjective likelihood of it being at least 30°C tomorrow while decreasing the subjective likelihood of it being at least 25°C tomorrow. JEFFREY IMAGING is a method of belief revision which does violate this intuitive rule. To see why, consider a case where $\boldsymbol{W} = \{r, g, b\}$ represents a space of alternative colour flashes from an RGB diode. At first, Alice thinks the flashes are equally likely, such that each world is assigned probability $\frac{1}{3}$. Then, she (subjunctively) supposes that a local miracle has subtly yet indubitably changed the probabilistic profile of these events so that the likeliness of $\{g\}$ slightly increases to 0.4, with all else held as equal as possible. If she achieves this supposition through JEFFREY IMAGING then her adoptive probability measure will be $Pr^{(q,\mathcal{X})}$ where q is such that $q(\{g\}) = 0.4$ and $q(\{r, b\}) = 0.6$ and $\mathcal{X} = \{\{g\}, \{r,b\}\}$. If, for some reason, Alice reckons r is more similar to g than b (perhaps due to mild red-green colourblindness) and $\sigma(g, \{r, b\}) = \{r\}$, then her adoption of this new probability measure will violate the MONOTONICITY CONDITION, since $Pr(\{g\}) = \frac{1}{3} < Pr^{(q,\mathcal{X})}(\{g\}) = 0.4$ but $Pr^{(q,\mathcal{X})}(\{g, b\}) = 0.6 < Pr(\{g, b\}) = \frac{2}{3}$. So, Alice would increase the probability of the light flashing green but decrease the probability of it flashing green *or* blue, relative to her initial probability measure Pr.

A recent paper from Zhang (2024) argues that while JEFFREY IMAGING does violate MONOTONICITY CONDITION, it only does so while obeying another closely related rule for belief revision. This alternative kind of monotonicity suggests that, instead of using the initial probability measure Pr as a benchmark, one should use what is called the *status quo probability measure* which, given a (indicative or subjunctive) change of Pr to some new probability measure Pr' using input values q and partition $\mathcal{X}$, is either $Pr^{(p,\mathcal{X})}$ or $Pr_{(p,\mathcal{X})}$ (depending on the mood used in deriving Pr'), where $p(X_i) = Pr(X_i)$ for all $X_i \in \mathcal{X}$. In other words, the status quo probability measure is that probability measure which you would adopt upon supposing that your current measure Pr is correct concerning the probability of all events in $\mathcal{X}$. To aid with the presentation of Zhang's surrogate for MONOTONICITY CONDITION, we will use the notation $Pr(\cdot||q, \mathcal{X})$ to mean an arbitrary revision or update of $Pr(\cdot)$ using input probabilities q and partition $\mathcal{X}$. Here, for comparison, is monotonicity *and* status quo monotonicity in Zhang's terms:[29]

---

[29] This alternate presentation of the monotonicity condition is also required to account for revisions where $\mathcal{X} \neq \{X, \overline{X}\}$. As Fusco (2023, pp. 453—454) observes, monotonicity

MONOTONICITY CONDITION (MC) Let $\mathcal{G} \subseteq \mathcal{X}$ be such that $X \in \mathcal{G}$ iff $q(X) \geq Pr(X)$. Furthermore, let G denote the event $\bigcup \mathcal{G}$. For any Y such that $G \subseteq Y$, it follows that $Pr(Y) \leq Pr(Y||q, \mathcal{X})$.

STATUS QUO MONOTONICITY (SMC) Let $\mathcal{G} \subseteq \mathcal{X}$ be such that $X \in \mathcal{G}$ iff $q(X) \geq Pr(X)$. Furthermore, let G denote the event $\bigcup \mathcal{G}$. For any Y such that $G \subseteq Y$, it follows that $Pr(Y||p, \mathcal{X}) \leq Pr(Y||q, \mathcal{X})$.

As Zhang (2024, p. 585) notes, (SMC) reduces to (MC) whenever $\cdot||\cdot$ is a *conservative* update rule, meaning that $Pr(\cdot||p, \mathcal{X}) = Pr(\cdot)$ for all partitions $\mathcal{X}$. That is, if revising your probability measure with the constraints $p, \mathcal{X}$ fails to change it, then the benchmark for (SMC) is just the initial probability as it is in (MC). Jeffrey conditioning obeys (SMC) and since it is conservative, it also obeys (MC). On the other hand, since it is not conservative, Jeffrey imaging only obeys (SMC). Using the previous example of monotonicity violation, one can see that while Alice does decrease $Pr^{(q,\mathcal{X})}(\{g, b\}) = 0.6$ relative to $Pr(\{g, b\}) = \frac{2}{3}$, she increases $Pr^{(q,\mathcal{X})}(\{g, b\})$ relative to $Pr^{(p,\mathcal{X})}(\{g, b\}) = 0.\dot{5}$. So, (SMC) appears to validate Alice's use of Jeffrey imaging by showing how her increase of g's probability also increases $\{g, b\}$'s probability relative to a close derivative of Pr which has only been transformed by the assumption that Pr is correct concerning the values of events in $\mathcal{X}$.

While this surrogate rule does capture some of the intuitive force behind MONOTONICITY CONDITION, it also invites some difficult-to-answer questions. The first question, which Zhang (2024, p. 586) acknowledges, asks whether it is reasonable for any rule of belief revision to be non-conservative. Surely, the smallest change to Pr under the supposition that Pr is correct concerning the probability of events in $\mathcal{X}$ is simply to retain Pr! Zhang (2024, p. 587) responds to this worry by emphasising that this might not be the case for minimal change in an imagistic sense, which prioritises minimal revision in matters of fact. In essence, the idea here is that you can picture (full) imaging as a process by which one goes through every $\neg X$-world that has been assigned positive probability and changes the smallest number of

---

is concerned with the largest event which raises in probability overall, as opposed to any arbitrary event whose probability is raised. This is why the constraints in (MC) and (SMC) are imposed on supersets of G.

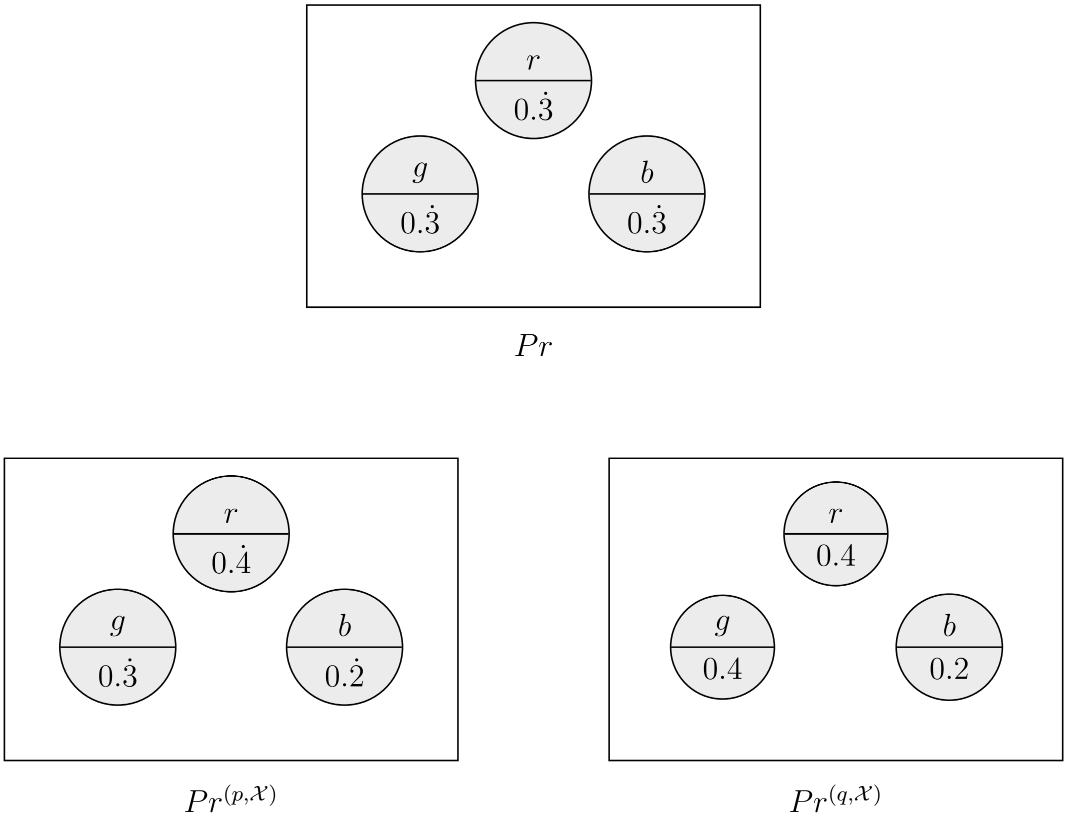

Figure 1.4: Alice's probability measures when considering the RGB diode and using Jeffrey imaging, where a number below the line at each world w indicates the probability of $\{w\}$.

facts required to make it an X-world. Redistribution of probability amongst worlds which are already accepted as possible is still minimal in this sense since it doesn't require any revisions of fact amongst worlds in $\boldsymbol{C}_{Pr}$—i.e., the set of worlds deemed possible by Pr. However, under Lewis' gloss on imagistic minimality recounted above—where imaging seeks the smallest migration of probabilistic mass required to give X probability 1—changing the probability of an event relative to the constraint that it has that probability *is* gratuitous. Furthermore, full imaging is conservative in the regular sense since Pr^X is Pr when $Pr(X) = 1$, so it's not unreasonable to expect a similar property from partial imaging.[30]

Another question is why the derived probability $Pr^{(p,\mathcal{X})}$ should be the yardstick for monotonicty, as opposed to the initial probability Pr. Even the above statement of (SMC) is concerned with the relationship between the initial probability of an event X and its revised probability. Why shouldn't (SMC) also be concerned with this same relationship between the initial probability of some $Y \supseteq X$ and its revised probability?

These concerns take us (in reverse chronological order) to Fusco's (2023) proposal for a version of partial imaging which, unlike Jeffrey imaging, obeys MONOTONICITY CONDITION. Instead of taking a sequence of images (one for each member of the partition $\mathcal{X}$) and combining them with respective weights according to q, Fusco imaging works by "skimming" the amount of probability required to raise $Pr(X_i)$ to $q(X_i)$ uniformly off of worlds in $\overline{X_i}$ which lose overall probability. This way, images on $X_j \subseteq \overline{X_i}$ aren't used in the process, meaning that these subsets of $\overline{X_i}$ don't get assigned disproportionately higher probabilities which is what ultimately leads to violations of monotonicity. Fusco imaging (FI) begins with the same ingredients as partial imaging: an initial probability measure Pr, a partition $\mathcal{X}$ on $\boldsymbol{W}$ and a series of new input probabilities $q(X_i)$ for each $X_i \in \mathcal{X}$. Next, the cells in $\mathcal{X}$ are rearranged into the form $\{L_1, ..., L_k, G_1, ..., G_m\}$ where $q(L_i) < Pr(L_i)$ (so-called "loser" cells) and $Pr(G_j) \leq q(G_j)$ (so-called "gainer" cells). Furthermore, let L and G denote $\bigcup_i L_i$ and $\bigcup_i G_i$ respectively. (FI) is defined as follows given this set-up:

---

[30] Note that there is a different notion of conservativity employed in Gärdenfors (1982, p. 754) which is *not* satisfied by imaging. Here, I am stating that full imaging is conservative in Zhang's (2024, p. 585) sense—whenever $Pr(X) = 1$, Pr is the image of Pr on the constraint that event X has probability 1.

$$Pr^{(q,\mathcal{X})}(\{w\}) = \begin{cases} Pr(\{w\}) \cdot \dfrac{q(L_i)}{Pr(L_i)} & \text{if } w \in L_i \\ Pr(\{w\}) + \lambda_i \left[\displaystyle\sum_{v \in L | w \in \sigma(v,G_i)} (1 - \dfrac{q(L)}{Pr(L)}) \cdot P(v) \cdot T_{v,G_i}(w) \right] & \text{if } w \in G_i \end{cases} \tag{FI}$$

where the value λ_i is for each G_i the proportion of probability taken from L which is needed to raise $Pr(G_i)$ to $q(G_i)$, defined as follows:

$$\lambda_i := \frac{q(G_i) - Pr(G_i)}{Pr(L) - q(L)}$$

As shown in Fusco (2023, p. 454), (FI) obeys (MC) and this can be seen in the case of Alice and the RGB diode. Were she to use (FI) instead of Jeffrey imaging, then Alice's new probability measure $Pr^{(q,\mathcal{X})}$ would increase the probability of $\{g, b\}$ from $\frac{2}{3}$ to 0.7, where $\{g, b\}$ is a superset of the "gainer event" $G = \{g\}$. Figure 1.5 shows the uniform siphoning of probability from $L = \{r, b\}$, which occurs since (FI) doesn't incorporate the information that $\sigma(g, \{r, b\}) = \{r\}$. This is because, according to (FI), the only values of σ which are relevant are of the form $\sigma(v, G_i)$ where $v \in L$, leading to a unidirectional transfer of probabilistic mass.

Despite (FI)'s success concerning monotonicity, Zhang (2024, p. 582) points out that there is a tradeoff to be considered. While (FI) satisfies (MC), it does not satisfy the following condition for update procedures:

Convexity A probabilistic update rule $(\cdot||\cdot)$ is said to be convex just in case for all q, $\mathcal{X}$ and Pr,

$$\min_{\mathrm{X} \in \mathcal{X}} Pr(\cdot||X) \leq Pr(\cdot||q, \mathcal{X}) \leq \max_{X \in \mathcal{X}} Pr(\cdot||X)$$

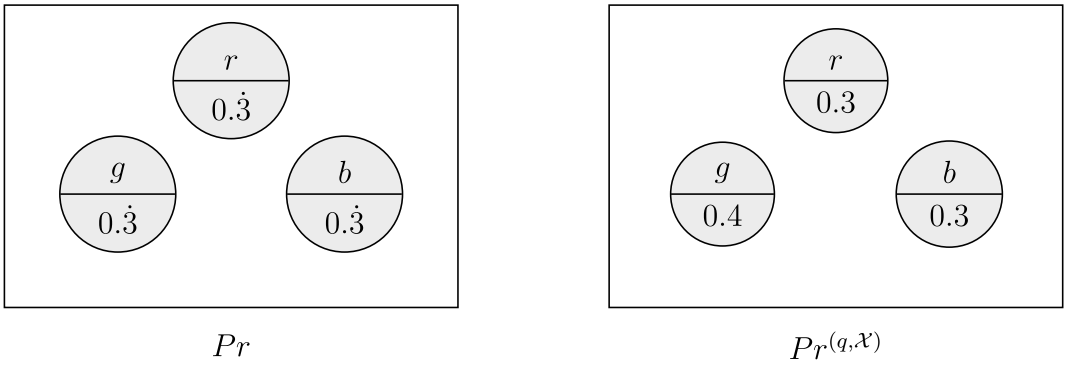

Figure 1.5: Alice's probability measures when considering the RGB diode under Fusco imaging (FI).

As Zhang (2024, p. 589) points out, the only update rule (.||.) which is both convex and conservative is Jeffrey conditioning. As such, those in the market for a type of partial imaging are forced to choose between convexity and conservativity. Being in this market, I choose to adopt (FI) in favour of conservativity (and (MC) alongside it). This is partly due to the plausibility of conservativity as a constraint of belief revision (especially given the Lewisian gloss on imaging as 'minimal shifting of probabilities'), but also due to fact that convexity doesn't carry the same intuitive force as monotonicity and, as previously stated, I'm skeptical that (SMC) recoups the bulk of these intuitions.

1.4.5 An Accommodation Rule for Intensive and Probabilistic Qualification

At this stage, when an agent is informed that the events $X_1, ..., X_n \in \mathcal{X}$ have the respective probabilities $q(X_1), ..., q(X_n)$ there are two kinds of update they can perform to their probability measure Pr in order to accommodate that information: Jeffrey conditioning and Fusco imaging. The current task is to integrate these two procedures into an account of how contexts change when accommodating qualified assertions – which can be understood here as attempts to introduce the constraints set by q and $\mathcal{X}$ to the common background. As a guiding principle, we try to develop further the same division of labour carved out by CONDITIONAL REVISION and IMAGING REVISION, where accommodation should be mediated via Jeffrey conditioning wherever possible and (FI) otherwise. This is in contrast to Eva and Hartmann's (2021) project, which associates indicative supposition with conditionalization and subjunctive supposition with imaging. The reason for this departure is that assertions, as opposed to suppositions, are typically made in relation to propositions which are to be taken for granted in the indicative mood. When a speaker asserts that P, listeners accommodate by acting as if they have all *learned* P from an authoritative source. They then attempt to add P to the common background, which preferably goes by conditionalization due to aforementioned concerns of informational economy but imaging becomes necessary when P is incompatible, for some reason or other, with this common background. The same goes for cases of probabilistic qualification. If a speaker says that the probability of X is q, then listeners accommodate this fact by acting as if they have learned the probability of X is q, by Jeffrey conditioning if possible and (FI) otherwise (when the initial probability of X is zero, for some $X \in \mathcal{X}$). Intensive qualification is pursued as a special type of unqualified probabilistic assertion, following Smith (2008).

Extending the division of labour established by CONDITIONAL REVISION and IMAGING REVISION, let us say that a probability measure $Pr[(q, \mathcal{X})]$ is the revision of Pr, which is defined as:

$$Pr[(q, \mathcal{X})] = \begin{cases} Pr_{(q,\mathcal{X})} & \text{if } Pr(X) \neq 0 \text{ for all } X \in \mathcal{X} \\ Pr^{(q,\mathcal{X})} & \text{otherwise} \end{cases} \tag{Revision}$$

Agents who wish to revise their epistemic state with the information $(q, \mathcal{X})$ do so by replacing Pr with $Pr[(q, \mathcal{X})]$. This allows them to preserve probabilistic evidence when possible and avoid gratuitous shifting of probabilistic mass otherwise.

Now we turn our attention to how this individualistic principle of revision might be used for the accommodation of qualified assertions, but before addressing these new dynamics for qualified assertion, we must supplement contexts with additional details required to facilitate imaging—namely, selection functions σ and transfer functions T. Now, we let contexts $\mathfrak{C}$ be 7-tuples of the form $\langle G, \mathbb{P}, \mathbb{S}, \mathbb{T}, \boldsymbol{W}, @, \nu\rangle$ where all else is held the same except,

$\mathbb{S}$ is a set of selection functions σ_i for each $i \in G$.[31]

$\mathbb{T}$ is a set of transfer functions $T_{w,X,i}$, one for each $w \in \boldsymbol{W}$, $X \in \wp(\boldsymbol{W})$ and $i \in G$.

These two agent-relative functions are new variables along which contexts might be found defective. Recall that the hallmark of a defective context is instances of miscommunication. If agents fail to coordinate their selection and transfer functions, then they will respond to an assertion that P by inflating the probability of entirely different propositions. In a conversation where it is presupposed that Jules the calico is in the living room, when someone asserts that Jules isn't in the living room then, given certain choices of σ for each agent, one listener might come to believe that Jules is in the bathroom while another comes to believe that she is in the kitchen. If the context $\mathfrak{C}$ is such that $\sigma_i = \sigma_j$ for all $i, j \in G$ and $T_{w,X,i} = T_{w,X,j}$ for all $i, j \in G$, $w \in \boldsymbol{W}$ and $X \in \wp(\boldsymbol{W})$, then we say that $\mathfrak{C}$ is *non-σ-defective*, and we can drop the agent-wise index on σ and $T_{w,X}$. If not, then we call $\mathfrak{C}$ *σ-defective.*

First, let's consider the case of a context accommodating a probabilistically qualified assertion. We take these assertions to be of the form $\langle q, \mathcal{X} \rangle$ where q is a series of input probabilities for the cells of partition $\mathcal{X}$. We treat simple assertions like "The likelihood of X is x" as a special instance of this form where $\mathcal{X} = \{X, \overline{X}\}$ and $q(X) = x$ and $q(\overline{X}) = 1 - x$. Unqualified assertions, where a speaker simply asserts X, are treated as further instances of this form where $q(X) = 1$ and $q(\overline{X}) = 0$. How do listeners accommodate the assertion $\langle q, \mathcal{X} \rangle$? Well, they do so by all coming to believe that it is a common belief that all speakers $i \in G$ have a probability measure which agrees with q on the probabilities of $X_1, ..., X_n$—i.e., by coming to have a series of (vagueness free) partial presuppositions about events in $\mathcal{X}$. To see how this kind of accommodation might work, let's first consider the simplest case of a context which is neither defective nor σ-defective. This means that there is a unique context set $\boldsymbol{C}_{\mathfrak{C}}$, which is the set of all worlds which speakers in G entertain as live possibilities. Furthermore, it means that $\mathbb{S}$ contains a single, mutually agreed upon, selection function σ and $\mathbb{T}$ is restricted to having one transfer function per world-event pair. Suppose that someone

---

[31] Earlier, $\mathbb{S}$ was used to denote a system of spheres. We are now repurposing this notation for a set of functions σ which play the same role as spheres re: selecting a "closest" set of alternatives relative to some condition X.

makes the assertion $\langle q, \mathcal{X}\rangle$ relative to this non-defective context $\mathfrak{C}$. Speakers in G can accommodate $\langle q, \mathcal{X}\rangle$ by presupposing X to degree $q(X)$, for each $X \in \mathcal{X}$—i.e., by making it true that $\mathbf{B}_i\mathbf{C}(q(X))_G X$ for each $i \in G$ and $X \in \mathcal{X}$.[32] This can be achieved using a tool from Fagin et al.'s (1995, p. 24) epistemic logic. There, they have recourse to a notion called G-reachability. Let a world v be *G-reachable* from another world w just in case there is a chain of worlds $w_1, ..., w_n$ where $w = w_1$, $v = w_n$ and for any pair $\langle w_n, w_{n+1}\rangle$ in the chain $Pr^i_{w_n}(\{w_{n+1}\}) \neq 0$ for some speaker $i \in G$. As Fagin et al. (1995, p. 24) note, $\mathbf{C}_G P$ is true at a world w if and only if P is true in all worlds which are G-reachable from w.[33]

Now, the goal for accommodating say, X, is to make $\mathbf{C}_G X$ true at each world in the context set, which thereby makes speakers in G presuppose X since speakers believe what is true across their context set. Given G-reachability, there appears to be a procedure for making $\mathbf{C}_G X$ true at w: simply take all worlds which are G-reachable from w and revise an agent's probability measure at w with the information that X. However, the matter is slightly more complicated than this. Changing a probability measure Pr^i_w changes which worlds are G-reachable from w, so merely revising Pr^i_v for all v which are G-reachable from w might fail to make it true that there is a common belief that X at w. Instead, accommodation must be facilitated by a step-by-step procedure. This leads us to the accommodation rule for probabilistic qualifications.

Qualification Accommodation If at time t a speaker successfully asserts $\langle q, \mathcal{X}\rangle$ and just before t the context is $\mathfrak{C}$, then at t $\mathfrak{C}$ becomes $\mathfrak{C}'$, which is exactly the same as $\mathfrak{C}$ except $\mathbb{P}$ is replaced with $\mathbb{P}(q, \mathcal{X})$ which is defined as follows:

> Let $\mathbb{P}^0$ be exactly like $\mathbb{P}$ except all measures of the form $Pr^i_{@}$ are replaced with $Pr^i_{@}[(q, \mathcal{X})]$. Furthermore, let $s(\mathbb{P}^0)$ be the subset of measures in $\mathbb{P}^0$ of the form $Pr^i_{@}[(q, \mathcal{X})]$.

---

[32] Recall that I am ambiguously using uppercase italic letters like X to mean either proposition letters in some language $\mathcal{L}$ or events in $\wp(\boldsymbol{W})$. When X appears in $q(X)$ it is playing the role of an event, whereas when X appears in a formula like $\mathbf{C}_G(0.5)X$ it is playing the role of a proposition.

[33] Fagin et al. (1995, p. 24) also point out that this result holds even when dropping the requirement, which we have dropped, that R_i be an equivalence relation.

> Let $\mathbb{P}^1$ be exactly like $\mathbb{P}^0$ except all measures of the form Pr^i_w are replaced with $Pr^i_w[(q, \mathcal{X})]$ where w is such that $Pr(\{w\}) \neq 0$ for some $Pr \in s(\mathbb{P}^0)$. Furthermore, let $s(\mathbb{P}^1)$ be the subset of measures in $\mathbb{P}^1$ of the form $Pr^i_w[(q, \mathcal{X})]$.
>
> In general, let $\mathbb{P}^n$ be exactly like $\mathbb{P}^{n-1}$ except all measures of the form Pr^i_w are replaced with $Pr^i_w[(q, \mathcal{X})]$ where w is such that $Pr(\{w\}) \neq 0$ for some $Pr \in s(\mathbb{P}^{n-1})$. Let $s(\mathbb{P}^n)$ be the subset of measures in $\mathbb{P}^n$ which are of the form $Pr^i_w[(q, \mathcal{X})]$.
>
> Then, $\mathbb{P}(q, \mathcal{X}) = \mathbb{P}^m$ for the smallest number m such that $\mathbb{P}^m = \mathbb{P}^{m+1}$

The idea here is that speakers accommodate $\langle q, \mathcal{X} \rangle$ by first revising their probability measures, then revising the measures held in all worlds accessible from @, then revising the measures accessible from these worlds and so on until the procedure fails to generate new set of probability measures. The procedure will always terminate (i.e., there will always *be* a smallest m such that $\mathbb{P}^m = \mathbb{P}^{m+1}$) because (i) $\boldsymbol{W}$ is finite, so there will not be an endless chain of probabilities to revise for each world in an infinitely large set and (ii) probabilistic revision (as we've defined it) is conservative, so when this process reaches a probability measure which has already been revised there's no possibility of expanding the list of measures which need to be revised at a later point. This iterative process creates the presuppositions necessary for accommodation by ensuring that all probability measures along the newly generated chain of worlds which are G-reachable from the context set in $\mathfrak{C}'$ obey the constraint $(q, \mathcal{X})$, thus making all of G believe that there is a common $q(X)$-strength belief that X for all $X \in \mathcal{X}$.

An example of QUALIFICATION ACCOMMODATION is in order. Suppose that a group G of speakers are speculating about the weather tomorrow. The realm of possibilities $\boldsymbol{W} = \{@, i, ii, iii\}$—where @ is a world where the temperature fluctuates between 20 and 30 degrees, i between 10 and 20, ii between 5 and 10 and iii between 0 and 5. Furthermore, suppose that the speakers have initial probability measures which are all in agreement across $\boldsymbol{W}$, such that we can speak *as if* $\mathbb{P}$ contains one probability measure Pr_w for each world $w \in \boldsymbol{W}$. Let these measures be as indicated in Figure 1.6, where they split their credence between @ and i if $\{@, i\}$ holds and they are certain that $\{ii\}$ or $\{iii\}$ if either of them hold. A trusted meteorologist in G suggests to the group that there's a 20% chance of a cold front tomorrow which would cause the event $\{ii, iii\}$ obtain. Thus, the meteorologist asserts $\langle q, \mathcal{X} \rangle$ where

$\mathcal{X} = \{\{@, i\}\{ii, iii\}\}$, $q(\{@, i\}) = 0.8$ and $q(\{ii, iii\}) = 0.2$. Let $W = \{@, i\}$ and $C = \{ii, iii\}$, for *warm* and *cold* respectively. Furthermore, let's stipulate that σ always returns a set of worlds with the smallest possible deviation in terms of temperature—e.g., $\sigma(ii, \{@, i\}) = \{i\}$, $\sigma(i, \{@, ii, iii\}) = \{@, ii\}$ and so on. Given this set-up, QUALIFICATION ACCOMMODATION deems that when the speakers accommodate the meteorologist's assertion they will adopt a new context $\mathfrak{C}'$ which is exactly like $\mathfrak{C}$ except it contains $\mathbb{P}(q, \mathcal{X})$ in place of $\mathbb{P}$. In this case, $\mathbb{P}(q, \mathcal{X})$ is $\mathbb{P}^1$ since $\mathbb{P}^0 \neq \mathbb{P}^1$ and $\mathbb{P}^1 = \mathbb{P}^2$. Relative to this new context $\mathfrak{C}'$, it can be seen that the speakers in G have an 0.8-strength presupposition that W and a 0.2-strength presupposition that C, since $\nu(\langle @, \mathbf{B}_i\mathbf{C}(0.8)_G W\rangle) = \nu(\langle @, \mathbf{B}_i\mathbf{C}(0.2)_G C\rangle) = 1$ for all $i \in G$. Figure 1.7 shows the details of this revised set of measures $\mathbb{P}(q, \mathcal{X})$.

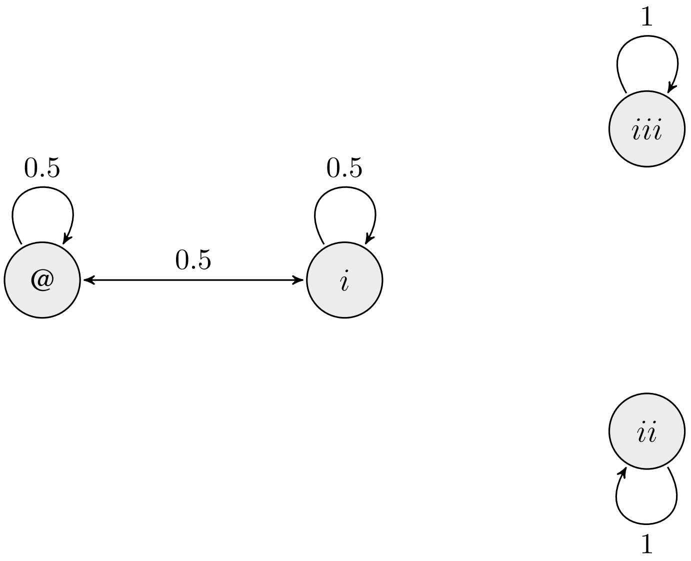

Figure 1.6: $\mathbb{P}$ for the weather speculators

Expanding this accommodation rule to defective and σ-defective contexts complicates the explanation slightly, but doesn't necessitate any change to QUALIFICATION ACCOMMODATION. When running the same process on a defective context, a $q(X)$-strength presupposition that X is still fostered amongst all worlds in each speaker's context set, despite these context sets potentially failing to overlap. This is because the initial step takes *all* probability measures of the form $Pr^i_@$, revises them and then places them in $s(\mathbb{P}^0)$,

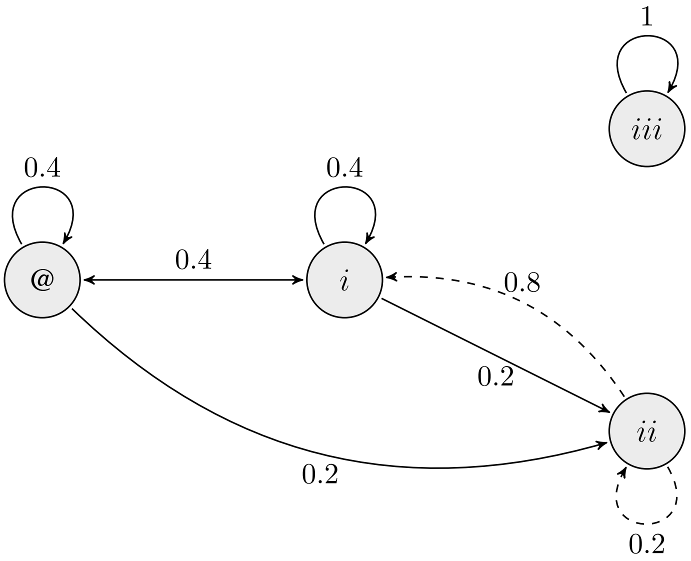

Figure 1.7: $\mathbb{P}(q, \mathcal{X})$ for the weather speculators, which in this case is $\mathbb{P}^1$

meaning the set of measures which are to be revised at the next step are those which are held in *any* world that is in at least one agent's context set. Then, the process continues as it does in the non-defective case, making it true at each of these worlds that there is a common $q(X)$-strength belief that X.

The only complication added by σ-defective contexts is that the values of $Pr_w^{i(q,\mathcal{X})}$ are determined by $\sigma_i \in \mathbb{S}$, rather than some shared selection function. Note that this means accommodation in σ-defective contexts can take agents from a non-defective context to a defective one. If speakers i and j in a non-defective context presuppose $\neg X$ but then accommodate X, their disagreement concerning σ_i and σ_j could make their context sets diverge in the initial step. This is what characterises the aforementioned kind of miscommunication which occurs when speakers fail to coordinate selection functions.

As for intensive qualification, this is best understood as a special type of unqualified assertion. As suggested in Smith (2008, p. 248), when a speaker makes an n-strength assertion that P, they are asking for the conversation to take the event P_n for granted, where $w \in P_n$ iff $\nu(\langle w, P\rangle) = n$. This can be achieved through the accommodation of the unqualified assertion $\langle q, \mathcal{X}\rangle$

where $\mathcal{X} = \{P_n, \overline{P_n}\}$ and $q(P_n) = 1$. Unlike Smith's (2008) account, using probabilistic qualification as developed here allows for the accommodation of n-strength assertions that P even in cases where P is presupposed to be true to some degree other than n. This way of characterising intensive qualification also allows us to explain how combinations of probabilistic and intensive qualifications work. When a speaker says something like "It is 70% likely that tomorrow it will rain to degree 0.5", they are making an assertion of the form $\langle q, \mathcal{X}\rangle$ where $\mathcal{X} = \{R_{0.5}, \overline{R_{0.5}}\}$, $q(R_{0.5}) = 0.7$ and $q(\overline{R_{0.5}}) = 0.3$.

1.4.6 Partial Accommodation

Partial accommodation, discussed in §1.2.2, is distinct from but follows the same kind of contextual dynamics as probabilistic qualification. When speakers have limited confidence in someone who makes an unqualified assertion, they might act *as if* they have made a probabilistic qualification and accommodate accordingly. Take the example from §1.2.2 where, in a game of pub trivia, two speakers give an answer (without qualification) but the trivia group responds by equivocating their confidence between the two. Here, the effect is the same as if they took the speakers to be making a kind of "distributed" probabilistic qualification of the form $\langle q, \mathcal{X}\rangle$ where the answers $A_1, A_2 \in \mathcal{X}$ are given the respective probabilities $q(A_1)$ and $q(A_2)$, which indicate the group's respective confidence in each. This way, QUALIFICATION ACCOMMODATION is also able to model (part of) the "certain limits" proviso which Lewis (1969, p. 340) includes in his definition of accommodation. A conversation will only ever accommodate an assertion that X to a degree which is proportional to the amount of confidence they have in the assertor's judgements concerning matters related to X.

Of course, this doesn't capture the entire spectrum of limitations to accommodation. If speakers have no confidence whatsoever in someone's assertions concerning X, it wouldn't be prudent for them to accommodate any assertion of the form $\langle q, \mathcal{X}\rangle$ where $X \in \mathcal{X}$ (even if $q(X) = 0$) since doing so would treat the assertor as if they have strong probabilistic information concerning X. Similar concerns arise when X is asserted by someone who isn't a legitimate party to the conversation. While these cases deserve attention, this project will only concern itself with assertions which have at least some amount of purchase in the context being considered. This is why rules like QUALIFICATION ACCOMMODATION are concerned with what happens to a context once a *successful* assertion of $\langle q, \mathcal{X}\rangle$ is made.

Chapter 2

Partial Proof

In this second half of the thesis, the theory of partial presupposition developed in Chapter 1 is utilised to give an analysis of proof where a proof that P's degree of strength corresponds to how successful it is at raising the audience's confidence in P according to their common ground. As a first pass, this definition does well at explaining how proofs are evaluated by (and enact changes to) the context in which they are performed. However, this first pass definition also leads to multiple issues. It results in the undesirable consequence that proof is closed under entailment—that is, whenever P entails Q, a proof that P automatically constitutes a proof that Q. To avoid this result, I expand the sample space $\boldsymbol{W}$ so that agents can have diverging presuppositions *about* such entailments, thereby saving this definition of proof from these problems of logical closure. Furthermore, the first pass definition misses an auxiliary Gricean kind of requirement on proofs, which is that they must be *appropriate* as construed in various ways. Proofs, like other contributions to communicative exchanges, are often subject to expectations about informativeness and relevance. So, a proof which strongly confirms P might still be inappropriate insofar as it exhibits *excessive* or *insufficient* amounts of information or fails to be relevant. This Gricean program gives us an opportunity to enrich the contexts from Chapter 1 with additional components that track these conversational standards.

In this chapter, I also consider the idea of a standard of proof, a concept whose domain ranges from courts of law to idle chitchat. These standards of proof are taken to be a species of accommodation rule which determine how strongly an assertion's content must be confirmed before speakers are willing to accept it. In a discussion of legal standards of proof, I consider some well-known objections to probabilistic interpretations of these standards and ultimately suggest that they should be interpreted using non-probabilistic degrees of proof. After this discussion, I use the notion of verisimilitude to explain how, in some conversations, the associated standard of proof is so low that speakers are willing to accept assertions of propositions which they presume to be false.

2.1 Strength of Proof

According to *relevant alternatives* theories of knowledge, of the kind espoused by Stine (1976), Goldman (1976), Dretske (1970, 1981a) and Lewis (1996), coming to know that P is a matter of ruling out a series of possible scenarios where P isn't true. For instance, you know that the deck of cards on the table contains no aces because (i) you have two aces in your hand and (ii) you've peeked at the hands of two other ace-holding players. Observations (i) and (ii) rule out two possible ways the deck could've had at least one ace. There are other ways the deck could have an ace. It might have an additional ace if, for instance, the deck was accidentally supplied with five aces at the card factory or if someone absentmindedly added cards from another deck when shuffling. However, the relevant alternatives theorist suggests that you nonetheless know the deck has no aces because you are able to dismiss these far-fetched possibilities which your observations haven't ruled out as irrelevant to your knowledge. Things may have been different if you'd read a headline earlier that week about increased clumsiness at playing card factories or if you noticed a suspiciously thin deck of cards somewhere near the main deck. In these circumstances, the two possibilities your observations failed to rule out become relevant, making it no longer true that you know the deck is ace-less.

What does all this talk of relevant alternatives have to do with a proof's strength? Well, I am proposing an account of proof where something only counts as a proof of P if it is able to eliminate a series of alternative possibilities where P is not true. For a speaker to prove P in a given conversation, they must perform some conversational move which eliminates all alternatives to P which other speakers entertain. In brief, proofs must allay doubts. I will begin to spell out this account of proof by first considering some different ways of determining what counts as a relevant alternative to P.

Let an alternative to P be a world in $\boldsymbol{W}$ where P is not true. In other words, alternatives to P are members of $\overline{P}$. There are two broad kinds of fact which make an alternative to P relevant to an agent's knowledge of P:

A Facts about the putative knower's beliefs, presuppositions and evidence.

B Facts about the circumstance which the putative knower finds themselves in.

A-type facts include things such as the putative knower having reason to believe an alternative is possible (Stine, 1976, p. 252) or presupposing the alternative is possible in a given conversation (Lewis, 1996, p. 554). B-type facts include things such as whether the putative knower is in a scenario where the alternative to P *could have*, in some objective sense, occured (Dretske, 1970, p. 1021). If you pin the relevance of an alternative on A-type facts, then one of your opponents merely saying that they read an article about printing errors in card factories could be sufficient to overrule your knowledge that the deck is ace-less. Their assertion suddenly makes the possibility of a fifth ace being in the deck relevant. Whereas, if you pin the relevance of an alternative on B-type facts, there must be some sense in which it really was possible for the deck to have an additional ace—e.g., if another deck of cards on the shelf next to the one which was purchased contained five aces. Some relevant alternative theories of knowledge use some combination of A and B-type facts in their analysis of relevance. For instance, Lewis's (1996) theory suggests an alternative is relevant just in case speakers *legitimately* presuppose it is possible, where legitimacy of presupposition is constrained by B-type facts about the circumstance of the speakers. For Lewis, the assertion about misprints at the card factory may overrule my knowledge about the remaining cards in the deck but only if the presuppositions introduced by this assertion are ones which could've been true in the objective sense.

If we turn our attention to proofs instead of knowledge, what counts as a relevant alternative to P? For the purposes of this project, where I am attempting to explain the conditions under which a proof is successful at introducing presuppositions, it is primarily A-type facts which determine relevance. This is because proving P successfully is a matter of changing the beliefs of speakers in the context where the proof is being performed. If one of your opponents openly harbours the presupposition that the deck has an inexplicable fifth ace, then—no matter how irrational this presupposition is—revealing that there are four aces in circulation does nothing to prove, to them, that the deck is ace-less. You will have to perform some other proof, such as going through the deck card-by-card and showing that each is not an ace, to convince the skeptical opponent of this fact.

This is not to deny that a complete analysis of proof might also incorporate some B-type facts as well. Perhaps such an analysis could be given in the style of Lewis's (1996) hybrid account of knowledge. However, the primary aim of this thesis is to provide a characterisation of proofs in terms of their effect on a context. My target is the sort of proof which Moore (1993[1939], p. 167) discusses when defending his infamous proof of the external world, which consisted in holding up his hands and intoning 'Here is one hand, and here is another':

> Are there any other conditions necessary for rigorous proof, such that perhaps it did not satisfy one of them? Perhaps there may be; I do not know; but I do want to emphasise that, so far as I can see, we all of us do constantly take proofs of this sort as absolutely conclusive proofs – as finally settling certain questions, as to which we were previously in doubt.

Moore's proof of the external world makes it evident that proofs which successfully allay doubts might still be lousy when considered for their rigour or veridicality. A proof's strength, according to this metric, is simply the extent of its ability to convince others.

The A-type facts which determine whether an alternative is relevant to a putative proof that P will be facts about the context $\mathfrak{C}$ in which the attempted proof is performed. In particular, I suggest that w is a relevant alternative to P just in case,

(i) $w \in \overline{P}$

(ii) w is G-reachable from at least one context set $\boldsymbol{C}_i$ such that $i \in G$

Condition (i) merely ensures that w is an alternative to P, whereas condition (ii) ensures relevance. If $\mathfrak{C}$ is non-defective, then condition (ii) amounts to the claim that w is G-reachable from the sole context set $\boldsymbol{C}_{\mathfrak{C}}$. Defining relevant alternatives this way captures the idea that, at the level of a conversation, a doubt about P persists when parties to that conversation can't take P for granted. If someone in G thinks it is possible that P isn't true, then P cannot be taken for granted. Furthermore, if someone in G thinks it is possible that someone in G thinks it is possible that P isn't true, then P cannot be taken for granted. The only way P can be taken for granted amongst the speakers in G is if there are no alternatives to P which can be reached by a chain of subjective possibility judgements from the speakers in G—i.e., if there is a presupposition that P.

Given this characterisation of relevant alternatives, a speech act counts as a successful proof of P just in case it is able to remove all relevant alternatives to P which populate a context. That is, when an agent proves or demonstrates P relative to a context $\mathfrak{C}$, they present some reasoning or evidence which has the effect of revising $\mathfrak{C}$ with a constraint $(q, \mathcal{X})$ that eliminates (reduces to probability zero for each speaker at each G-reachable possibility) all relevant alternatives to P. This framework gives a preliminary pass/fail account of a proof's success:

Successful Proof A successful proof of the proposition P is some speech act which replaces the context $\mathfrak{C}$ in which it is performed with a context $\mathfrak{C}'$ which contains no relevant alternatives to P, thereby making it a presupposition that P amongst all $i \in G$.

This pass/fail criterion for a proof's success captures two important features of proofs: (i) they make some change to the context in which they are performed and (ii) whether they are successful depends on the presuppositional structure of the context in which they are performed. Once P is successfully proved, speakers take P for granted in their conversation; this is accounted for by the fact that successful proofs have the same effect as successful assertions of P. As for (ii), suppose I attempt to prove that the city of Bolzano is in Italy by showing that it isn't in Austria. Perhaps I achieve this by procuring a map of Austria and pointing out that Bolzano is nowhere to be found. This proof will be successful in contexts where it is presupposed that Bolzano is either in Italy or Austria, but unsuccessful in contexts where it is presupposed that Bolzano is in Italy, Austria or Switzerland.

As it currently stands, successful proof counts unchallenged assertions of P as proofs that P. To make room for a distinction between the two, it's helpful at this point to think of proofs in relation to Austin's (1962, p. 5) notion of a *performative*, which he defines as utterances of sentences which:

A. do not 'describe' or 'report' or constate anything at all, are not 'true or false'; and

B. [are, or are] part of, the doing of an action, which again would not *normally* be described as saying something.

Proofs can be thought of as any utterance *or* act which has the effect of showing that P to some audience, in contrast with an assertion that P, which clearly *does* describes the world as being a certain way and *is* part of what would normally be called saying something. Of course, proving P might involve saying or describing things which are used to support P, but the proof itself—the act which is constituted by these assertions, descriptions and other behaviour—is something which we would hesitate to call an assertion. Here are some illustrative examples of proof along these lines,

> A teacher proves that Pythagoras' theorem holds for a particular right-angle triangle by drawing it on the whiteboard and attaching squares to each of the sides—showing that the area of the square attached to the hypotenuse is the same as the combined area of the other two.
>
> Someone shows that the word "obsolescence" has four syllables by uttering it slowly, emphasising each syllable.
>
> A private detective demonstrates that Mr. X was at the Copacabana last night by showing pictures of Mr. X dining with Madame Y at the Copacabana.

Despite not saying or describing the world being a certain way, these proofs have the same effect on conversations as assertions. They add presuppositions to the common ground. So, what is the exact relationship between the two? Proofs can be seen as an opportunity for those who are held in epistemic disrepute with respect to P to ensure the success of an assertion that P. That is, proofs are often given in contexts where a speaker's claim has been implicitly or explicitly rejected by others. One can naturally imagine each of the foregoing examples as following a failed assertion attempt—e.g., Mr. X challenges the detective to substantiate the claim that he was at the Copacabana, or the students ask how it could possibly be that c^2 is the sum of a^2 and b^2 when the hypotenuse c is *this* short.

This idea that proofs are offered to support unsuccessful assertions fits quite naturally with an idea from the medieval practice of disputation where a 'burden of proof' for any given claim is placed upon its proponent, as detailed by Rescher (2007, p. 17). The same allotment of responsibilities more or less holds in regular conversation. If you attempt to assert that P and this assertion is unsuccessful, then insofar as you would like your assertion to be accommodated it is your responsibility as the one proposing that P to show its truth. Lewis's (1979, p. 347) proposal that conversations are naturally accommodating suggests that this responsibility is met as soon as it is created. However, we have already explained that if listeners hold you in epistemic disrepute with respect to P, you will not be able to shirk this burden of proof so easily. This is the role which proofs play in conversations: they are opportunities to meet the un-waived burden of proof which threatens failed assertion attempts.

In meeting this challenge, the proof which an assertor provides might also be more-or-less convincing to the audience which they are trying to convince. This can be reflected by allowing proofs, like assertions, to come in varying degrees of strength, leading to the following graded criterion of success for proofs:

Strength of Proof When a proof is performed relative to some context $\mathfrak{C}$ which is neither defective nor σ-defective, its degree of strength with respect to the proposition P is the degree of presupposition that P in the resulting context $\mathfrak{C}'$ which contains the revised set of probability measures $\mathbb{P}(q, \mathcal{X})$—i.e., the strength of a proof that P is the unique n (if there is one) such that, on $\mathfrak{C}'$, $\mathbf{B}_i\mathbf{C}(n)_G P$ is true at @ for all $i \in G$.

If there is no degree n to which the proposition P is presupposed after a proof has been performed, then this proof is not assigned a degree of strength in relation to the proposition P. In cases like these, the proof hasn't established P to any particular degree across all worlds which are G-reachable from @. This might be the case because at some of these worlds speakers have an x-strength presupposition that P whereas in others they have a y-strength presupposition that P, where $x \neq y$. Since proofs like this don't bring speakers to a consensus regarding their confidence in P, they don't count as proofs of P at all.

When dealing with a defective (or σ-defective) context $\mathfrak{C}$, then a proof's degree of strength is always relative to some speaker in G. This is due to the fact that, in these contexts, speakers can have diverging degrees of presupposition with respect to the proposition being proved. So a proof that might bring someone to the presupposition that P might fail to bring others to that same presupposition. A graded notion of proof can be preserved for these defective contexts, so long as we drop the assumption that there is a singular degree of proof and relativise this degree to each speaker in G:

Strength of Proof (for i) When a proof is performed relative to some context $\mathfrak{C}$, its degree of strength with respect to the proposition P for some agent i in the resulting context $\mathfrak{C}'$ is the unique n (if there is one) such that on $\mathfrak{C}'$, $\mathbf{B}_i\mathbf{C}(n)_G P$ is true at @.

In other words, the strength of a proof that P for i is the strength of i's presupposition that P once the proof has been performed. As in the non-defective case, if a proof fails to take i to a belief about some consensus confidence in P amongst speakers in G, then the speech act doesn't count as a proof of P at all.

2.1.1 Logical Omniscience and Probative Omnipotence

There is a problem engrained in possible-worlds accounts of epistemic logic which also threatens the foregoing account of proofs. Hintikka's (1962) landmark work in epistemic logic begins with a presentation of some criteria for the consistency of sets of propositions featuring epistemic operators like $\mathbf{K}_i$ and $\mathbf{B}_i$. Hintikka (1962, p. 30) points out that, according to these criteria, the set $\{\mathbf{K}_iP, \neg\mathbf{K}_iQ\}$ is inconsistent wherever P implies Q (that is, whenever the set $\{P, \neg Q\}$ is inconsistent). However, the inconsistency of $\{\mathbf{K}_iP, \neg\mathbf{K}_iQ\}$ seems to conflict with ordinary usage of the term "knowledge", since an agent i may fail to know complicated or obscure consequences of P even when they know that P. The inconsistency of this set isn't a quirk of Hintikka's system, either. In the epistemic logic we built upon in Chapter 1, from Fagin et al (2003), we have the same ruling on the inconsistency of $\{\mathbf{K}_iP, \neg\mathbf{K}_iQ\}$ whenever P entails Q: there's no valuation ν which assigns 1 to both propositions at any world. This problem has come to be known as the problem of *logical omniscience*. It challenges us to find an account of knowledge which doesn't treat all agents as knowers of all consequences of their knowledge.

Hintikka's (1962, p. 31) initial solution to this problem was to reframe the project of epistemic logic so that, instead of giving an account of the consistency/inconsistency of sets of propositions, it gives an account of which sets are *defensible* or *indefensible*. So, while it might be possible for $\mathbf{K}_iP$ and $\neg\mathbf{K}_iQ$ to both be true when P implies Q, it is not a position which i can be expected to hold onto in the long run. If someone were to patiently and carefully show i that Q follows from what they already know, then presumably i would eventually reach a position where the incompatibility of $\mathbf{K}_iP$ and $\neg\mathbf{K}_iQ$ is made transparent to them. Thus, while an agent can know P without knowing P's consequences, these consequences are implicitly contained within the epistemic commitments which grant their knowledge that P. In Hintikka's (1962, p. 31) own words:

> suppose a man says to you "I know that p but I don't know whether q" and suppose that p can be shown to entail logically q by means of some argument which he would be willing to accept. Then you can point out to him that what he says he does not know is already implicit in what he claims he knows.

This solution might work for the problem of logical omniscience which threatens epistemic logic,[1] but it's entirely inadequate for solving a correlate of this problem as applied to proofs. This is because proving P and proving that Q follows from P might require entirely different acts which aren't clearly related, even when Q is a consequence of P.

Furthermore, this idea that an agent can be *shown* some of the consequences of their knowledge by an argument which they'd accept inevitably involves some kind of proof. So, taking a proof that P and performing an additional proof that Q is a consequence of P doesn't mean that a proof of Q was already 'contained' in the initial proof of P, much in the same way that mixing blue paint with yellow paint doesn't show that green paint was already there, implicitly 'contained' within the blue paint.

For example, suppose we are at a fair and I demonstrate that the prize jar contains 82 jellybeans by pulling them out one-by-one and counting to 82 as I do so. Unless you already knew that 82 is even, this demonstration alone would do very little to convince you that the jar contains an even number of jellybeans. I might be able to prove the evenness of 82 with a different demonstration (e.g., by making 41 pairs of jellybeans which exhaust the jar's contents) but there is no sense in which this auxiliary proof is somehow contained within my one-by-one counting of the jellybeans. Showing the cardinality of some set in terms of an integer and showing the parity of that integer may require different courses of action, even if facts about that set's cardinality entail facts about the corresponding integer's parity.

As it currently stands, however, the proposed theory of proof doesn't allow me to prove that there are 82 jellybeans in the jar *without* proving that there are an even number of them. Let J_{82} be the proposition that the jar contains 82 jellybeans and let J_E be the proposition that the jar contains an

---

[1] That being said, Hintikka (1979, p. 370) ultimately shied away from this kind of response to the problem of logical omniscience as applied to epistemic logic as well.

even number of jellybeans. If I successfully prove J_{82} with some act, then I make it so that every world which is G-reachable from the context set (or *sets* plural, if the context is defective) is one at which J_{82} is true. Every possibility in which there are 82 jellybeans in the jar is also one in which there are an even number of jellybeans in the jar. So, if I manage to prove J_{82} in a given context, I will also inadvertently prove J_E. Let's call this proof-wise version of logical omniscience the problem of *probative omnipotence.* This problem challenges us to find an account of proof which allows for agents to prove some proposition without proving everything which that proposition entails.

Earlier, in §1.3, I forewent the selection of a particular logic underlying the valuation functions ν. That is, I chose to give neither (i) some constraint on how ν assigns truth values to compound propositions whose main connective is $\neg, \wedge, \vee, \rightarrow$ or $\leftrightarrow$ nor (ii) some account of what it takes for an argument to be valid in terms of ν so constrained. This was done to demonstrate that the structure of a context will be more or less the same, irrespective of how truth values are assigned to complex propositions or how one unpacks the notion of a proposition *following* from others. In this section however, it will be crucial to choose a logic (and more specifically, an account of validity), since the problem of probative omnipotence is always relative to a particular account of logical consequence. As will become clear later in this section, the particular logic I have chosen is only used to illustrate one way the problem of probative omnipotence might be broached. The solution I provide is compatible with a wide range of logical intuitions and I encourage readers to consider what is said below in relation to their logic of choice.

The semantics which we choose will be the slightly modified Łukasiewicz style semantics for fuzzy logic given in Smith (2008). First, each function ν assigns arguments of the form $\langle w, P\rangle$ some value in $\mathfrak{T}$, where $w \in \boldsymbol{W}$ and P is a propositional constant. There is no constraint concerning what truth values these propositional constants take at a world—that is, any combination of truth values can be assigned to all basic propositions at any world. Once the ν assigns a truth value to all world-constant pairs, the truth values of compound propositions at all worlds are determined recursively. Wherever α and β are well-formed sentences of the propositional language $\mathcal{L}$ and $w \in \boldsymbol{W}$:

$$(\neg)\ \ \nu(\langle w, \neg\alpha\rangle) = 1 - \nu(\langle w, \alpha\rangle)$$

$$(\wedge)\ \ \nu(\langle w, (\alpha \wedge \beta)\rangle) = \min\{\nu(\langle w, \alpha\rangle), \nu(\langle w, \beta\rangle)\}$$

($\vee$) $\nu(\langle w, (\alpha \vee \beta)\rangle) = \max\{\nu(\langle w, \alpha\rangle), \nu(\langle w, \beta\rangle)\}$

($\rightarrow$) $\nu(\langle w, (\alpha \rightarrow \beta)\rangle) = \max\{1 - \nu(\langle w, \alpha\rangle), \nu(\langle w, \beta\rangle)\}$

Furthermore, given the epistemic operators defined in §1.3.1-2 above, we can also expand our semantics to assign truth values to compound propositions whose main operators are these epistemic operators like $\mathbf{B}_i$ and $\mathbf{C}(n)_G$. While the foregoing rules for propositional connectives were extensional, the rules for these epistemic operators are intensional in the sense that their truth value depends on factors outside of ν, such as the structure of $\mathbb{P}$. Given a group of speakers G, some set of probability measures $\mathbb{P}$ and a formula α from $\mathcal{L}$, here are the semantic clauses for each of the four epistemic operators introduced in §3:

($\mathbf{B}_i$)

$$\nu(\langle w, \mathbf{B}_i\alpha\rangle) = \begin{cases} 1, & \text{if } \nu(\langle v, \alpha\rangle) = 1 \text{ for all } v \text{ such that } Pr^i_w(\{v\}) = 1 \\ 0, & \text{otherwise} \end{cases}$$

($\mathbf{B}(n)_i$)

$$\nu(\langle w, \mathbf{B}(n)_i\alpha\rangle\rangle) = \begin{cases} 1, & \text{if } E^i_w(\alpha) = n \\ 0, & \text{otherwise} \end{cases}$$

($\mathbf{C}_G$)

$$\nu(\langle w, \mathbf{C}_G\alpha\rangle) = \begin{cases} 1, & \text{if } \nu(\langle v, \alpha\rangle) = 1 \text{ for all } v \text{ which are } G\text{-reachable from } w \\ 0, & \text{otherwise} \end{cases}$$

($\mathbf{C}(n)_G$)

$$\nu(\langle w, \mathbf{C}(n)_G\alpha\rangle) = \begin{cases} 1, & \text{if } \nu(\langle w, \mathbf{C}_G\mathbf{B}(n)_i\rangle) = 1 \text{ for all } i \in G \\ 0, & \text{otherwise} \end{cases}$$

One might wonder why, given the infinite cosmos of truth-values in $\mathfrak{T}$, propositions like $\mathbf{B}(n)_iP$ which use epistemic operators are only ever true (1) or false (0) at a world given some valuation ν. The idea behind this enforced bivalence is that ascriptions of degrees of belief are all-or-nothing even if the beliefs which they describe come in degrees. For comparison, it is either completely true or completely false to say that someone is taller than 180cm, even though that person's height comes in degrees.

Now that these valuation functions assign truth-values to all world-proposition pairs in $\boldsymbol{W} \times \mathcal{L}$, we can present a definition of logical consequence. The fuzzy semantics given above is compatible with a wide spectrum of intuitions about what it takes for an argument to be valid. Even paraconsistent sympathisers have used Łukasiewicz semantics to provide a definition of validity which doesn't allow arbitrary conclusions to follow from jointly inconsistent premises.[2] However, in attempt to avoid controversy for what is ultimately a demonstrative example, we will follow Smith (2008, p. 222) and use these semantics to define a consequence relation which coincides with classical logic—or rather, in our case, one which is at least classical for the fragment of $\mathcal{L}$ which contains no epistemic operators. To do so, let an argument from some set of premises Γ to a conclusion β (denoted $\Gamma \therefore \beta$) be valid just in case there is no world w and valuation ν such that $\nu(\langle w, \alpha \rangle) > 0.5$ for all $\alpha \in \Gamma$ while $\nu(\langle w, \beta \rangle) < 0.5$. That is, all valuations which assign all premises a degree of truth greater than 0.5 at some world must assign the conclusion a degree of truth of *at least* 0.5 at that same world in order for the argument to be valid. Let us call the resulting logic which combines the given semantics with this classical consequence relation Fuzzy Epistemic Logic (**FEL**).

Now we can clearly state the problem of probative omnipotence as it applies to a particular choice of logic, like **FEL**. The problem is that all worlds in $\boldsymbol{W}$ are maximally obedient with respect to the semantic rules listed above. Suppose an agent wants to assign some credence to the possibility that a proposition α takes on the truth-value x. To do so, they can make their probability measure such that $Pr(\alpha_x) > 0$, where $\alpha_x = \{w : \nu(\langle w, \alpha \rangle) = x\}$. However, ν doesn't assign values to propositions in a vacuum, so this agent will also end up assigning positive probability to events of the form β_y where β is some proposition which must take on value y given the semantics of **FEL** and the supposition that α has truth-value x. Thus, agents in any given context $\mathfrak{C}$ are depicted as omniscient insofar as they can never give credence to a proposition's having some truth-value without also giving credence to

---

[2] Consult Priest (2002, p. 307) for a paraconsistent logic which uses Łukasiewicz fuzzy semantics.

a network of truth-value assignments supplied by the semantics of **FEL**. A consequence of this omniscience is that, sometimes, when a proposition α is proven to a degree greater than 0.5 and the argument $\alpha \therefore \beta$ is valid, then β ends up being proven to a degree of at least 0.5 as well. In particular, when propositions are proven to degree 1, some of their consequences are proven to degree 1 as well. Agents are unable to prove $((P \to Q) \wedge \neg Q)$ to degree 1 without proving $\neg P$ to degree 1 also.[3] This is counterintuitive because, at least in principle, one could be shown that the proposition $((P \to Q) \wedge \neg Q)$ is true without being shown that $\neg P$ is true also. Indeed, *modus tollens* is a rule which one comes to learn by studying some logic or other, like **FEL**. It would be very surprising if, given any instance of this argument form, all agents are already able to appreciate the truth of its conclusion merely given the truth of the premise.

Any solution to this problem of probative omnipotence must tread a fine line. On the one hand, we want to avoid a situation where proofs that P inadvertently prove anything which is implied by P. On the other hand, we don't want proofs which are entirely impotent. Sometimes proofs are "indirect" in the sense that they assume some level of logical competence for their success. Take the example of my proving that Bolzano is in Italy by proving it isn't in Austria. In this case, I prove a proposition of the form P by showing that $\neg Q$ against a background where it is already presupposed that $(P \vee Q)$. Solving the problem of probative omnipotence shouldn't involve the *en masse* prohibition of these indirect sorts of proof. This is a deficit which Skipper and Bjerring (2020, p. 100) have identified in some proposed solutions to logical omniscience, saying that "logical anarchy not only eliminates logical omniscience, but also breeds logical incompetence."

For those who appreciate this gulf between the logically proficient and the logically anarchic, there is a tendency to place agents onto a sliding scale of logical competence. The idea is that we should expect agents to make fairly simple deductions, but that these expectations become unreasonable as the complexity of these deductions increases. Skipper and Bjerring (2020)

---

[3] To see why, suppose that the premise has been proven in some context set $\mathfrak{C}$ and $\nu(\langle w, ((P \to Q) \wedge \neg Q)\rangle) = 1$ for all w which are G-reachable from the context set(s) in $\mathfrak{C}$. By $(\wedge)$, this means that $\nu(\langle w, (P \to Q)\rangle) = \nu(\langle w, \neg Q\rangle) = 1$. By $(\to)$, this also means that $\nu(\langle w, P\rangle) = 0$, from which it follows by $(\neg)$ that $\nu(\langle w, \neg P\rangle) = 1$. Now, since $\neg P$ is true to degree 1 at all worlds which are G-reachable from the context set(s) in $\mathfrak{C}$, it has also been proven to degree 1.

have attempted to construct such a scale by interpreting complexity in terms of length of proof in an axiomatic system. A problem with this variety of approach, noted as early as Hintikka (1979, p. 377), is that complexity is then tied to a particular choice of axioms and inference rules. Take an argument which has a conjunction of 100 unique basic propositions as its premise and one of these basic propositions as its conclusion. In an axiomatic system which allows α or β to occur on line n of a proof if $(\alpha \wedge \beta)$ appears on line $k < n$, a proof of this argument's validity might be 100 lines in length—if, for instance, the basic proposition being derived is nested in 99 pairs of parentheses. However, in an axiomatic system which allows α_i (with $1 \leq i \leq m$) to occur on line n of a proof wherever there is a conjunction of the form $\alpha_1 \wedge ... \wedge \alpha_m$ (regardless of the distribution of parentheses) on some line $k < n$, then a proof of the argument's validity could be as little as 2 lines long. Thus, the very same deduction is either completely trivial or significantly complex depending on which rules of inference one permits in their axiomatic proof system.

Hintikka (1979), building on Rantala (1975), has an alternate solution to the problem of logical omniscience which invokes a scale of complexity that doesn't rely on particular axiomatisations of logic. There, he has recourse to an epistemic logic where worlds in $\boldsymbol{W}$ are represented by nonstandard models of first-order logic called 'urn models'.[4] These models are a subspecies of game-theoretic models, first devised in Hintikka (1979, p. 27—47), where valuation functions are replaced with games. One of the players in these games, called the Verifier, has the goal of producing a true sentence. The Verifier plays against an opponent, called Nature, whose goal is to produce a false sentence. At the beginning of each game, the players are supplied with three things: (i) a starting formula α_0, whose truth-value is being determined, (ii) a domain D of objects and (iii) an interpretation function $\boldsymbol{I}$ which takes names to objects in D (i.e. $\boldsymbol{I}(a) \in D$) and takes n-ary predicates to sets of n-tuples made from D (i.e. $\boldsymbol{I}(P^n) \subseteq D^n$). This way, before the game is even played, all atomic formulae are assigned a truth-value as follows: $P^n a_1 ... a_n$ is true if $\langle \boldsymbol{I}(a_1), ..., \boldsymbol{I}(a_n) \rangle \in \boldsymbol{I}(P^n)$ and false otherwise. Each round n of the game for α_0 is played as follows:

1. The players consider some proposition α_n
2. Depending on the main operator of α_n, either Nature or the Verifier is

---

[4] See French (2015) for a contemporary presentation of these urn models.

given some rule to select a formula α_{n+1}

3. The players begin round $n+1$

The game ends after round $m-1$ just in case α_m is an atomic formula. If α_m is true, then the Verifier wins. If α_m is false, then Nature wins. The initial formula α_0 is true on one of these game-theoretic models if and only if there is a game-winning strategy available to the Verifier. Here's an example of one of the rules which Hintikka (1979, p. 34) provides for the selection process in step 2 of each round n:

(G.E) If α_n is of the form $\exists x\beta$, then the Verifier chooses some object $o \in D$ and gives it some new name (unless it already has one). Round $n+1$ then begins with the consideration of $\alpha_{n+1} = \beta(b/x)$, where $\beta(b/x)$ is a formula exactly like β except all free occurrences of the variable x are replaced with occurrences of the new (or existing) name for o, b.

On one of these models, a simple quantified proposition like "There are red things" ($\exists xRx$) is true just in case $\boldsymbol{I}(R) \neq \varnothing$. This gives the Verifier a simple winning strategy, which can be completed in the very first round: start the game by picking a red thing from the domain and produce the true atomic proposition Rb.

Hintikka (1979, pp. 372) points out that "ordinary models ('logically possible worlds') can be thought of as a subset of the class of all urn models." In particular, they are urn models whose domain D doesn't change after each round. In these models, whenever the Verifier/Nature picks some object, it is then placed back in to the domain and nothing else is added nor removed. On the other hand, worlds which are *merely* epistemically possible ('impossible possible worlds') can be thought of as urn models whose domain changes after a certain number of rounds. In these models, the Verifier/Nature might not replace the objects they draw and things may be added or removed from the domain after the end of each round. An agent's logical acumen is then measured by the smallest number n such that they consider possible (or assign positive probability to) an impossible world whose domain is invariant until round n, after which the domain changes.

While these urn models are an interesting way to grade logical skill which doesn't rely on something as arbitrary as an axiomatisation of logic, they are unhelpful for the problem of probative omnipotence which threatens propositional logics like **FEL** because their complexity depends on them being models *of* first-order logic. Impossible worlds whose domain varies after zero rounds are still well-behaved concerning the semantics for connectives, so a logical simpleton who assigns positive probability to these *most* anarchic worlds is still depicted as competent enough to know he validity of *modus tollens* and all other propositionally valid argument forms.

Instead of extending this graded notion of logical competence to cover propositional logic, I reject the idea that solving the problem of probative omnipotence requires a story about how to measure degrees of logical proficiency. What matters is not whether an agent is competent enough to know complex entailments, but whether they know the entailments which they are required to know for the purposes of the task at hand. To play chess with someone, it is required that you and your opponent presuppose the permissibility of certain movements for the chess pieces but it isn't required that you be able to simulate at least five moves ahead in the game. Likewise, in any given conversation we might presuppose the validity of certain arguments, or presuppose that a connective has such-and-such semantic rules and use these presuppositions to form a common ground. My proof that Bolzano is in Italy which goes indirectly via a proof that it isn't in Austria only works because I take for granted the validity of disjunctive syllogism. I couldn't give the same proof in a conversation where the validity of this argument form is disputed—say, if I am trying to prove Bolzano's location to someone who subscribes to a relevant logic in which disjunctive syllogism is invalid. To represent these semantic laws as something which agents could in principle have presuppositions about, I will introduce the notion of a pseudo-valuation and use it to make a distinction between possible and impossible worlds in $\boldsymbol{W}$.

Let a pseudo-valuation φ be an unconstrained function from $\boldsymbol{W} \times \mathcal{L}$ to $[0, 1]$. Since these functions are unconstrained, they can violate the semantic laws of **FEL** and provide any distribution of truth-values amongst sentences of $\mathcal{L}$ at a given world. By replacing valuations with pseudo-valuations in contexts, we get a natural distinction between possible and impossible worlds in $\boldsymbol{W}$. Let $\boldsymbol{W}_{\mathbf{FEL}}, \boldsymbol{W}_{\overline{\mathbf{FEL}}} \subseteq \boldsymbol{W}$ be defined as follows:

> $\boldsymbol{W}_{\mathbf{FEL}} = \{w \in \boldsymbol{W} : \varphi(\langle w, \alpha\rangle) = \nu(\langle w, \alpha\rangle)$ for some **FEL** valuation ν and all $\alpha \in \mathcal{L}\}$
>
> $\boldsymbol{W}_{\overline{\mathbf{FEL}}} = \{w \in \boldsymbol{W} : \varphi(\langle w, \alpha\rangle) = x$ and $\varphi(\langle w, \beta\rangle) = y$ for some

> $\alpha, \beta \in \mathcal{L}$ and there is no **FEL** valuation ν such that $\nu(\langle w, \alpha\rangle) = x$ and $\nu(\langle w, \beta\rangle) = y\}$

Thus, relative to some set of worlds $\boldsymbol{W}$ and a pseudo-valuation φ, a world which is logically possible (with respect to **FEL**) is one which agrees with some valuation of **FEL** concerning all sentences in the language $\mathcal{L}$ and an impossible world is one which disagrees with all valuations of **FEL** concerning the truth values of at least one pair of sentences in the language $\mathcal{L}$. An agent i believes the semantic laws of **FEL** just in case their context set $\boldsymbol{C}_i$ is contained within the set $\boldsymbol{W}_{\mathbf{FEL}}$, but the problem of logical omniscience shows us that it is unreasonable to expect that all agents will always be able to help themselves from keeping worlds from $\boldsymbol{W}_{\overline{\mathbf{FEL}}}$ outside their context set, which is exactly what we were assuming when using a sample space $\boldsymbol{W}$ which contains nothing but logically possible worlds.

So now we can represent agents who believe or fail to believe whole logics like **FEL**. However, the main benefit of adding pseudo-valuations to contexts is that we can now show agents who believe or fail to believe *parts of* these logics too. For instance, an agent might believe that $(A \wedge B) \therefore B$ is valid by having a context set $\boldsymbol{C}_i$ which is such that $\boldsymbol{C}_i \cap (A \wedge B)_n \cap B_m = \varnothing$ for all $n > 0.5$ and $m < 0.5$, but fail to believe that $((P \to Q) \wedge \neg Q) \therefore \neg P$ is valid by having a context set $\boldsymbol{C}_i$ where $\boldsymbol{C}_i \cap ((P \to Q) \wedge \neg Q)_n \cap \neg P_m \neq \varnothing$ for at least one $n > 0.5$ and one $m < 0.5$. Additionally, an agent might believe or disbelieve particular semantic rules of **FEL** too. For instance, an agent i doesn't believe in the conjunction rule from **FEL** if and only if there is some $w \in \boldsymbol{C}_i$ where $\varphi(\langle w, (\alpha \wedge \beta)\rangle) \neq \min\{\varphi(\langle w, \alpha\rangle), \varphi(\langle w, \beta\rangle)\}$ for some $\alpha, \beta \in \mathcal{L}$. By adding failures of these logical laws into the space of possibilities which agents form beliefs about, we get a non-trivial model according which agents can be ascribed beliefs and presuppositions about these laws.

Those with a keen nose for sniffing out inconsistency might be picking up a strong whiff at this point. If contexts are to be modelled using these pseudo-valuations, doesn't this mean that all semantic considerations, including the considerations that determine when $\mathbf{B}_i\alpha$ is true, go out the window? That is, since $\varphi(\langle w, \mathbf{B}_i\alpha\rangle)$ can be any value whatsoever from $[0, 1]$, how could there any connection between what worlds are in an agent's context set and what they believe? The important thing to remember here is that $\varphi(\langle w, \mathbf{B}_i\alpha\rangle)$ can take on any value when and *only* when $w \in \boldsymbol{W}_{\overline{\mathbf{FEL}}}$, otherwise we must look at the worlds which agent i deems possible at w in order to know whether $\mathbf{B}_i\alpha$ is true. In particular, if we help ourselves to the relatively uncontroversial assumption that what's actual is possible,[5] then @ will always be a member of $\boldsymbol{W}_{\mathbf{FEL}}$. So, all truth values of the form $\varphi(\langle @, \alpha\rangle)$ are beholden to the semantics of **FEL** and context sets like $\boldsymbol{C}_i$ do constrain the truth of propositions like $\mathbf{B}_i\alpha$ at the actual world.

Given this framework where agents can form beliefs about logical laws, we can now explicate what it takes for an indirect proof to be effective. Roughly speaking, an indirect proof only works for agents who presuppose the appropriate logical laws and furthermore hold that *had things been different*, these laws would still hold. A pair of examples will make the need for this counterfactual proviso clear. Let us return to my indirect proof of Bolzano's location. In that case, I attempt to prove I (that Bolzano is in Italy) against a background where it is presupposed that $(I \vee A)$ (Bolzano is either in Italy or Austria) by showing that $\neg A$. This proof won't work if I perform it for Relevance Bill, who actively entertains the possibility of $(I \vee A)$ and $\neg A$ both being true while I is false. In this case, Relevance Bill's context set contains some world $w \in \boldsymbol{W}_{\overline{\mathbf{FEL}}}$, according to which $\varphi(\langle w, (I \vee A)\rangle) = \varphi(\langle w, \neg A\rangle) = 1$ but $\varphi(\langle w, I\rangle) = 0$.[6] This way, even if my proof manages to convince Bill he should eliminate (assign zero probability) to all worlds where $\neg A$ is not true, Bill could remain in doubt as to whether I is true, so I won't succeed in adding this claim to Bill's beliefs, let alone the common ground. You might think that this problem goes away when I perform the same proof to someone who presupposes **FEL**'s semantic rules for $\neg$ and $\vee$, but even then, there are cases where the indirect proof still fails. To see why, suppose instead that I show that $\neg A$ to Łukasiewicz Sue, who fully presupposes **FEL**—all worlds which are G-reachable from Sue's context set are members of $\boldsymbol{W}_{\mathbf{FEL}}$. Even if we both presuppose $(I \vee A)$, my proof that $\neg A$ might lead Sue to abandon

---

[5] Priest (1997, p. 581) rejects this assumption. He suggests that there is no *a priori* reason for believing that contradictions are false and thus that the actual world could, for all we know, be an impossible one. However, I think someone like Priest can maintain that actuality is possible *and* allow for purported contradictions of **FEL** to be true by adopting a different logic according to which they aren't contradictions and nominating that logic as a guide to possibility.

[6] When Bill balks at the validity of $(I \vee A), \neg A \therefore I$ this way, it could either be due to Bill's skepticism about the semantic rule for $\vee$ or his skepticism about the semantic rule for $\neg$ depending on the truth-value of A at this impossible world in Bill's context set.

some of her presuppositions in the semantic rules of **FEL**. This will happen if, for instance, there are no worlds in her context set at which $\neg A$ is true and her selection function σ_s picks some set outside of $\boldsymbol{W}_{\mathbf{FEL}}$ when providing the closest $\neg A$-worlds. For an indirect proof to work for some agent i, that agent needs to both presuppose the semantic rules which that indirect proof exploits *and* possess a selection function σ_i which respects these rules; never taking a pair of the form (w, α) to a set which contains worlds where they fail to hold. This is the sense in which agents must not only believe the relevant logical laws, but also believe that they would hold had things been otherwise relative to their current beliefs. Let us call this stronger epistemic state *commitment*, and say that an agent i is committed to some semantic rule R if and only if (i) there is no world w which is G-reachable from $\boldsymbol{C}_i$ at which R fails to hold and (ii) for all $w \in \boldsymbol{W}$ and $\alpha \in \mathcal{L}$, $\sigma_i(w, \alpha) \subseteq \boldsymbol{R}$, where $\boldsymbol{R} \subseteq \boldsymbol{W}$ is the set of worlds where the rule R holds.

This notion of commitment is what ultimately offers a solution to the problem of probative omnipotence. As the foregoing examples show, indirect proofs will only ever go as far as the logical commitments of interlocutors allow. I am not at risk of incidentally proving that Bolzano is in Italy to Relevance Bill when I show him that Bolzano is not in Austria. More generally, I am not at risk of incidentally proving *all* consequences of a given proof in contexts where people aren't committed to the semantic rules which guarantee these consequences. It should also be clear at this point why nothing in this solution depends on the choice of **FEL** in particular. The pseudo-valuations introduced allow us to represent varying logical commitments across agents, regardless of whether **FEL** correctly demarcates which of the worlds in $\boldsymbol{W}$ are genuinely possible. The problem of probative omnipotence emerges in the naïve theory of proof because these agent-relative commitments aren't respected. All agents were depicted as logically omniscient relative to some background logical framework, leaving no room for disagreement with or ignorance about its semantic rules.

One might also worry about the nature of the impossible worlds which pseudo-valuations introduce. If $\boldsymbol{W}$ is meant to be a compendium of ways things could've turned out, then what makes the worlds in $\boldsymbol{W}_{\overline{\mathbf{FEL}}}$ impossible—that is, what stops them in particular from being ways things could've turned out? The best way to circumvent this anxiety is to think of worlds in $\boldsymbol{W}$ not as alternative states of affairs, but as placeholders for what might be alternative states of affairs. When an agent weighs up the probabilities of events $E \subseteq \wp(\boldsymbol{W})$, we should think of points in $\boldsymbol{W}$ as representations or descriptions which are specific enough to determine for each event E whether or not it holds. Some of these points correspond to possibilities (ways things could have been) whereas others are impossible insofar as they fail to correspond to a way things could've been. For the agent who is considering these points in $\boldsymbol{W}$, the task of separating the possible from the impossible proves extremely difficult: indeed, they will often assign positive probability to events like $\{w\}$ where w is impossible because of a failure to appreciate some metaphysical or logical fact which rules out $\{w\}$ occurring. To see how this happens, think of $\boldsymbol{W}$ as a manila folder filled with glossy photographs depicting a backyard. Without knowing which, the agent is informed that one of the photographs is a picture *of* the backyard in the house on King Street they find themselves in. The envelope contains two types of photograph: (i) pictures which were taken of *some* backyard in King Street and (ii) superficially convincing A.I.-generated images of a backyard in King Street. The agent then goes through and discards any photographs which, for some reason or other, they believe couldn't possibly be photographs of the backyard they are in.

Whether the agent thinks there might be a Jacaranda in the backyard depends on whether they have left any photographs in the folder which feature Jacarandas. On the other hand, whether it *is* possible that there is a Jacaranda in the backyard depends on whether one of the non-A.I.-generated pictures in the original envelope features a Jacaranda. The agent is logically or modally competent insofar as they are able to avoid discarding the genuine photographs. Some of the A.I.-generated photos will be easy to spot, with random geometric patterns in the grass and floating blobs of colour. However, slightly more insidious impossibilities will require diligent observation. An agent might *think* there could be a Jacaranda in the backyard simply because they fail to notice that all photos of Jacarandas are depicted such as to require small peripheral branches which are suspended in mid-air.

By including impossible worlds in $\boldsymbol{W}$, there is no risk that we will somehow lose a distinction between what's possible and impossible. $\boldsymbol{W}$ contains a number of representations of ways things could be—different assignments of truth-values to all propositions—but it might contain *more* of these representations than there are ways things could be. If this is the case, then necessarily some of the worlds in $\boldsymbol{W}$ will have to be fakes—i.e., impossible worlds. However, whether one of these representations is fake depends on metaphysical and logical facts. Just like you cannot make it true that there might be a Jacaranda in the backyard by failing to discard an A.I.-generated photo, you cannot make a world possible just by assigning it some positive probability.

2.1.2 Standards of Proof

Now that we have some idea of how proofs can come in degrees of strength, we can consider the question of whether a proof is strong enough to 'settle the question' in a given context. Proofs come in varying degrees of strength, but in certain activities we might have some criteria for when a proof is strong enough to take a proposition as true for current intents and purposes. For instance, you might need to prove your identity in order to renew a driver's license. It would be unreasonable for the service centre to expect that you are able to prove this to degree 1, since that would presumably require mountains of evidence that no-one has the time to rummage through. So, for the purposes of renewing your license, they settle for proofs of identity which are strong enough to meet some threshold. We may not be able to give a precise numerical characterisation of this threshold, but intuitively it is whichever threshold is met by producing a passport or a few pieces of mail with your name and address on them.

Nowhere is this phenomena more prevalent than in matters of law, where practitioners of jurisprudence have explicitly associated different kinds of trial with different so-called *standards of proof.* Leibniz (1996[1765]; Book IV, Chapter XVI, §9) lays out a cornucopia of these standards for legal evidence, such as *complete proofs*, *more than full proofs*, *presumptions*, *more than half full proofs*, *less than half full proofs* and so on. Leibniz (ibid.) also alleges that in criminal trials of the time there were standards in place which determine when evidence is strong enough to warrant torture of the accused as opposed to when evidence is merely strong enough to justify displaying instruments of torture without actually using them. In our slightly more civilised civil and criminal courts today, there are two important standards of proof which are commonly invoked. In civil disputes, a plaintiff must prove their claim *on the balance of probabilities* in order to succeed. In criminal disputes, on the other hand, the accused's guilt must be proven *beyond reasonable doubt* in order for a conviction to be made. In this section, I will consider whether these standards of proof can be interpreted in terms of degrees of strength and then explore the relationship between standards of proof in the legal setting and standards of proof in a more generalised conversational setting.

Legal Standards

The main problems facing a reduction of legal standards of proof to degrees of proof as I've defined them come from a series of challenges which Cohen (1977) identified for those who interpret the probabilities in a phrase like 'on the balance of probabilities' as standard probability measures obeying the Kolmogorov axioms. These objections from Cohen (1977) can be broken up into two main varieties: concerns about how the probability of compound propositions relates to the probabilities of their components and concerns about how interpretations of probability square with interpretations of legal doctrine. We'll consider each type of problem in order and consider whether they can be translated into objections about a comparable reduction using strength of proof in place of probability.

As for the first variety of problem, Cohen's (1977, pp. 58-67) first complaint is about conjunction. In a civil case, the plaintiff typically needs to defend a multitude of small claims in order to defend their case overall. For instance, if a plaintiff wishes to claim that a soft-drink manufacturer breached their duty of care by allowing a snail to be bottled in her ginger beer, she might need to demonstrate both that (N) negligence from the manufacturer is what caused the snail to be bottled and (D) that drinking from the bottle with the entombed snail caused her to be ill. Assuming, momentarily, that N and D aren't vague propositions so that $Pr(N)$ and $Pr(D)$ are identical to the expected truth-values of the propositions N and D respectively, the problem about conjunction is that the probabilities of minor claims in a civil case can be sufficient for proof whereas the probability of their conjunction falls short. To see why, notice that N and D are probabilistically independent events, such that $Pr(N|D) = Pr(N)$ and $Pr(D|N) = Pr(D)$. This means that the probability of their intersection $Pr(N \cap D)$ (and by extension, the expected truth-value of $N \wedge D$) is the product $Pr(N) \cdot Pr(D)$ of their respective probabilities. If we take an event to be proven 'on the balance of probabilities' when its probability is greater than 0.5, as those who reduce standards of proof to probabilistic thresholds are wont to do (Hedden and Colyvan, 2019, p. 457, Urbaniak and Di Bello, 2021), then the plaintiff will be able to prove N and D to the fact finder's satisfaction without proving their conjunction. This will be the case if, for instance, $Pr(N) = Pr(D) = 0.6$, whereupon $Pr(N \cap D) = 0.36$. Cohen (1977, p. 60) and more recently Allen (2020, p. 123) note that this problem compounds in accordance with the amount of minor claims required for a plaintiff to make their overall case.

Hedden and Colyvan (2019, p. 458) downplay the significance of Cohen's conjunction problem, suggesting that instead of being an obstacle, the conjunction problem offers an important prescriptive moral for the legal system: its "concern should be with whether the case as a whole, the entire allegation, meets the relevant standard of proof." In their minds, the plaintiff should be required to establish $N \wedge D$ on the balance of probabilities, regardless of how strongly N and D must be confirmed to ensure this result. However, by their own admission, this response does little to help descriptive legal theorists, who wish to model standards of proof *in situ*, as they are actually applied. Furthermore, Allen (2020, p. 123) suggests that this subtle shift in the target of legal evidence would lead to drastic changes in legal practice, including situations where the exact same fact would require different degrees of evidence depending on how many additional facts need to be established as part of the trial.

Cohen's (1977, pp. 74-81) other concern when it comes to the probabilities of compound propositions has to do with negation. Relative to some sample space $\boldsymbol{W}$, the probability of an event is always equal to one minus the probability of its complement, i.e., for all $X, \overline{X} \subseteq \boldsymbol{W}, Pr(\overline{X}) = 1 - Pr(X)$. Take X to be the proposition that some plaintiff is trying to establish. If we assume, once again, that X is not vague, this means that $E(\neg X) = 1 - E(X)$.[7] Under a framework where the strength of a case is measured by probabilistic mass, this complementation principle is problematic insofar as jurors are forced to give an amount of merit to the plaintiff (or defendant) merely on the basis of how much merit they have given to the defendant (or plaintiff). Intuitively, however, the strength of a case does not need to vary inversely with the strength of its contrary. Indeed, Cohen (1977, p. 80) points out that lawyers sometimes say such things as, "the defendant's case is equally good on the facts in both of two similar lawsuits, while the plaintiff's case is better in one than the other." Furthermore, Haack (2014, p. 62) points

---

[7] Recall that probability measures range over events, which are subsets of $\boldsymbol{W}$ (like $\overline{X}$), whereas degrees of belief (or expectations) apply to propositions (like $\neg X$).

out that this principle fails to reconcile with the intuition that a case for X and a case for $\neg X$ might both have zero merit if, for instance, there is no evidence whatsoever for either claim or the available evidence is sufficiently weak. It's clear from these considerations that whatever quantity is used to evaluate the strength of legal evidence must not obey the same kind of complementation principle which applies to the probabilities of propositions and their negations.

Cohen's other variety of issues, besides these two involving compound probabilities, have to do with the awkwardness of providing an interpretation of legal standards of proof using probability measures. For instance, Cohen's (1977, pp. 82-3) issue with using probabilities to represent the standard of proof in criminal cases is that whether it is reasonable to doubt some claim X depends not on the magnitude of X's probability but rather whether "there is a particular, specifiable reason for doubting it." A very similar objection has also been made recently in Moss (2023), who offers a framework where legal proof requires knowledge of the claim being proven. In this account, whether a doubt in X is reasonable depends on whether there is some uneliminated relevant possibility at which X fails to hold. Moss (2023, p. 178) argues that, in general, whether these epistemic possibilities are relevant has no correlation with the probability which they are assigned. Consider a prison riot, where $n-1$ out of n prisoners start attacking guards and the remaining prisoner tries to stop the riot. While increasing the number n decreases the probability of a randomly chosen prisoner being innocent, this has no effect on the reasonableness of doubting that this randomly chosen prisoner is guilty. Whatever quantity is measured by probability appears to have no bearing on the quantity which matters for legal evidence.

The disparity between these quantities is also made clear by considering how received interpretations of probability could be used to provide accounts of legal evidence. A frequentist interpretation, which takes $Pr(P|E)$ to be $\frac{n}{m}$ just in case there are n cases where P across the total number (m) of cases where E, is a non-starter according to Cohen (1977, pp. 87-88) because judges explicitly deny that such frequencies are sufficient for legal proof. The randomly selected prisoner is never found to be guilty simply in virtue of belonging to a group with a large proportion of guilty members. What about a more subjective interpretation of probability, such as the one provided in §1.3 above? According to this view, an agent's probability in an event P is constituted by the strength of their tendency to behave as if P. Following de Finetti (2017[1970], p. 74), this strength can be cashed out in terms of the value p according to which you would be willing to enter any bet whose payoff is $\$S(1-p)$ if P is true and $\$S(0-p)$ if P is false, where S is any number, positive or negative. The p which an agent chooses is called their *betting quotient* for P. Cohen's (1977, p. 89) issue with interpreting legal evidence in terms of these betting quotients comes from the absurdity of gambling on a proposition whose truth is up for dispute in a way that couldn't be validated by some forthcoming observation. We can bet on whether a particular candidate will be elected, whether a team will win the soccer match and whether a horse will finish first only because in each case there is some anticipated future observation which will decide how the bet is to be paid out—be it the televised election results, the performance of each team in the soccer match or Phar Lap bounding ahead to cross the finish line. Members of a jury cannot gamble on the result of a trial, Cohen insists, because they are in a position where all of the evidence which could possibly be used for settling this bet is already public knowledge to them. Thus, Cohen doubts that legal evidence can be understood using these subjective probabilities either.

Now, we will take stock and consider whether these arguments against a probabilistic interpretation of legal evidence find the same purchase as arguments against an analysis which treats legal standards as degrees of proof. First, let's consider the problems for compound propositions. Do the problems of conjunction and negation remain when we talk of degrees of proof instead of degrees of probability? It depends on the nature of the proof at hand. If two claims P and Q are proven to degree n and m respectively by a probabilistic proof (i.e., one where the induced constraint $(q, \mathcal{X})$ is not an intensive qualification to an intermediate degree) and furthermore there is no G-reachable world from the context set(s) where some $i \in G$ takes P and Q to be probabilistically dependent, then this proof will also be an $(n \times m)$-strength proof of the conjunction $P \wedge Q$—that is, assuming all G-reachable worlds are contained within $\boldsymbol{W}_{\mathbf{FEL}}$. Similarly, if a probabilistic n-strength proof is given for P, then $\neg P$ is also proven to degree $1-n$. The result is that, unsurprisingly, probabilistic proofs inherit the same problems facing a probabilistic interpretation of legal proof.

However, if we don't assume that legal claims are precise, this opens up room for an *intensive* interpretation of standards of proof. According to this view, jurors are asked to make a judgement on how *true* they take a claim to be given the admitted evidence: in civil trials this evidence will be sufficient if it makes the plaintiff's claim more true than false (i.e., at least true to degree 0.5) and in criminal trials it will be sufficient if it makes the defendant's claim to innocence true to some very high degree. Putting to one side the plausibility of this style of interpretation, which I'll address in a moment, its immediate advantage over the probabilistic account concerns compound propositions. The fact that juries don't take a conjunction to be less supported than either of its conjuncts can be explained by treating juries as if they are committed to a Łukasiewicz style conjunction rule, where $\varphi(\langle w, \alpha \wedge \beta\rangle) = \min\{\varphi(\langle w, \alpha\rangle), \varphi(\langle w, \beta\rangle)\}$. Relative to this semantic constraint, an intensive proof of a conjunction representing a plaintiff's overall case will be as strong as its weakest conjunct, not as strong as the multiplicative product of these conjuncts. This way, there would be no difficulty above and beyond giving a sufficient proof of each component in an overall case and no prima facie difficulty is imposed by adding further components to a given case.

In fact, these intensive proofs not only resolve the conjunction issue but do so in a way which is at least in one respect superior to Cohen's (1977) proposed solution, which interprets standards of proof in terms of degrees of *inductive support*. As Cohen (1977, p. 266) points out, these degrees of inductive support are only assigned to conjunctions whose conjuncts are about the same subject matter. This is due to the fact that a hypothesis' inductive support is graded by the most complex experiment which confirms that hypothesis. If two hypotheses require different courses of experimentation, then their degrees of inductive support may be incommensurable insofar as it is absurd to consider the most complex experiment which confirms their conjunction. No such incommensurability threatens the intensive interpretation of legal proof: if two propositions are true to a particular degree then, regardless of what they are about, so is their conjunction. This is not a moot point either, since it is often true that a plaintiff's overall case will contain components with different subject matters. Consider the plaintiff who wishes to prove that the soft-drink manufacturer breached their duty of care by allowing a snail to be bottled. The kind of experiments and evidence needed to show the negligence of the manufacturer are clearly unrelated to those which would be needed to show that a snail contamination can cause illness, so the overall claim $(N \wedge D)$ would end up not being assigned a degree of inductive support. On the intensive account, however, the degree to which $(N \wedge D)$ is proven will simply be the truth-value of N or D, depending on which is lower. Cohen can't escape this problem by suggesting that the standard of inductive support required for the overall case will be met when the components N and D have high enough inductive support on their own, because he has already blockaded this exit for probabilistic reductions of legal proof. Cohen (1977, p. 66) suggests that ignoring the overall case in this fashion leads to an unacceptable disconnect between the judgements of the formal model of legal proof and the carrying out of justice. There is something patently odd about maintaining that inductive support is the sole measure of legal proof while also allowing $(N \wedge D)$ to be legally proven when it is assigned no degree of inductive support whatsoever.

When it comes to the issues surrounding negation, the intensive interpretation of legal proof also fares well. The complementation principle for negation can be avoided by having a jury which is committed to a bivalent rule for negation where $\varphi(\langle w, \neg\alpha \rangle) = 0$ when $\varphi(\langle w, \alpha \rangle) > 0$ and $\varphi(\langle w, \neg\alpha \rangle) = 1$ otherwise.[8] The resulting framework is in effect quite similar to Cohen's (1977, p. 177), where the degree of inductive support for a negated hypothesis is zero whenever the hypothesis has some positive degree of inductive support. However, unlike Cohen's account, this bivalent negation rule also requires that $\neg\alpha$ be true to the highest degree whenever α is true to degree zero. This difference appears to create a roadblock given that I want to maintain, following Haack (2014), that there may be no degree of legal proof whatsoever for a claim and its negation. However, this problem is resolved if one keeps in mind a distinction between (i) evidence that some claim P is true to degree zero and (ii) a lack of convincing evidence about P. If a jury is presented with evidence that leads them to believe P is completely false, without any doubt in their minds, then they are well within their proper functioning to infer that $\neg P$ is true to the highest degree and this would constitute a finding against P in the trial. However, if the evidence presented to that jury is weak enough to leave them uncertain as to whether P is completely false, such that $Pr(P_0) < 1$, then they will not be certain that $\neg P$ is true to degree 1. Keeping this possible distinction in mind, the intensive account of legal proof in combination with this bivalent rule for negation does succeed at avoiding the same problems as a probabilistic reduction.

Despite this intensive account's success with compound propositions, this victory would ring completely hollow if it could not also be used to provide

---

[8] This is the negation rule from the Gödel and product fuzzy logics, which is called *strict* negation (See Běhounek, Cintula and Hájek, 2011, p. 8).

a reasonable interpretation of legal proof. However, I'll now argue that it can. Expanding on this interpretation, it takes the strength of a legal proof with respect to some claim P to be the degree of truth which a jury takes P to have once that proof is performed (or evidence delivered). In contextual terms, we will say that a legal proof's degree of strength with respect to the claim P is n wherever a jury presupposes that P is true to degree n after that proof is delivered and undefined otherwise. Although this interpretation still depends on the subjective opinions of the jury, it doesn't face the same challenges that Cohen outlined for interpretations which treat jurors as gamblers. One can consistently possess all possible evidence concerning P while also believing P to an intermediate degree, provided that they are certain P has an intermediate degree of truth. This way, there is nothing absurd about insisting a jury be more or less confident that P even though there is no forthcoming observation which will settle whether P.

Another benefit of the intensive account is that degrees of truth are a quantity which *do* appear to track reasonableness of doubt, leading to a much more suitable interpretation of the standard in criminal trials. Imagine an example which is exactly like the prisoner riot case, except that instead of arbitrarily decreasing the probability of a given prisoner being innocent we gradually decrease the degree to which they are innocent. What I have in mind here is a sorites series of possible behaviours of the prisoner, ranging from clearly opposing the riot to clearly instigating it. Suppose a jury is presented with a series of conclusive pieces of evidence, each of which points to the prisoner acting in the way specified by the corresponding behaviour in the sorites series. At first, it will be incredibly reasonable to doubt the prisoner's guilt, citing things such as the prisoner's insistence on signalling other guards, but as one progresses through this series the reasonableness of these doubts will wane. Evidence of the prisoner's guilt will be strong enough for conviction when it makes it certain that the prisoner's behaviour is one which leaves no room for a clear and specifiable reason for doubting their guilt—e.g., when it points to a scenario where the prisoner is the first one to attack one of the prison guards. This is how an intensive account can interpret the criminal standard of 'beyond reasonable doubt'.

One might be quite hesitant to accept the intensive account because, intuitively, it appears as if guilt (or innocence) of the accused is an all-or-nothing affair. It is the kind of thing which is either completely true or completely false. The prisoner on trial was either part of the riot or one of the riot's detractors. I think there are two sources of misunderstanding in such an objection to the intensive view. First of all, it *is* the case that trials will only ever end in one of two ways—ignoring mistrials, the defendant will either be served a verdict of guilt or innocence. While there might not be bivalence in the deliberation stage, there is no denying that a different kind of bivalence occurs at the end of a trial. Secondly, what is required to come in degrees according to the intensive account isn't the actual degree of truth of a proposition such as, 'this accused prisoner was rioting'. Rather, it is the degree of truth which a jury *takes* such a proposition to have given the admitted evidence. The jury is tasked with constructing what is, in their minds, a complete account of what happened in a given case using nothing more than the provided evidence. It is this mental construction according to which claims of the accused's guilt are required to be true to a greater or lesser degree. This framework is made plausible by the fact that jurors are explicitly instructed *not* to procure additional evidence. In the phrasing of the New South Wales Criminal Trial Courts Bench Book (2002, p. 105),

> A criminal trial is [...] not a search for the truth. Therefore neither the judge nor the jury has any right to make investigations or inquiries of any kind outside the courtroom and independent of the parties. The verdict must be based only upon assessment of the evidence produced by the parties.

The degrees of support which claims gain through the opinion of jurors are not supposed to be based on knowledge of how strong the claim is as a matter of fact. If this were the case, there would be no issue presented by triers of fact conducting their own investigations. Rather, it is up to these triers of fact to determine how strong a case is on the basis of the narrative they construct using nothing beyond the admitted evidence. The intensive interpretation respects the structure of this kind of inquiry by treating jurors as certain of a particular reconstruction of the events, but nevertheless uncertain of whether guilt can be truly ascribed relative to this reconstruction.

These considerations show that legal standards of proof can be carved out quite respectably using partial proofs, so long as their effect on presupposed degrees of truth is emphasised. This model of what goes on in the minds of triers of fact succeeds because it treats them not as weighing up how probable a series of claims are, but as weighing up how strong these claims are according to the state of affairs depicted by the evidence before them.

Conversational Standards

Moving from jurisprudence back to conversational pragmatics, one might also wonder whether standards of assertion can be modelled using degrees of proof. These are standards which establish the degree to which a proposition must be supported in order for it to be taken as true for the purposes of a given conversation. This way, when a speaker attempts to assert that P, they might only need to prove that P to some degree short of the maximum before their assertion is accommodated. Importantly, these thresholds are not ways of re-describing what counts as accommodation. Having a greater than n-strength presupposition in P is not the same thing as accommodating P in a context where the standard of assertion is n. Rather, it means that when a speaker asserts P and the speakers achieve a degree of presupposition in P which is n or greater, they will accommodate the assertion that P and adopt a full-strength presupposition that P thereafter. These standards of assertion are also called "standards of precision" by Lewis (1979, p. 352), which he posits as an additional variable tracked by the conversational score. This is supposed to explain why some contexts are relaxed enough to accommodate assertions like, "Harry is bald" even though his head has a few patches of hair whereas in others we might call someone up on such assertions, retorting with, "Harry's no skinhead, look at his flowing locks!"

A preliminary point I wish to make about these standards of assertion is that I do not, pace Cohen (1977, p. 50), believe that whatever structure standards of proof have in legal contexts ought to match the standards of assertion found in conversational contexts. Standards of proof are technical notions which are entrenched in legal practice, whereas standards of assertion are rough and ready, non-institutional rules that a group of speakers can decide to adopt or abandon more or less at will. Speakers might be prepared to admit a fairly low standard of assertion for grave matters, whereas jurors are not permitted to do so with legal standards. This is in spite of the fact that jurors have historically been informed that they should deliberate on cases as they would on the serious matters of their everyday life. Cohen (1977, p. 50) cites the following instruction as an example of this practice, given to a jury by Baron Pollock in the late 19th century:

> If the conclusion to which you are conduced be that there is a degree of certainty in the case which you would act upon in your own grave and important concerns, that is the degree of certainty

> which the law requires and which will justify you in returning a verdict of guilty.

However, despite this instruction, it's just not true that there is a unique degree of certainty which is sufficient for action in people's important affairs. In their everyday lives, people can and do establish any standard they wish. An angry mob with pitchforks might be prepared to take justice into their own hands when there are merely rumours of the town mayor's corruption and a Cartesian skeptic might only leave their heated room when they are convinced in the existence of the outside world beyond *any* doubt. Judges noticing this problem have been critical of instructions which urge jurors to use their 'everyday' standards in courtrooms. In *Regina v. Hepworth and Fearnley*,[9] Lord Goddard made this point while expressing concerns about the notion of a reasonable doubt:[10]

> to tell [the jury] that a reasonable doubt is such a doubt as to cause them to hesitate in their own affairs never seems to me to convey any particular standard; one member of the jury might say he would hesitate over something and another might say that that would cause him not to hesitate at all.

These concerns suggest that there are conventions in law which go beyond the comparatively un-regimented conventions of our daily lives, so we shouldn't feel compelled to give the same analysis of legal and conversational standards. Indeed, in this section, I will develop a theory of conversational standards by exploring various ways in which they depart from legal standards.

Taking this difference seriously, the first thing which sets standards of assertion apart from legal standards is their flexibility with respect to the kinds of evidence they require. Unlike legal standards, standards of assertion can sometimes be met by presenting mere statistical evidence or probabilistic support. Consider a group that is planning what they will do on the weekend. One of the speakers asserts, "It's going to rain on Sunday, so we should go to the movies instead of the beach." Such assertions are often accommodated but rarely are they accommodated from a position of absolute certainty about

---

[9] *Regina v. Hepworth and Fearnley*, 1955, Volume 3, *Weekly Law Reports*, p. 333.

[10] See Heydon (2024, p. 392) for a discussion of this opinion.

whether it will rain on Sunday. Usually this accommodation will be made on the basis of a weather forecast which ascribes high probability to the event of rain on Sunday. Alternatively, such conversational standards may be met by intensive evidence. Consider the assertion that it is raining (R), made in a context where there is an observable, heavy drizzle outside. In a context like this, speakers might be prepared to accommodate R based on their presupposition that the degree of truth possessed by R is very high. The fact that these standards of assertion can be met using diverse kinds of evidence is most naturally explained by a framework which treats standards of assertion as conditional offers of accommodation, which only apply to an assertion when speakers are sufficiently confident in its content.

How this confidence is garnered isn't generally at issue, only whether it *is* garnered. Of course, specific conversations might establish rules about what kinds of confidence warrant an assertion's accommodation. One can think of legal proceedings as a specific kind of conversation where probability-based degrees of confidence are insufficient on their own to warrant the accommodation of claims. Conversely, we can imagine scenarios where speakers only accommodate P when they have a sufficiently strong probabilistic confidence that P. Suppose a pair of gamblers are in the habit of rounding probabilities up or down, so that they fully accommodate α whenever they take it that $Pr(\alpha) \geq 0.5$ and fully accommodate $\neg\alpha$ whenever they take it that $Pr(\alpha) < 0.5$. The gamblers then use these rounded probabilities when deciding what bets they are willing to take part in. Showing these gamblers that α is true to a high but intermediate degree would lead them to round down their confidence in α and accommodate $\neg\alpha$, no matter how high that degree of truth is. This is because $Pr(\alpha)$ is zero whenever $Pr(\alpha_x) = 1$ and $x < 1$. No degree of intensive evidence short of the maximum will satisfy the standard of assertion for α set by this pair of gamblers. Perhaps this is because they will treat all such evidence as signalling that α does not hold for the purposes of a bet; if Phar Lap has somehow won the race to an intermediate degree then Phar Lap hasn't won the race at all. These kinds of conversation, where speakers favour a particular kind of evidence for their standards of assertion, should be seen as exceptions to a general relationship between these standards and degrees of confidence.

Taking this general relationship between standards of assertion and degrees of confidence to hold, we can construct a straightforward interpretation of standards of assertion in terms of degrees of proof. Let us say that whenever there is a strength n standard of assertion in place, an assertion that P

will be accommodated just in case speakers have an m-strength presupposition that P, where $m \geq n$. This way, a speaker who asserts P can ensure that P is taken for granted by performing an n-or-greater-strength proof that P. As a dynamic rule for conversations, this account of conversational standards will be stated as follows:

Standard of Assertion If at time t a speaker asserts that P in a context $\mathfrak{C}$ with an n-strength standard of assertion in place, then just after t—provided that the group G has a degree of presupposition of at least n that P— the context will become $\mathfrak{C}'$, which is exactly like $\mathfrak{C}$ except that $\mathbb{P}$ is replaced with $\mathbb{P}(q, \mathcal{X})$ where $\mathcal{X} = \{P, \overline{P}\}$ and $q(P) = 1$.

While I think this framework is close to the mark when it comes to representing standards of assertion, it also involves an idealisation. In most conversations, standards of precision are unevenly distributed amongst different topics under discussion. Surgeons performing a liver transplant might have extremely high standards of assertion with respect to topics like the patient's heart-rate, blood pressure and current dosage of anaesthetics but low standards of assertion with respect to topics like the music playing during the operation, weekend plans and other small talk. There are two ways one might go about representing the topic-relevance of standards of assertion. One way would be to imagine that multiple different conversations are occurring in the same spatiotemporal location, each with a different context containing a different standard of assertion. Another way, which I believe is more plausible, is to distinguish topics *within* a conversation and posit that conversations maintain different standards of assertion for each topic.

To achieve this, let a topic be any non-empty set $\mathcal{T}$ of propositional constants from $\mathcal{L}$. For instance, the topic of the liver transplant might be represented with a set $\mathcal{T}_t$ which contains basic propositions of the form 'The patient's heart-rate is x bpm' (H_x) and 'The patient has been given x mL of propofol' (P_x), whereas the small talk between surgeons might be covered by a topic $\mathcal{T}_s$ which contains basic propositions like 'It is warm today' (W), 'This is a Foo Fighters song' (F) and so on. These topics will allow us to group assertions in terms of their relevance to a given topic, allowing for different standards of assertion in different avenues of discussion.

Following this strategy, what does it take for an assertion of $\alpha \in \mathcal{L}$ to be relevant to a to the topic $\mathcal{T}$? One can answer this question by considering a more general question, which is; what are the conditions under which two propositions α and β are relevant to each-other? This broader question is foundational for relevant logicians, who seek to explicate a notion of validity where arguments of the form $\Gamma \therefore \alpha$ are invalid whenever α is, in some way, irrelevant to the premises in Γ. An early sketch of this kind of relevance is given by Anderson and Belnap (1971, p. 33), who suggest that the sharing of propositional constants is a necessary condition. Say that α and β are *syntactically relevant* if and only if there is some propositional constant P which is a subformula of both α and β. Syntactical relevance is insufficient for relevance in an intuitive sense since, as noted by Makinson (2009, p. 378), it is overly dependent on syntax. Whether an assertion is *on topic* should depend on the truth-conditions of the proposition it expresses, not on the symbols which that proposition happens to contain. A good case in point is that, although $(W \wedge (\neg F \vee W))$ is syntactically relevant to F, $(W \wedge (\neg F \vee W))$ is also logically equivalent to W, which isn't syntactically equivalent to F.[11] In light of this problem, Makinson (2009, p. 378) proposes a revised conception of propositional relevance which he calls *essential relevance.* To define essential relevance, Makinson uses a "fewest constants set" for each proposition α, denoted $\mathrm{FCS}[\alpha]$, which is the smallest set of propositional constants required to express a proposition which is equivalent to α. These fewest constants sets inform us what capital letters are redundant for the purposes of expressing a proposition. For example, $\mathrm{FCS}[(W \wedge (\neg F \vee W))] = \{W\}$, which shows us that the proposition's use of the constant F is superfluous. This concept can, in turn, be used to characterise essential relevance as follows:

Essential Relevance Any two propositions α and β are *essentially relevant* if and only if $\mathrm{FCS}[\alpha] \cap \mathrm{FCS}[\beta] \neq \varnothing$

The idea behind this definition is that two propositions are only relevant to one another when they have some non-superfluous propositional constant in common. Using this kind of relevance, I will say that an assertion of α is a contribution to the topic $\mathcal{T}$ just in case α is essentially relevant to some $P \in \mathcal{T}$. This captures the idea that an assertion can contribute to more than one topic of discussion: if the anaesthetist asserts $P_{20} \wedge W$, then this can be seen as a contribution to both $\mathcal{T}_t$ *and* $\mathcal{T}_s$, since $\mathrm{FCS}[P_{20} \wedge W] = \{P_{20}, W\}$

---

[11] At least, $(W \wedge (\neg F \vee W))$ and W are equivalent in classical logic and **FEL**. Feel free to substitute this pair of propositions with another if they aren't equivalent in your preferred logic.

which overlaps with the fewest constants-sets of basic propositions in both topics. By attaching standards of assertion to topics, as opposed to contexts, we can relativise the standard required for an assertion's accommodation to the topic (or topics, plural) to which it contributes. So, the surgeons will only accommodate assertions of H_{180} when their degree of presupposition in it is high, whereas they will be willing to accommodate assertions of W even when their degree of confidence in W is comparatively low.

Let's spell out how this proposal works in more detail. We begin by adding to each context $\mathfrak{C}$ a set of topics, denoted $\mathbf{TOP}_{\mathfrak{C}} = \{\mathcal{T}_1, ..., \mathcal{T}_n\}$, which are up for discussion. Then, we let $S_{\mathfrak{C}}(\mathcal{T}) \in [0,1]$ denote the standard of assertion required for topic $\mathcal{T} \in \mathbf{TOP}_{\mathfrak{C}}$. Adding these two features to contexts provides them with enough information to determine what standards of assertion apply when α is related to a single topic, but not enough to determine what standards apply when α is related to a range of topics. These latter standards can be found by associating standards of assertion with *sets* of topics, alongside topics individually. For any non-empty $\mathbf{TOP} \subseteq \mathbf{TOP}_{\mathfrak{C}}$, we define the standard of assertion for $\mathbf{TOP}$, $S_{\mathfrak{C}}(\mathbf{TOP})$, as $\max\{S_{\mathfrak{C}}(\mathcal{T}_1), ..., S_{\mathfrak{C}}(\mathcal{T}_n)\}$. Taking the maximum standard ensures that no assertions are permitted to sneak through under the radar. We wouldn't want the anaesthetist's assertion of $P_{20} \wedge W$ to be accommodated as if it were small talk, since it requires introducing presuppositions about the liver transplant, which is a topic with a much higher standard of assertion. These features all come together to give the following topic-sensitive dynamic rule for standards of assertion, where $\mathbf{TOP}_{\alpha}$ denotes the set of all topics in $\mathbf{TOP}_{\mathfrak{C}}$ which the proposition α contributes to:

TOPIC-SENSITIVE STANDARDS OF ASSERTION If at time t in a context $\mathfrak{C}$ some speaker asserts that α, which is a contribution to the range of topics $\mathbf{TOP}_{\alpha}$, then just after t—provided that the group G has a degree of presupposition in α which is equal to or exceeds $S_{\mathfrak{C}}(\mathbf{TOP}_{\alpha})$—the context will become $\mathfrak{C}'$, which is exactly like $\mathfrak{C}$ except that $\mathbb{P}$ is replaced with $\mathbb{P}(q, \mathcal{X})$ where $\mathcal{X} = \{\alpha, \overline{\alpha}\}$ and $q(\alpha) = 1$.

With this rule, we are able to represent the surgeon's conversation accurately by making $S_{\mathfrak{C}}(\mathcal{T}_t)$ sufficiently high and $S_{\mathfrak{C}}(\mathcal{T}_s)$ sufficiently low. They will then be prepared to accommodate assertions of F even if they aren't entirely sure whether the song currently playing is by the Foo Fighters or Nirvana, but won't be prepared to accommodate assertions of H_{100} unless they can confirm this with the heart monitor, do not suspect there is electrical interference and so on.

Another degree of flexibility which these conversational standards have over their legal counterparts concerns the shifting nature of the standards themselves. Lewis (1979, p. 352) observes there is a *sui generis* kind of accommodation which applies to standards of precision as an item on the conversational score. Suppose we are in a context where the standard of precision is too high for me to successfully assert that Harry (with his few patches of hair) is bald. If, later in this conversation, someone successfully asserts that Larry David (who has more than a few patches of hair) is bald, then the standard of assertion seems to shift, now rendering it acceptable to say that Harry is bald, too. Lewis (1979, p. 352) also proposes that these standards can shift in the other direction. We can raise the standards of assertion by rejecting someone's assertion that Larry David is bald, thereby making it unacceptable to say that Harry is bald. We can track both directions of standard shifting by introducing the following two accommodation rules for contexts:

STANDARD LOWERING If at time t the proposition α is presupposed to degree n and an assertion of α goes unchallenged, then just after t the context is changed so that, for each $\mathcal{T} \in \mathbf{TOP}_\alpha$, $S_\mathfrak{C}(\mathcal{T})$ is lowered to n (if it is not lower than n already).

STANDARD RAISING If at time t the proposition α is presupposed to degree n and an assertion of α is rejected, then just after t the context is changed so that, for each $\mathcal{T} \in \mathbf{TOP}_\alpha$, $S_\mathfrak{C}(\mathcal{T})$ is raised to be greater than n (if it is not already greater than n).

Note that there is an asymmetry between both rules. STANDARD LOWERING shifts the standard for assertion down to some determinate degree n, but STANDARD RAISING lifts the standard for assertion to an indeterminate degree which is 'greater than n'. How much greater than n this standard will be raised is presumably a function of how vehemently the assertion of α is rejected, although I am not in a position here to give some precise account of how one determines the 'strength of a rejection.'

Loose Talk and Verisimilitude

In many conversations, the bar for assertion is set so low with respect to some topics that speakers are willing to presuppose things which, according to them, are certainly false. This is the kind of standard which is invoked when numerical rounding occurs—e.g., when speakers accommodate the assertion that John is 2 metres tall after observing he is, in fact, one centimetre shorter than this. Let's call such locutions 'loose talk', referring to their tendency to be accommodated even when they are understood to be false. Loose talk cannot be modelled using the standards of assertion developed in the preceding section. Keeping track of a conversation's tolerance with respect to propositions which are presumed to be false involves making finer distinctions than these standards of assertion allow. If we were to set the standard of assertion with respect to a topic including John's various heights at zero, then not only would we be willing to accommodate the assertion that John is 2 metres tall, we would be willing to accommodate the assertion that he is any height whatsoever. If we were to set the standard any higher, then we wouldn't be willing to accommodate the assertion that John is 2 metres tall, since we have a zero strength presupposition that he is this height. In this section, I propose that we should treat loose talk using a concept from the philosophy of science: verisimilitude.

Verisimilitude is a quantity which was posited to make distinctions between propositions which are on par with respect to their degree of truth. If two propositions T_1 and T_2 are false, then they are both true to degree zero. However, we might still wish to say that T_1 is, in some respects, closer to the facts of the matter than T_2. Suppose T_1 is the claim that water is made up of H_2O_2 and T_2 is the claim that water is made up of C_6H_6. Then, T_1 clearly gets something right which T_2 doesn't—namely, that water is made up of a molecule which contains both hydrogen and oxygen. More generally, if T_1 and T_2 have the same truth value, we might still be able to distinguish them in terms of how close they are to being a correct theory. This quantity was first suggested by Popper (2002[1963], p. 315), who stressed this point, that degrees of verisimilitude are not the same thing as degrees of truth:

> every statement or theory is not only either true or false but has, independently of its truth value, some degree of verisimilitude [and this] does not give rise to any multi-valued logic [...]

The first task for any theory of verisimilitude is to describe in precise terms exactly what is meant by "the matters of fact" which propositions are being judged as closer to or further from. Oddie (1986, p. 11) suggests that when analysing verisimilitude, the "concept under study is that of closeness to the truth, in the sense of the *whole* truth, or the whole truth of some matter at hand." In other words, a proposition's verisimilitude is a measure of its closeness to some complete theory concerning a particular area of empirical inquiry. This framework builds on an analysis of verisimilitude given by Tichý (1978), who suggests that inquiry involves (i) a domain of discourse (denoted D) and (ii) what he calls an *intensional base*, which is a set of n-ary properties/relations/conditions/attributes that the inquirer is interested in. As a matter of fact, properties from the intensional base are distributed amongst members of the domain in a particular way; some of the things are red, some are not red, some of them are taller than others and so on. The aim of inquiry is to discover the way these properties are distributed in the actual world. Inquirers make a bunch of observations and put forward hypotheses about how the properties are distributed based on what they've observed. What verisimilitude is a measure *of* is the distance between these hypotheses and the true distribution of properties.

The kind of properties which Tichý (1978) envisages for these intensional bases are functions from n-tuples in D^n to truth values. For our purposes, the interesting case is when the intensional base is made up of 0-ary (or *medadic*) properties. Then, the domain of discourse becomes irrelevant and a distribution of properties is essentially just an assignment of truth values to basic propositions in the intensional base. Using some terminology from the preceding section, these intensional bases play the same role as topics—they outline a set of basic propositions which can be grouped under some common interest. Let us use $w(\mathcal{T})$ to denote the distribution of $\mathcal{T}$—the assignment of truth values to each proposition in $\mathcal{T}$—which holds at $w \in \boldsymbol{W}$. *The whole truth* with respect to $\mathcal{T}$ is the assignment of truth values to each of its basic propositions which holds at the actual world—i.e., $@(\mathcal{T})$. Where $W \subseteq \boldsymbol{W}$ is an event, we use $W(\mathcal{T})$ to denote the set of all distributions $w(\mathcal{T})$ such that $w \in W$. These sets of distributions are what play the role of the inquirer's hypotheses about $\mathcal{T}$.

Tichý (1978, p. 180) begins his analysis of verisimilitude with a definition of distance between $w(\mathcal{T})$ and $v(\mathcal{T})$ for any two possible worlds w and v.[12] This distance, denoted $\epsilon(w(\mathcal{T}), v(\mathcal{T}))$, is taken to be the number n of

---

[12] This is a simplification, made in order to focus on cases where the intensional base

propositions in $\mathcal{T}$ which $w(\mathcal{T})$ and $v(\mathcal{T})$ disagree on. That is, the number of propositions that are either true according to $w(\mathcal{T})$ and false according to $v(\mathcal{T})$, or true according to $v(\mathcal{T})$ and false according to $w(\mathcal{T})$. Since I am using a many-valued logic, it is appropriate to replace this measure of distance between distributions with the following measure, where $\mathcal{T} = \{P_1, ..., P_k\}$:

$$\epsilon(w(\mathcal{T}), v(\mathcal{T})) = \sum_{i=1}^{k} |\varphi(\langle w, P_i\rangle) - \varphi(\langle v, P_i\rangle)|$$

This measure takes each basic proposition in the topic $\mathcal{T}$, calculates the absolute difference between the truth values which that proposition gets assigned in w and v and then sums the results. If one imagines $\mathcal{T}$ as a line-up of measuring cylinders, one for each proposition, filled with different amounts of truth serum, then $\epsilon(w(\mathcal{T}), v(\mathcal{T}))$ tells you how much siphoning and filling of the cylinders would need to be done in order to move from a situation where they are filled according to $w(\mathcal{T})$ to a situation where they are filled according to $v(\mathcal{T})$, or vice versa.

Tichý (1978, p. 187) expands this measure into a full-blown definition of verisimilitude by using it to introduce a measure of the distance between a given distribution $w(\mathcal{T})$ and a hypothesis $W(\mathcal{T})$. This distance is taken to be the average distance between $w(\mathcal{T})$ and $v(\mathcal{T})$, for each $v \in W$. Where k is the number of worlds in W, this average distance can be calculated as follows:

$$d(w(\mathcal{T}), W(\mathcal{T})) = \sum_{v \in W} \frac{\epsilon(w(\mathcal{T}), v(\mathcal{T}))}{k}$$

The verisimilitude of $W(\mathcal{T})$, denoted $\Delta(W(\mathcal{T}))$, is then just taken to be the value $d(@(\mathcal{T}), W(\mathcal{T}))$. In other words, the degree of truth-likeness for any hypothesis is just a matter of how far its worlds are, on average, from the truth of the matter concerning $\mathcal{T}$.

Before explaining how verisimilitude can be used to model loose talk, we should discuss two issues which might be worrying the reader at this point

---

only contains basic propositions. In actual fact, Tichý introduces a broader measure of distance between any two distributions of, potentially non-0-ary, properties or attributes. These distributions are represented using *constituents* from Hintikka (1973, p. 242), which are first-order sentences that describe a possible distribution of n-ary predicates across the domain of discourse.

in the reconstruction of Tichý's view. The first issue, which can be dealt with swiftly, is that the measure d is only well-defined when k, the number of worlds in W, is finite. However, since we are dealing with finite sample spaces, this does not undermine our ability to calculate the distance of any world-event pair in $\boldsymbol{W} \times \wp(\boldsymbol{W})$. The second issue comes from the fact that averaging isn't the only way to convert the between-worlds measure ϵ to a measure of the distance between propositions and worlds like d. Oddie and Cevolani (2022) call this the "extensional problem" for verisimilitude. An adequate theory of likeness to the truth needs to give some principled reason for favouring one way of combining the distances measured by ϵ. However, I am in agreement with Tichý (1978, pp. 180,187) that there is principled reason for favouring averaging. He suggests that inquirers are like archers in that, when they propose a theory T, each $w \in T$ can be viewed as a shot at the complete truth of the matter. The inquirer, like the archer, isn't judged more favourably with respect to one of these shots rather than another. Averaging ensures that each $w \in T$ is assigned the same weight, proportional to the amount of space it takes up in the theory T. Things might be different if, instead of proposing a theory T, the inquirer proposed some distribution of confidence over subsets of T. However, this is a different type of inquiry, in the same way that archery would be a different game if archers were permitted to decide how much each of their shots on target should be weighted when calculating their overall points.

If the inquirer were making probabilistic qualifications over a theory T, then it would still be fitting to measure verisimilitude of these qualifications by using the weighted average of each world in T's distance from @—however, in that case, it would be important to weigh each distance by the probability which the inquirer qualifies each event $\{w\}$ as having, rather than the proportion $^1/_k$. Suppose that the inquirer makes a qualification (what is essentially a qualified assertion) of the form $\langle q, \mathcal{X} \rangle$, where $\boldsymbol{W} = \{v_1, ..., v_n\}$ and $\mathcal{X} = \{\{v_1\}, ..., \{v_n\}\}$. Let $T = \{v \in \boldsymbol{W} : q(\{v\}) > 0\}$. Where this is the inquirers qualification about the theory T, we should use the following calculation to determine the distance of $\langle q, \mathcal{X} \rangle$ from a given world w (concerning matters of $\mathcal{T}$):

$$d(w(\mathcal{T}), \langle q, \mathcal{X} \rangle(\mathcal{T})) = \sum_{v \in \boldsymbol{W}} q(\{v\}) \cdot \epsilon(w(\mathcal{T}), v(\mathcal{T}))$$

When measuring d this way, the verisimilitude of the inquirer's opinions depends not just on what shots at the truth are contained in T, but how likely the inquirer thinks each of these complete theories are. They are 'scored' for verisimilitude in the same way as the archer who gets to pick what percentage of points each shot of an arrow is worth.

This notion of 'scoring' a probabilistic qualification is analogous to the use of *scoring rules* in accuracy-centred epistemology (See Joyce 1998, 2018 and Pettigrew 2016). Scoring rules and measures of verisimilitude are two different ways of determining, in two different senses, the accuracy of an agent's probability measure. Scoring rules measure the kind of accuracy which an agent achieves by ensuring their expectation of each proposition's truth-value is close to that proposition's truth-value—i.e., by ensuring that $E^i_w(\alpha)$ is close to $\varphi(\langle w, \alpha \rangle)$.[13] Verisimilitude, on the other hand, measures the kind of accuracy which an agent achieves by investing more confidence in possibilities which are close to the actual world, as measured by ϵ. Both measures agree on what counts as perfect accuracy: ensuring that your degree of belief in all propositions matches their truth value exactly is the same thing as assigning maximum probability to the actual world. However, when it comes to probability measures which fall short of this ideal, these measures rank the alternatives differently.

To see why, take the Brier score as an illustrative example of a scoring rule, used by Joyce (1998) and Pettigrew (2016). Given a topic $\mathcal{T}$ which contains the basic propositions which an agent i is interested in, this scoring rule calculates the inaccuracy of their probability measure Pr^i_w as follows:

$$\mathfrak{b}(w(\mathcal{T}), Pr^i_w(\mathcal{T})) = \sum_{\alpha \in \mathcal{T}} [\varphi(\langle w, \alpha \rangle) - E^i_w(\alpha)]^2$$

For comparison with this Brier score, we re-interpret Tichý's verisimilitude measure d as a score of probability measures, as opposed to probabilistic qualifications:

$$d(w(\mathcal{T}), Pr^i_w(\mathcal{T})) = \sum_{v \in \boldsymbol{W}} Pr^i_w(\{v\}) \cdot \epsilon(w(\mathcal{T}), v(\mathcal{T}))$$

Now, let's consider the kind of case where the two scores disagree. Let $Pr^a_{w_1}$ and $Pr^b_{w_1}$ be two probability measures over the sample space $\boldsymbol{W} =$

---

[13] In the words of Joyce (1998, p. 588), scoring rules will "gauge the extent to which the truth-value estimates sanctioned by [some credence] diverge from the truth-values that propositions would have were ω actual."

$\{w_1, w_2, w_3, w_4\}$. Suppose $\mathcal{T} = \{P, Q\}$ and let w_1 be a world where P and Q are both true, w_2 be one where P is true and Q is false, w_3 be one where P is false and Q is true and w_4 be one where they are both false. $Pr^a_{w_1}(\{w_2\}) = Pr^a_{w_1}(\{w_3\}) = Pr^a_{w_1}(\{w_4\}) = \frac{1}{3}$, whereas $Pr^b_{w_1}(\{w_2\}) = 1$. That is, a spreads their confidence evenly amongst w_2, w_3 and w_4 but b is certain that w_2 is actual. According to the Brier score, $\mathfrak{b}(w_1(\mathcal{T}), Pr^a_{w_1}(\mathcal{T})) < \mathfrak{b}(w_1(\mathcal{T}), Pr^b_{w_1}(\mathcal{T}))$, meaning $Pr^a_{w_1}$ is counted as less inaccurate than $Pr^b_{w_1}$. However, according to the verisimilitude measure, $d(w_1(\mathcal{T}), Pr^a_{w_1}(\mathcal{T})) > d(w_1(\mathcal{T}), Pr^b_{w_1}(\mathcal{T}))$, meaning $Pr^b_{w_1}$ is counted as less inaccurate than $Pr^a_{w_1}$. The disagreement between these two measures boils down to the difference between the two notions of accuracy they describe: one favours accurate estimates of the propositions in $\mathcal{T}$ at the expense of closeness to the whole truth, whereas the other favours being highly confident in a state of affairs which is close to conveying all the facts concerning propositions in $\mathcal{T}$. Since my goal is to explain how the accommodation of presumed falsehoods is facilitated, I will be focusing on measures of accuracy in the latter sense, which favour high confidence in near truths.[14]

Here is my proposal for how verisimilitude can be used to model loose talk. If a speaker asserts α and accommodating this assertion doesn't require transporting probability too far, as measured by d, then the revision will be carried out and α will be accommodated. What counts as too far? This depends on what topics α contributes to and the degree of approximation to the truth which speakers are willing to permit relative to each of them. Let us use $a(\mathcal{T}) \in [0, \infty)$ to denote the maximum distance from their presuppositions, calculated by d, which speakers are willing to permit concerning matters about $\mathcal{T}$. Intuitively, talk about $\mathcal{T}$ becomes looser as $a(\mathcal{T})$ increases and is maximally strict, on the other hand, when $a(\mathcal{T}) = 0$. To talk about where probability is shifted from or to after a context revision, we will use $G_{\mathfrak{C}}$ to mean, in any context $\mathfrak{C}$, the set of worlds which are G-reachable from @. Given these details, I submit the following dynamic rule for loose talk:

LOOSE TALK If at time t someone asserts α in context $\mathfrak{C}$ then just after t the context set becomes $\mathfrak{C}'$, which is exactly like $\mathfrak{C}$ except that $\mathbb{P}$ is replaced with $\mathbb{P}(q, \mathcal{X})$ where $\mathcal{X} = \{\alpha, \overline{\alpha}\}$ and $q(\alpha) = 1$—provided that, for

---

[14] See Oddie (2019) and Dunn (2019) for a discussion of some attempts to combine the two kinds of accuracy into a unified measure.

each $w \in G_{\mathfrak{C}}$ and $\mathcal{T} \in \mathbf{TOP}_{\alpha}$, $d(w(\mathcal{T}), G_{\mathfrak{C}'}(\mathcal{T})) < a(\mathcal{T})$

Let's consider how this loose talk rule might be applied in the example considered at the beginning of this section, where speakers accommodate the assertion that Harry is 2 metres tall despite it being proven beyond any doubt that he is 1.99 metres tall. Suppose that the speakers are measuring Harry's height with a ruler that is 3 metres long and they are working with one topic, $\mathcal{T}_H$, which contains two kinds of basic proposition for each x such that $1 \leq x \leq 300$: H_{tx} (Harry is taller than x cm) and H_{sx} (Harry is shorter than x cm). The speakers are able to learn the actual distribution of truth values across $\mathcal{T}_H$, i.e. $@(\mathcal{T}_H)$, by holding the ruler up to Harry, thereby seeing his exact height to the nearest centimetre. If they do this, they will learn that H_{tx} is true if $x < 199$ and false otherwise, whereas H_{sx} is true if $x > 199$ and false otherwise. If this context after the observation is $\mathfrak{C}$, then $G_{\mathfrak{C}} = \{@\}$, since speakers are certain of Harry's actual height. When someone asserts that Harry is 2 metres tall in this context, they are proposing to move all probabilistic mass from $\{@\}$ to the set $\{w_{200}\}$, where w_{200} is a world exactly like @ except H_{t199} is true instead of false and H_{s200} is false instead of true. Since two propositions from $\mathcal{T}_H$ have the greatest possible disagreement in terms of truth value at @ and w_{200}, the value of $\epsilon(@(\mathcal{T}_H), w_{200}(\mathcal{T}_H))$ is 2. Since, in this case, $G_{\mathfrak{C}'}$ is just the singleton set $\{w_{200}\}$, this also means that $d(@(\mathcal{T}_H), G_{\mathfrak{C}'}(\mathcal{T}_H)) = \epsilon(@(\mathcal{T}_H), w_{200}(\mathcal{T}_H)) = 2$. If $a(\mathcal{T}_H) > 2$, then speakers will accommodate the assertion that Harry is 2 metres tall via LOOSE TALK, whereas if $a(\mathcal{T}_H) < 2$, the standard of precision will be too strict too allow this instance of numerical rounding.

According to LOOSE TALK, there's nothing stopping the speakers from going the other way and accommodating an assertion that 'Harry is 1.98 m tall' relative to the same initial context. In terms of $\mathcal{T}_H$, the distance between @ and w_{198} is also 2, meaning that accommodating this assertion requires no more shifting of probabilistic mass than would be required by the accommodation of 'Harry is 2 m tall'. However, while such an accommodation would be permitted by LOOSE TALK, adopting the presupposition that Harry is 1.98 m tall may present the speakers with issues unrelated to accuracy, such as difficulty of calculating with 'non-round' numbers like 198 and needless prolixity when expressing Harry's height. So, an assertion that 'Harry is 198 m tall' may nonetheless be blocked by speakers due to these peripheral concerns. In the following section, I'll investigate in more detail some of the ways in which a proof or assertion might be rejected for reasons other than its strength or accuracy.

Besides modelling a group's tolerance when it comes to accommodating things which are presumed to be false, LOOSE TALK is also helpful for a general understanding of the extent to which a conversation is willing to change their opinions about $\mathcal{T}$. As discussed above at the end of §1.4.1, some people are skeptical the kind of context revision which is allowed by imaging because, like Rudin (2018, p. 372), they think that a successful assertion of some proposition α should not be followable with the accommodation of $\neg\alpha$ thereafter. I, instead, argued that this pattern of accommodation happens all the time. We are now in a position to construct a more moderate stance, in case your sympathies lie somewhere between mine and Rudin's: speakers will accommodate $\neg\alpha$ after accommodating α, but *only* if doing so does not spread the G-reachable worlds too far away, as indicated by $a(\mathcal{T})$, for any topics $\mathcal{T} \in \mathbf{TOP}_\alpha$. When playing this role, the values set by a can be thought of as tracking a conversation's openness to changing its mind. Where $a(\mathcal{T}) = 0$, the conversation won't be prepared to accommodate any assertions on the topic $\mathcal{T}$ which take them to a context $\mathfrak{C}'$ where $G_{\mathfrak{C}}(\mathcal{T}) \neq G_{\mathfrak{C}'}(\mathcal{T})$, since this will ensure that for at least one $w \in G_{\mathfrak{C}}$, the distance $d(w(\mathcal{T}), G_{\mathfrak{C}'}(\mathcal{T}))$ is positive. That is, when $a(\mathcal{T}) = 0$, speakers will be completely unwilling to change any of their full presuppositions concerning $\mathcal{T}$. However, when $a(\mathcal{T})$ is positive, the size of this positive margin indicates the extent to which speakers in G *are* willing to change these presuppositions about $\mathcal{T}$. If $a(\mathcal{T})$ is sufficiently large, then speakers in G will even be prepared to go from presupposing a proposition to presupposing its negation.

2.2 Some Different Measures for a Proof's Success

In the previous section, I was concerned with spelling out various ways a proof might be strong enough to guarantee an assertion's accommodation. However, whether a proof is appropriately employed in a given context depends on more than just whether it strongly confirms a proposition. Consider the following two scenarios which feature some proof which intuitively fails (or falls short) for some reason other than sufficiency of strength:

BANANA A storeowner wants to know how ripe the bananas are in a fresh shipment. As it happens, the bananas are fairly ripe, making B true to degree 0.6. While the storeowner doesn't know that $\varphi(\langle @, B\rangle) = 0.6$, he does know from past experience that the shipment could be ripe to any degree $^{n}/_{5}$, where $0 \leq n \leq 5$. One of the store clerks, Clive, gives 0.6-strength proof that B by providing the storeowner some facts about the frequency of various degrees of ripeness in past shipments; pointing out that (i) 25% of previous shipments have contained bananas which were ripe to degree 0.8, (ii) 50% of previous shipments have contained bananas which were ripe to degree 0.6 and (iii) the remaining 25% of shipments all contained bananas which were ripe to degree 0.4, concluding that the expected ripeness of the new shipment is $(0.25 \cdot 0.8) + (0.5 \cdot 0.6) + (0.25 \cdot 0.4) = 0.6$. Another one of the store clerks, Claude, gives a 0.6-strength proof that B by picking a banana from the shipment and showing it to the storeowner. Both Claude and Clive's demonstrations are convincing 0.6-strength proofs that the bananas are ripe, but clearly there is some sense in which Clive's proof is inferior.

CEMPEDAK Two friends are at a fruit market and they notice there is a spiky, tropical fruit on display. One of the friends asks if the fruit smells of garlic, which he had heard is the case with some tropical fruits. The other friend responds by showing all trademark features which can be used to identify the spiky fruit, such that it is proved beyond all doubt to be a cempedak, as opposed to a durian, breadfruit or jackfruit. Confused, the friend asks if cempedaks have the aroma of garlic and is met by a shrug. Here, the tropical fruit enthusiast gave a maximally strong proof, but it is inappropriate insofar as his friend didn't request to know whether the fruit was a cempedak.

These two cases are thinly-veiled allusions to a pair of Grice's (1989, p. 28) conversational maxims: BANANA violates maxim of quantity (which says to be as informative as required) and CEMPEDAK violates the maxim of relation (which says to make a contribution which is relevant to the conversation). Clearly, just like assertions, proofs are sometimes evaluated against standards which are orthogonal to their strength or 'quality'. So, a complete account of how proofs perform in a context also needs some way of showing how contexts set these non-qualitative standards and when proofs are able to meet them. In the following two sections, I will explore how one might expand context to include items which track the conditions which render these two proofs inappropriate.

2.2.1 Maxim of Quantity

The insight from BANANA is that two proofs might be on par with respect to their strength but diverge with respect to how much information they convey. Claude and Clive's proofs both confirm B to the same degree, but there is a sense in which Claude's proof, where a 0.6-ripe banana is shown, eliminates more uncertainty for the storeowner than Clive's proof does. This is made clear by considering that, after Clive's proof, where a statistical distribution of ripeness is given, the storeowner is not able to say how ripe the bananas are. He is only able to guess, based on his current estimations, that the bananas are ripe to degree 0.6. This is the sense in which Claude's proof contains more information than Clive's. It leaves the storeowner with less room for doubt. In this section, we will try to see if this intuitive notion can be sharpened using information theory, which prescribes a series of methods for measuring how much uncertainty is reduced by a given piece of data.

First, some terminological housekeeping. There is an important distinction in information theory between *quantitative* and *semantic* conceptions of information. Weaver (1949, p. 8) teases out this distinction by noting that, in its quantitative sense, the "word information [...] relates not so much to what you *do* say, as to what you *could* say." The amount of information which a speech act delivers to some listener depends on how much uncertainty it removes for them as compared to other speech acts which could've been performed, not on whether the speech act confirms something which is true or false. Claude would've provided the same quantity of information to the store-owner if he had convinced them of the false claim that the bananas are ripe to degree 0.8. In either case, the store-owner's overall uncertainty about the bananas is reduced by the same amount. This is in contrast to information in a semantic or qualitative sense, where we are tempted to say that Claude's 0.8-strength demonstration carries no information at all. Claude cannot provide the store-owner with 'information *that* the bananas are ripe to degree 0.8', since the bananas aren't ripe to that degree. This is no doubt the kind of information Grice (1989, p. 371) has in mind when he says that, "False information is not an inferior kind of information; it is just not information." Semantic information that P has the truth of P as a necessary condition, whereas there is still quantitative information attached to an assertion or speech act alleging that P even if P isn't true.

Which of these two notions of information is relevant for modelling the Gricean conversational maxims? In spite of Grice's comment, I suggest that—like the name suggests—the maxim of quantity, which holds that cooperative speakers should make appropriately informative contributions, is a decree about quantitative information. If we took this maxim to be about semantic information, we would then lose the ability to make a clear distinction between the maxims of quantity and quality. This is, in fact, why Grice (1989, p. 371) suggested that false 'information' is no information at all; he makes this comment in the context of suggesting that assertions can't be judged against *any* of the conversational maxims unless they are first taken to be truthful, or at least of sufficient quality. However, in spite of this suggestion, different conversations set different benchmarks for what counts as cooperative behaviour. Some speakers, like podcast hosts, might prioritise interesting and outrageous assertions over truthful ones. Being cooperative in these contexts involves raising regard for the maxim of quantity and lowering regard for the maxim of quality. This is why it is better to interpret the maxim of quantity without an assumption of quality, which is only possible when we take it to be a rule concerning information in the quantitative sense.

Shannon's (1949) theory of communication concerns how information is transmitted from a point of origin (called the *source*) to some point of reception (called the *receiver*).[15] In the simplest type of communication, the source consists in a finite set of possibilities $S = \{s_1, ..., s_n\}$ which is reduced to a particular member (s_i) by a process of selection; a television host draws some lottery numbers, someone types up a text-message or a pair of dice land in a certain way. Information is created when this reduction of possibilities occurs. According to Shannon (1949, pp. 48—49), the amount of information to be associated with the selection of s_i decreases proportional to the inevitability of that selection. For him, information is a measure of how surprising it would be to learn that s_i was selected from the class of possibilities S. If each member of S is given a probability of being selected, then the information (or *surprisal*) generated by selecting s_i, denoted $I(s_i)$ is,

---

[15] Here, I will follow quite closely the philosophical presentation of this theory given in Dretske (1981b, pp. 3—39).

$$I(s_i) = \log_2 \frac{1}{Pr(s_i)}$$

This information, generated by s_i at the source, is then transmitted to the receiver where a parallel process of selection occurs. The receiver has some stock of possibilities $R = \{r_1, ..., r_n\}$ under consideration and the arrival of some data from S facilitates the elimination of some of these possibilities; the TV set shows a lottery ticket holder which numbers were drawn, the awaited text-message is received or the gambler learns how the dice landed by looking at them. $I(r_i)$, the amount of information associated with the reduction of R to r_i, is also calculated as the surprisal of r_i.

Not-so-careful inspection of the formula for calculating surprisal reveals that the quantity is undefined whenever the probability of a selection being made ($Pr(s_i)$ or $Pr(r_i)$) is zero. Much like conditional probability, the concept of surprisal behaves poorly when the probabilities involved reach their lower limit. However, despite technically being undefined, it is conventional to talk of s_i as containing unlimited or infinite amounts of information whenever $Pr(s_i) = 0$ and s_i is impossible (see Carnap and Bar-Hillel, 1952, p. 31). This convention is motivated by the fact that the quantity ${}^1/_{Pr(s_i)}$ approaches infinity as $Pr(s_i)$ approaches 0 and the logarithm of x approaches infinity as x approaches infinity. If, as Weaver suggests (1949, p. 9), information is a measure of how much "freedom of choice" is involved in the selection of s_i, then this freedom appears to be *absolute* when $Pr(s_i) = 0$.

Faultless communication occurs when the information generated by the process S reaches the receiver R. As an example, suppose that $S = \{s_H, s_T\}$ consists in the flipping of a fair coin by Alice and $R = \{r_H, r_T\}$ consists in Bob hearing that the coin Alice flipped has landed either heads or tails. By the fairness of the coin, $Pr(s_H) = Pr(s_T) = 0.5$, so it follows that $I(s_H) = I(s_T) = 1$, i.e., one unit of information is generated by the coin flip. If Alice decides to say 'Heads!' whenever s_H is chosen and 'Tails!' whenever s_T is chosen, then the amount of information associated with r_H and r_T is also 1, since the probability of Bob hearing either message is tied to the probabilities of Alice's coin-flip. In this case, the information generated by Alice's coin flip reaches Bob. However, if Alice instead decided that she would say 'Heads!' no matter how the coin landed, then the information generated by the coin flip would not reach Bob. Indeed, since this second strategy from Alice raises $Pr(r_H)$ to 1, the amount of information which Bob receives from r_H is none at all, since then $I(r_H) = 0$.

For modelling the maxim of quantity, the key takeaway from Shannon's theory of communication is the notion of surprisal. As Dretske (1981b, p. 14) points out, we can re-describe the measuring of information at S and R as measures of the information associated with events (i.e., subsets of $\boldsymbol{W}$) in a general sense; there is some uncertainty about which events will come out true, and once one of them does, "we can take what did happen as a reduction of what *could* have happened to what *did* happen and obtain an appropriate measure of the amount of information associated with the result." However, a natural question which might be asked here is, whose uncertainty does the ruling out of particular events reduce? In other words, whose probability measure is to be used when calculating the surprisal of an event E? Hake (1959, p. 257) suggests there are two types of information which can be measured where the choice of probability measure here determines which:

> There is, first, the measurement based upon the actual probability of occurrence of the message. There is, second, the information measurement which is based upon the subjective notions of the receiver about the likelihood of occurrence of each of the possible messages in the set, as he sees them.

When discussing the information associated with the event of Bob hearing that the coin landed heads, $I(r_H)$, the relevant probability concerning r_H was taken to be its actual (or objective) probability. This is why $I(r_H)$ was taken to be zero when Alice adopts the lying strategy, because the message r_H was inevitable. However, like Hake suggests, we might also consider a subjective measure of information, denoted $I^B_w(E)$, which is calculated by taking the surprisal of E using Bob's subjective probability measure Pr^B_w, that is:

$$I^B_w(E) = \log_2 \frac{1}{Pr^B_w(E)}$$

If Bob *thinks* that hearing 'Heads!' is just as likely as hearing 'Tails!' and $Pr^B_w(r_H) = Pr^B_w(r_T) = 0.5$, then the subjective information generated by r_H (i.e., $I^B_w(r_H)$) is 1, even when Alice adopts her lying strategy. This subjective measure of information is what matters in BANANA, since Claude and Clive's respective proofs eliminate more or less of the storeowner's uncertainty, rather than more or less uncertainty in a metaphysical sense.

Proofs that P reduce a speaker's uncertainty by taking them from a position of doubt about P to a state of certainty about P. When a successful proof of P is performed, a speaker i's context set C_i is replaced with some new context set $C'_i \subseteq P$. The amount of information which i stands to gain from this successful proof is $I^i_{@}(C'_i)$. To appreciate how two proofs of the same proposition could differ with respect to how informative they are for i, consider the following two proofs:

DICE PROOF I Alice rolls a fair six-sided die which lands on the number 3 but Bob is too far away to see how the die landed. Alice proves to Bob that the die didn't land on 1 by placing an opaque sheet over the face of the die which has a window in the top-left corner. This informs Bob that the die could've landed on any number except 1.

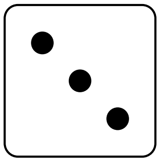
Alice sees

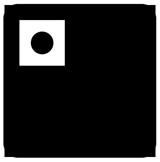
Bob sees

DICE PROOF II Alice rolls a fair six-sided die which lands on the number 3. Alice then proves to Bob that the die didn't land on 1 by showing Bob that the dice landed on 3.

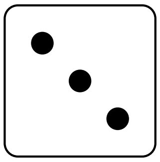
Alice sees

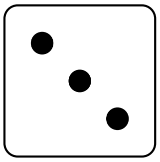
Bob sees

These proofs both begin with the sample space $\boldsymbol{W} = \{w_1, w_2, @, w_4, w_5, w_6\}$, where w_n is a world where the die lands on n and @ is the actual world, where the dice lands on 3. Just before either of the proofs is performed, Bob thinks each of the outcomes is equally likely, such that $Pr^B_{@}(\{w\}) = 1/6$, for all $w \in \boldsymbol{W}$. When Alice shows Bob the covered die face, as in DICE PROOF I, this takes Bob from the context set $C_B = \boldsymbol{W}$ to the context set

$C^I_B = \{w_2, @, w_4, w_5, w_6\}$. Because $Pr^B_@(C^I_B) = {}^5\!/\!_6$, the amount of information which this covered die coneys to Bob is $I^B_@(C^I_B) \approx 0.26$. When Alice shows Bob the uncovered die face, as in DICE PROOF II, this instead takes Bob to the context set $C^{II}_B = \{@\}$. Because $Pr^B_@(C^{II}_B) = {}^1\!/\!_6$, the amount of information which the uncovered die conveys to Bob is $I^B_@(C^{II}_B) \approx 2.6$. So, despite both of Alice's proofs being proofs *of* the event $\{w_2, @, w_4, w_5, w_6\}$, the second proof gives Bob at least two units of subjective information more than the first.

More information isn't always preferable. There are contexts where DICE PROOF II proves too much and DICE PROOF I is just right. It may have been stipulated, as part of the dice game Alice and Bob are playing, that Bob wins whenever the die lands on 1 but that he must guess the number correctly in order to win when the die has landed on any other side. When showing Bob that the dice hasn't landed on 1, Alice is required to do so without giving up the game, which she can only do if she gives Bob no more information than is provided by DICE PROOF I. When a context is concerned with providing some proof of P, whether a proof that P is appropriate depends on whether the amount of information it generates lies somewhere within an acceptable range. This goldilocks zone of informativeness was suggested by Grice (1989, p. 26), who introduced the maxim of quantity as a two-part rule for conversations:

1. Make your contribution as informative as is required (for the current purposes of the exchange).
2. Do not make your contribution more informative than is required.

The structure of this maxim is shaped by what we expect from polite conversation. When I ask what you had for breakfast, I anticipate an answer that's more informative than 'Some food' and less informative than a complete description of the bowl of porridge you ate, down to the number of oat grains ingested. We can represent the lower and upper bounds of expected informativeness by adding an *information range* $I[x, y]$, for some $x, y \in [0, \infty)$ to the n-tuple which forms a context. Intuitively speaking, when a context $\mathfrak{C}$ has the information range $I[x, y]$, each speaker $i \in G$ expects speech acts to convey no more information than y and no less information than x. A proof violates the maxim of quantity if it takes at least one agent i to a context set C_i' where $I^i_@(C_i') \notin [x, y]$.

Now we have a way of enshrining the maxim of quantity in contexts. So far, so good. However, there are two oversights which will need to be addressed before we can provide a wholly adequate model for BANANA and subjective information in general. Firstly, the quantity $I^i_w(C_i)$ isn't fine-grained enough to make distinctions between the amount of information conveyed by different distributions of probability over the proposed context set C_i. If a speaker first believes that a series of outcomes in some event E are equally likely, then we should expect they would be more surprised to hear that a particular outcome in E is extremely likely than they would to hear merely that E is true. However, $I^i_w(C_i)$ is the same value regardless of the probabilities which the post-proof measure $Pr^{i\prime}_@$ assigns to events in C_i, since it depends solely on the value of $Pr^i_w(C_i)$. The second issue with using I^i_w to measure subjective information is that $I^i_w(C_i)$ is infinite whenever $Pr^i_w(C_i) = 0$. If i believes it will rain tomorrow and $Pr^i_w(R) = 1$, then they stand to gain the same amount of information from a proof that it won't rain tomorrow as they would from a proof that it won't rain tomorrow because, in particular, it will be sunny. This will be the case because $I^i_w(E)$ is infinite whenever $Pr^i_w(E) = 0$, no matter how small the subset E is relative to the sample space $\boldsymbol{W}$. Additionally, when E conveys infinite amounts of information for at least one speaker i, then any proof of E will be in violation of the maxim of quantity. This is because, by stipulation, $I^i_@(E)$ exceeds all $x \in [0, \infty)$. Before addressing the example BANANA, we need to address these oversights for the provided model of subjective information.

First, I'll address the issue concerning infinite subjective information expressed by probability zero events. There is a certain respect in which, despite appearances, this is a good feature of the proposed model of subjective information. As a function of one's subjective probability, it's not obvious that there should be different amounts of subjective information conveyed by events which are assigned the same probability by a speaker. If i thinks that the event E is impossible, then hearing E will reduce the maximum possible amount of uncertainty, regardless of how many impossibilities it eliminates from i's perspective. The quantity I^i_w respects this intuition by making $I^i_w(X)$ and $I^i_w(Y)$ equal and infinite whenever $Pr^i_w(X) = Pr^i_w(Y) = 0$. However, on the other hand, there is also a clear and intuitive sense in which i might be more surprised to hear X than Y, even though i thinks both events are impossible. Alice, who thinks Phar Lap will win the race, would be more surprised to hear that the least favourite horse won than she would be to hear that the second-favourite after Phar Lap won. These two intuitions are in tension. How can we maintain that subjective information depends solely on an agent's subjective probability but also suggest that, of two events with the same probability, one might be more surprising to the agent than another? I suggest that we might navigate this tension by applying the same technique I used to allow for agents who have different levels of confidence in events which they nonetheless fully believe, discussed in §1.3.3 above. To do this, I will still take the subjective information contained in an event to be $I^i_w(E)$ but also introduce a separate way of determining, for any two events, which of the two is more surprising for i given their beliefs.

Recall that we took an agent to be more confident in X than Y just in case, were they to stop believing $X \wedge Y$, they would stop believing Y before they stop believing X. This was spelled out using systems of spheres centred on an agent's context set, but we can also describe this relation using an agent's selection function σ. Say that an agent i is more confident in X than they are in Y, denoted $X >_i Y$ if and only if, for all $w \in C_i$, their selection function is such that $\sigma_i(w, \overline{X} \cup \overline{Y}) \subseteq X \cap \overline{Y}$. That is, if they move all probabilistic mass to X-worlds where Y is not true upon learning that X and Y can't be true together. If $X \not>_i Y$ and $Y \not>_i X$, then we say that i has equal confidence in events X and Y. This will be the case whenever, upon revising their beliefs with the constraint that X and Y can't both be true, i stops believing both propositions.

To get a relation of *comparative surprise* between two events for an agent, we can utilise this confidence relation $>_i$. Let us say that i would be more surprised to hear that X than they would be to hear that Y if and only if $\overline{X} >_i \overline{Y}$, i.e. if they are more confident that X is not true than they are that Y is not true. This relation is able to represent Alice as being more surprised in the underdog winning than she is in the second-favourite winning, even when she thinks neither event is possible. To see how, let P, U and S respectively be the events of Phar Lap, the underdog and the second-favourite winning the race. Alice believes P, which excludes U and S being true, but suppose she is more confident in the non-truth of U than she is in the non-truth of S, i.e., $\overline{U} >_A \overline{S}$. This will be the case if $\sigma_A(w, U \cup S) \subseteq \overline{U} \cap S$ for each $w \in C_A$. That is, if she were to learn that Phar Lap lost, Alice would sooner believe the second-favourite horse won. This shows us that, while two events X and Y might be indistinguishable in terms of the amount of subjective information they convey given some probability measure Pr^i_w which assigns 0 to both, they can still be distinguished in terms of how surprised i would be to hear one or the other.

The resulting scale of comparative surprise is an ordinal one, so there's no way to say whether X is inside an information range $I[x, y]$ based merely on X's comparative surprise in relation to other events. However, by considering some of Grice's other remarks about conversational information, we can still salvage a respectable auxiliary to the maxim of quantity which applies to events which are assigned probability zero. A point which Grice (1989, p. 27) makes about the maxim of quantity is that one could potentially replace its upper bound for appropriate information with the following proviso: give as much information as you can, provided that giving this information does not violate any other conversational maxims. That is, a cooperative speaker should always lead with the most surprising thing they're prepared to say, provided that this potential assertion would not confuse interlocutors, waste their time or fail to be relevant. We can enshrine this rule as a separate maxim about probability zero events as follows:

MAXIM OF COMPARATIVE QUANTITY If a speaker in G is considering asserting X or Y and there is at least one $i \in G$ such that $Pr^i_{@}(X) = Pr^i_{@}(Y) = 0$ and furthermore $\overline{X} >_i \overline{Y}$ for such i, then the speaker should assert X—provided that doing so does not violate any other conversational maxims in force.

In other words, a cooperative speaker who is entertaining two assertions whose content are assigned probability zero by some speakers in G should lead with the assertion which is most news-worthy or surprising, provided that this does not lead to the speaker being non-cooperative in other ways. If Bob sees the race results and is considering whether to tell Alice that Phar Lap lost or that the underdog won, he should clearly tell her that the underdog won. That is, unless doing so would be irrelevant or otherwise confusing given the details of their conversation.

A speaker who does what the maxim of comparative quantity demands will also violate the regular maxim of quantity, since they will end up asserting a proposition whose subjective information falls outside of all information ranges for those speakers who assign the asserted proposition X probability zero. However, they can still minimise the impact of this violation by ensuring that they obey the maxim of quantity with respect to speakers in G who don't assign X probability zero. That is, by ensuring $I^i_@(X) \in I[x, y]$ for all i such that $Pr^i_@(X) \neq 0$. This can be enshrined by adding a small addendum to the maxim of comparative quantity: ...the speaker should assert X, provided that doing so does not violate any conversational maxims other than the maxim of quantity and this violation only occurs with respect to agents who assign X probability 0. A clear division of labour is carved out by the two quantitative maxims. You should always keep your assertions within the context's information range whenever possible, but you should give the most informative and otherwise cooperative assertions whenever doing so is not possible.

As for the other problem, of measuring subjective information in a way which is fine-grained enough to pick up on differences in probability measures over some proposed new context set C_i, this can be achieved by using *Kullback-Leibler divergence* (or KL divergence),[16] which is a way of measuring the differences in amounts of information between two probability measures with respect with respect to a series of events. If we suppose that an agent i has the probability measure Pr^i_w, then the amount of subjective information they can expect to gain by revising Pr^i_w with the constraint $(q, \mathcal{X})$ is taken to be the following:

$$I^i_w(q, \mathcal{X}) = \sum_{X \in \mathcal{X}} Pr^i_w[(q, \mathcal{X})](X) \cdot \log_2 \left(\frac{Pr^i_w[(q, \mathcal{X})](X)}{Pr^i_w(X)} \right)$$

Intuitively, KL divergence tracks how much mass from the initial probability measure must be re-distributed into other events in order to accommodate the constraint $(q, \mathcal{X})$. Thus, the reduction of uncertainty which $(q, \mathcal{X})$ provides for i at world w is related to how much of Pr^i_w must be transformed in order to reach $Pr^i_w[(q, \mathcal{X})]$. Note that whenever $Pr^i_w = Pr^i_w[(q, \mathcal{X})]$, the value of $I^i_w(q, \mathcal{X})$ is 0, which reflects the idea that trivial revisions of one's probability carry no subjective information whatsoever. Furthermore, the KL divergence between Pr^i_w and $Pr^i_w[(q, \mathcal{X})]$ is infinite whenever $Pr^i_w(X) = 0$

---

[16] KL divergence was first presented in Kullback and Leibler (1951). I came across this particular way of applying KL divergence in Skyrms (2010, p. 36), who used it to measure the divergence between conditional and prior probabilities over some set of events. The main difference here is that I am considering the divergence between partial revisions of probabilities (i.e., Jeffrey conditioned or Fusco imaged probabilities) and their priors.

and $q(X) > 0$, suggesting that all revisions which are made via Fusco imaging carry infinite amounts of information for the speaker i at w.

The main benefit which comes from using KL divergence to measure the quantity $I^i_w(q, \mathcal{X})$ is its ability to factor in the shape of a probability distribution over some proposed context set C_i, assigning them different amounts of subjective information. This will become clear by considering its application in a model of BANANA, which I will now provide. This example begins with a sample space $\boldsymbol{W} = \{b_0, b_{0.2}, b_{0.4}, b_{0.6}, b_{0.8}, b_1\}$, where b_n is a world where the bananas are ripe to degree n, i.e. $\varphi(\langle b_n, B\rangle) = n$. The actual world in this context is $b_{0.6}$, since that's how ripe the shipment of bananas is. At first, the store-owner thinks all amounts of ripeness are equally likely, such that $Pr^S_{b_{0.6}}(\{b_n\}) = 1/6$. Clive's proof, where he presents the store-owner with a statistical distribution of ripeness in past shipments, has the effect of revising the context with the constraint $(q, \mathcal{X})$, where $\mathcal{X} = \{\{b_0, b_{0.2}, b_1\}, \{b_{0.4}\}, \{b_{0.6}\}, \{b_{0.8}\}\}$ and $q(\{b_0, b_{0.2}, b_1\}) = 0$, $q(\{b_{0.4}\}) = 0.25$, $q(\{b_{0.6}\}) = 0.5$ and $q(\{b_{0.8}\}) = 0.25$. By KL divergence, the amount of information which the store-owner stands to gain from this proof is $I^S_{b_{0.6}}(q, \mathcal{X}) \approx 1.08$. Claude's proof, on the other hand, where he plucks a banana from the shipment and shows it to the store-owner, has the effect of revising the context with the constraint $(r, \mathcal{X})$, where $\mathcal{X} = \{\{b_{0.6}\}, \overline{\{b_{0.6}\}}\}$ and r is such that $r(\{b_{0.6}\}) = 1$ and $r(\overline{\{b_{0.6}\}}) = 0$. By KL divergence, the amount of information conveyed by Claude's proof is $I^S_{b_{0.6}}(r, \mathcal{X}) \approx 2.58$. This is the precise sense in which Claude's proof is capable of eliminating more uncertainty than Clive's, despite both being 0.6-strength proofs of the proposition B. Claude's proof takes all of the store-owner's probability and concentrates it on a single outcome, whereas Clive's proof leaves a greater portion of the store-owner's probability unmoved.

By using KL divergence, we can give a full statement of the maxim of quantity which takes into consideration different distributions of probability:

MAXIM OF QUANTITY A speaker who wishes to perform some speech act which revises the context $\mathfrak{C}$ with a constraint $(q, \mathcal{X})$ should first ensure that, for each $i \in G$ such that $Pr^i_@(X) > 0$ for all $X \in \mathcal{X}$, $I^i_@(q, \mathcal{X}) \in I[x, y]$

The inappropriateness of Clive's proof can be understood as a violation of this maxim. The storeowner expects to receive more information about the shipment of bananas than Clive's report of various frequencies gives.

This is not to say that Clive was wrongheaded in his entire approach; if the frequencies reported delivered a 'pointy' enough distribution, they could eliminate enough certainty for the storeowner to be satisfied. As $q(\{b_{0.6}\})$ increases, with all else being held equal, the value of $I^i_w(q, \mathcal{X})$ increases too, giving Clive's proof a better chance of being appropriate insofar as $I^i_w(q, \mathcal{X})$ lies within the information range $I[x, y]$. If Clive were to give a proof which showed that 98% of previous shipments contained bananas which were ripe to degree 0.6, then the storeowner would stand to gain approximately 2.4 units of information, bringing it closer in line to Claude's proof with respect to informativeness for the storeowner.

2.2.2 Maxim of Relevance

What does it take for a proof to be relevant? As a naïve first suggestion, we might say that a proof is relevant if and only if it changes the presuppositions of at least one speaker in G. This mirrors a classic Bayesian approach to the idea of relevance, where a proposition E is counted as evidence which is relevant to a proposition P just in case the conditional probability $Pr(P|E)$ is different to the unconditional probability $Pr(P)$.[17] That is, E counts as evidence which is relevant to P just in case conditioning on E would change the probability you assign to P. Due to issues which were discussed at length in §1.4, this notion of relevance is unhelpful in cases where $Pr(E) = 0$, since then the conditional probability $Pr(P|E)$ is undefined. We might charitably extend this Bayesian notion of relevance and suggest that a proposition E is relevant to P just in case $Pr(P) \neq Pr[(q, \mathcal{X})]$ in a revision where $\mathcal{X} = \{..., E, ...\}$ and $q(E) = 1$. According to this extension, E is relevant to P just in case learning E would change the probability which you assign to P, either through conditioning or imaging as the case demands. In this same vein, we might suggest that a proof is relevant to some context just in case its being performed induces some change to that context. However, the problem with this approach is that a proof like CEMPEDAK would still be considered relevant, given that it changes the speakers' presuppositions by adding the fruit's being a cempedak to the common ground. What we need to account for the Gricean maxim of relevance is some framework which allows us to distinguish relevant changes to the context from irrelevant changes to the context.

---

[17] See Horwich (1982, p. 51)

This notion of relevance is one sort among many. There is relevance between evidence and hypotheses, relevance between the premises and conclusions of an argument, relevance between actions and some goal, and so on. Cohen (2002) breaks these varieties of relevance into two main categories: conversational and non-conversational relevance. Conversational relevance is our target phenomena and it concerns whether a speech act makes a meaningful contribution to the goals of an inquiry. Non-conversational relevance, on the other hand, includes things like the relevance of evidence to a particular hypothesis, which is a relation we might think exists even in the absence of agents who believe or appreciate this connection. For instance, footprints which the murderer left outside the shed are relevant to the hypothesis that the gardener is the murderer, even if no-one is aware that these footprints exist. Cohen (2002, p. 284) makes the suggestion that all kinds of relevance can be reduced to the relevance of a proposition *to* some question, problem or issue. A speech act is irrelevant insofar as it doesn't make progress towards the goal of the conversation in which it is performed and a piece of evidence is irrelevant to some hypothesis insofar as it fails to confirm or undermine that hypothesis.

What is wrong with the proof in CEMPEDAK is that a specific goal of inquiry has been established by the two friends (i.e., whether the fruit smells of garlic) and the proof confirms a proposition which contributes little to no progress towards that goal of inquiry. If we can find some way of representing the 'open questions' under a given context's consideration, then we will have some tool with which we can judge whether a proof makes a significant contribution to answering these questions. In this section, I will critique and build upon a classic theory of questions from Groenendijk and Stokhof (1984, G&S hereafter) and show how it can be used to service this very aim.

G&S's (1984, pp. 213-214) central contribution to the analysis of questions comes from representing them as functions from possible worlds to sets of worlds called answers. The idea here is that every world gets associated with a set of worlds (or proposition) which is the correct answer to the question at that world. For instance, the question "Will it rain tomorrow?" is viewed as a function which takes possibilities where it rains to the set R of all possibilities where it rains and possibilities where it doesn't rain to the set $\overline{R}$ of all possibilities where it doesn't rain. This function can also be represented as a partition on the space of possible worlds $\boldsymbol{W}$, denoted $\boldsymbol{W}/\boldsymbol{Q}$, where each cell $A \in \boldsymbol{W}/\boldsymbol{Q}$ is an answer to the question $\boldsymbol{Q}$. If $w \in A$, then

A is called the 'semantic answer' to the question partition $\boldsymbol{W}/\boldsymbol{Q}$ at world w. If an answer A contains the actual world, @, then A is called the 'true semantic answer' to the question partition $\boldsymbol{W}/\boldsymbol{Q}$.

Adding a question partition to each context will allow us to represent the aims of a conversation's inquiry at any given time. This will be the basis for understanding whether a proof makes a relevant contribution to the conversation. If a context $\mathfrak{C}$ contains the set of worlds $\boldsymbol{W}$, then the question under consideration in $\mathfrak{C}$ is some question partition of the form $\boldsymbol{W}/\boldsymbol{Q}$. When the question partition associated with a context is $\boldsymbol{W}/\boldsymbol{Q} = \{A_1, ..., A_n\}$, this means that speakers in G are wondering which event out of the list $A_1, ..., A_n$ holds. Much like other members of the n-tuple which form a context, we should expect that a question partition changes as a conversation evolves. New questions might be posed and old questions might be taken out of consideration. For instance, a detective might first be wondering whether Mr. X or Madame Y stole the Maltese Falcon and then late into the investigation of this theft they might also begin to wonder whether the stolen falcon is genuine or a facsimile. We can represent changes such as these by adding dynamic rules which dictate how a context's question partition changes relative to further questioning. To introduce these dynamic rules, I will utilise *G&S*'s (1984, p. 219) two-place operations on question partitions, which are defined as follows:

$$\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R} = \{X \cap Y : X \in \boldsymbol{W}/\boldsymbol{Q} \text{ and } Y \in \boldsymbol{W}/\boldsymbol{R} \text{ and } X \cap Y \neq \varnothing\}$$

$$\boldsymbol{W}/\boldsymbol{Q} \sqcup \boldsymbol{W}/\boldsymbol{R} = \{Z : Z \neq \varnothing \text{ and for some } \boldsymbol{X} \subseteq \boldsymbol{W}/\boldsymbol{Q} \text{ and } \boldsymbol{Y} \subseteq \boldsymbol{W}/\boldsymbol{R}, Z = \bigcup \boldsymbol{X} = \bigcup \boldsymbol{Y} \text{ and there is no } Z' \neq \varnothing \text{ such that for some } \boldsymbol{X}' \subseteq \boldsymbol{W}/\boldsymbol{Q} \text{ and } \boldsymbol{Y}' \subseteq \boldsymbol{W}/\boldsymbol{R}, Z' = \bigcup \boldsymbol{X}' = \bigcup \boldsymbol{Y}' \text{ and } Z' \subset Z\}$$

The operator $\sqcap$ takes any two question partitions $\boldsymbol{W}/\boldsymbol{Q}$ and $\boldsymbol{W}/\boldsymbol{R}$ and returns the question partition ($\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$) whose cells are all non-empty intersections of cells from $\boldsymbol{W}/\boldsymbol{Q}$ and $\boldsymbol{W}/\boldsymbol{R}$. Imagine that $\boldsymbol{W}$ is a sheet of paper covered in small dots representing possible worlds and each question $\boldsymbol{Q}$ is a mesh which can be used to screen print some question partition $\boldsymbol{W}/\boldsymbol{Q}$ on to the paper, breaking the dots up into a series of sections. The process described by $\sqcap$ involves printing both input meshes on the sheet of paper (one after the other, in either order) so that they overlap. The operator $\sqcup$ takes any two question partitions $\boldsymbol{W}/\boldsymbol{Q}$ and $\boldsymbol{W}/\boldsymbol{R}$ and returns a partition whose cells are the smallest unions of cells common to both. Extending the printmaking metaphor, $\sqcup$ describes a process where the printmaker uses two meshes ($\boldsymbol{Q}$ and $\boldsymbol{R}$) to make a new mesh which only allows ink to flow through lines which are common to both input meshes and then uses this new mesh to print a pattern onto $\boldsymbol{W}$, yielding the question partition $\boldsymbol{W}/\boldsymbol{Q} \sqcup \boldsymbol{W}/\boldsymbol{R}$

When a new question $\boldsymbol{W}/\boldsymbol{R}$ is successfully posed in a context $\mathfrak{C}$ which has $\boldsymbol{W}/\boldsymbol{Q}$ as its question partition, let this result in a new context $\mathfrak{C}'$, which is exactly like $\mathfrak{C}$ except its question partition is $\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$. This rule captures the idea that adding a question involves creating distinctions within answers. Take the detective who first wonders whether X or Y stole the falcon, and then begins to wonder whether the falcon is real (R) or fake ($\neg R$). This compounding of questions creates finer-grained goals for inquiry, since events like $X \cap \neg R$ become answers in the combined question $\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$ (see Figure 2.1).

Sometimes, during the course of an inquiry, questions which were previously taken as central might be dropped. When a change like this happens, we should expect that the question partition becomes coarser-grained: erasing distinctions between answers which previously existed. The detective might stop wondering whether X or Y stole the falcon in the case where it's a facsimile, perhaps because there's no need to identify the thief of a false falcon. If an agent successfully calls for an end to the inquiry as to whether $A_1, ..., A_n$ in a context $\mathfrak{C}$ whose question partition $\boldsymbol{W}/\boldsymbol{Q}$ contains $A_1, ..., A_n$ as cells, let this result in a new context $\mathfrak{C}'$ which is exactly like $\mathfrak{C}$ except its question partition is $\boldsymbol{W}/\boldsymbol{Q} \sqcup \boldsymbol{W}/\boldsymbol{Q}_{\cup_i A_i}$, where $\boldsymbol{W}/\boldsymbol{Q}_{\cup_i A_i}$ is exactly like $\boldsymbol{W}/\boldsymbol{Q}$ except that it contains the union $A_1 \cup ... \cup A_n$ in place of $A_1, ..., A_n$ as cells. This dynamic rule captures the idea that calling for an end to some inquiry involves erasing distinctions between answers. The detective ends their inquiry into whether X or Y stole the fake falcon by replacing the answers $X \cap \neg R$ and $Y \cap \neg R$ with the broader answer $\neg R$. This stage in the detective's inquiry is represented by the bottom-right partition in Figure 2.1.

I should also note that while G&S (1984) introduced the operators $\sqcap$ and $\sqcup$, they did not introduce them for the purpose of conversational kinematics. Indeed, G&S (1984, p. 220) didn't envision any kind of role at all for the operator $\sqcup$ claiming that "it has no straightforward linguistic analogue". However, using these two operators to outline the dynamics of questions in a context shows that $\sqcup$ has a crucial linguistic role when it comes to the practice of *removing* the relevance of a given set of answers. It explains why it is possible for conversations to evolve in such a way that propositions which were once relevant are no longer of great importance to the inquirers, like how the detective stops caring whether X or Y is the thief *per se* and only whether X is the thief of a real falcon, Y is the thief of a real falcon or the falcon is fake.

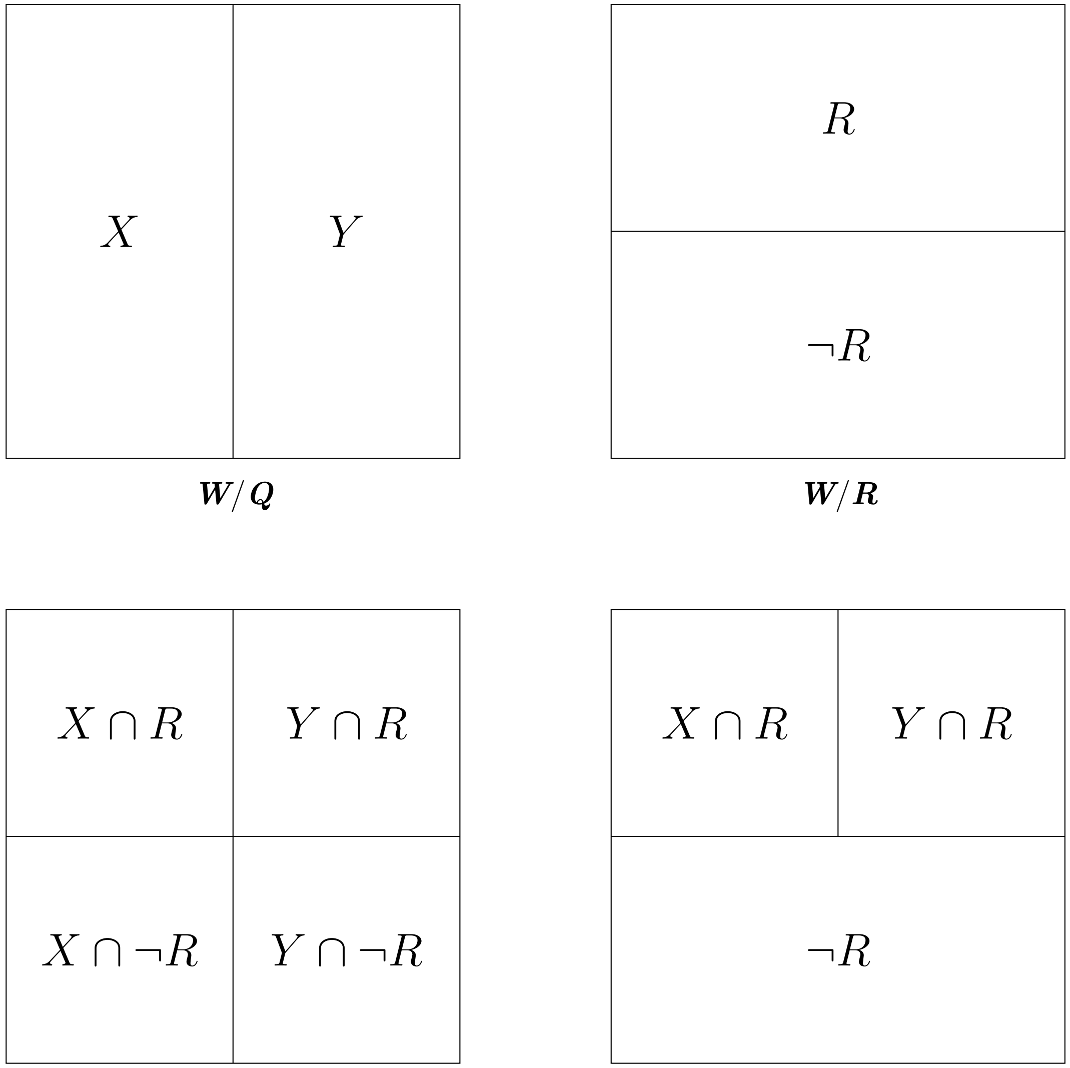

$\boldsymbol{W}/\boldsymbol{Q}$ $\boldsymbol{W}/\boldsymbol{R}$

$\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$ $(\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}) \sqcup (\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R})_{(X \cap \neg R) \cup (Y \cap \neg R)}$

Figure 2.1: A series of question partitions for the detective.

On its own, this picture where question partitions break the sample space $\boldsymbol{W}$ into sets of answers isn't sufficient to determine whether a speech act is relevant or appropriate in a given context. As G&S (1984, p. 217) suggest, there is a difference between being an answer to the question $\boldsymbol{W}/\boldsymbol{Q}$ and *answering* that question for someone. In the words of Hintikka (1978, p. 290), the notion of an answer "does not depend only on the logical and semantical status of the question and its putative answer [...] but also on the state of knowledge of the questioner at the time he asks the question." An answer to $\boldsymbol{W}/\boldsymbol{Q}$ is merely one of its members, but answering $\boldsymbol{W}/\boldsymbol{Q}$ for an agent i involves taking i from a state of ignorance about these answers to a state of belief where they take one of them to be true. G&S (1984, p. 224) use $\boldsymbol{D}_{i,w} \subseteq \boldsymbol{W}$ to denote the doxastic state of an agent i at a given world w. Intuitively, these doxastic states are the set of worlds are consistent with the beliefs which i has at the possible world w. Placing these doxastic states within the probabilistic epistemic logic developed above, I will take an agent i's doxastic state at some world w to be $\boldsymbol{D}_{i,w} = \{w : Pr^i_w(\{w\}) > 0\}$. An agent's doxastic state at the actual world @ is just their context set $\boldsymbol{C}_i$, which is the set of worlds that are consistent with i's actual beliefs. G&S (1984, p. 223) say that a doxastic state $\boldsymbol{D}_{i,w}$ offers an answer the question $\boldsymbol{W}/\boldsymbol{Q}$ just in case $\boldsymbol{D}_{i,w} \subseteq A$ for some cell $A \in \boldsymbol{W}/\boldsymbol{Q}$. In other words, at world w an agent i offers an answer to some question just in case they believe one of its answers is true.

To answer $\boldsymbol{W}/\boldsymbol{Q}$ for an agent i, you must say or show them something that leads them to a doxastic state $\boldsymbol{D}_{i,w}$ which offers an answer to $\boldsymbol{W}/\boldsymbol{Q}$. In particular, G&S (1984, p. 226) suggest that a proposition P is an answer to the question $\boldsymbol{W}/\boldsymbol{Q}$ for the agent i at some world w just in case revising the doxastic state $\boldsymbol{D}_{i,w}$ with the information that P would offer an answer to $\boldsymbol{W}/\boldsymbol{Q}$. The updating rule for doxastic states provided by G&S (1984, p. 227) here is a slight variation on Stalnaker's intersective rule for updating context sets, defined as follows:

$$\boldsymbol{D}_{i,w}(P) = \begin{cases} \boldsymbol{D}_{i,w} \cap P, & \text{if } \boldsymbol{D}_{i,w} \cap P \neq \varnothing \\ \boldsymbol{D}_{i,w}, & \text{otherwise} \end{cases}$$

In adopting this Stalnakerian updating rule, G&S commit themselves to the same dogma which has impeded theories of context revision. Updates to some doxastic state which require the accommodation of incompatible information are cast aside, despite the fact that we regularly change our beliefs in light of incompatible evidence. The position is especially problematic in a theory of questions or goals for inquiry. As it currently stands, the theory allows for an agent to accept that P would answer some question $\boldsymbol{W}/\boldsymbol{Q}$ (if $P \subseteq A \in \boldsymbol{W}/\boldsymbol{Q}$) but nonetheless dismiss P *as* an answer because they don't believe it (if $\boldsymbol{D}_{i,w} \cap P = \varnothing$). Thus, if the detective believes that X stole the falcon, they are able to dismiss all countervailing evidence, including any evidence which bears on the authenticity of the falcon but also implies that Y is the thief, such as photographs of Madame Y painting a layer of bronze-tinted paint on a clearly fake falcon. This is absurd, especially given that the detective may still be actively investigating whether the stolen falcon is a fake.

We can resolve this flaw in G&S's theory by re-invoking the more general account of belief revision developed in the first half of this thesis. Recall that we took the revision of a probability measure Pr^i_w with the information that P (denoted $Pr^i_w[P]$) to be either (i) the image of Pr^i_w on P or (ii) the conditionalization of Pr^i_w on P in accordance with whether the latter is well-defined. In either case, the selection function σ_i which is used to facilitate this revision can also be used to characterise the change from a doxastic set $\boldsymbol{D}_{i,w}$ corresponding with Pr^i_w to the new doxastic state $\boldsymbol{D}_{i,w}(P)$ corresponding with $Pr^i_w[P]$:

$$\boldsymbol{D}_{i,w}(P) = \bigcup_{w \in \boldsymbol{D}_{i,w}} \sigma_i(w, P) = \{w : Pr^i_w[P](\{w\}) \neq 0\}$$

In other words, the revision of $\boldsymbol{D}_{i,w}$ with respect to P is the subset of P-worlds which, according to σ_i, are willed probability from worlds in $\boldsymbol{D}_{i,w}$. In light of this adjustment to G&S's theory, we can say that the proposition P answers $\boldsymbol{W}/\boldsymbol{Q}$ for i at world w just in case $\boldsymbol{D}_{i,w}(P)$ offers an answer to $\boldsymbol{W}/\boldsymbol{Q}$. Thus, the detective won't dismiss evidence that the falcon is fake in cases where he is convinced that it is real but still cares whether it is or isn't. Due to the (Success) constraint on all selection functions, which holds that $\sigma(w, P) \subseteq P$, it will always be the case that $\mathbf{D}_{i,w}(P) \subseteq A$ whenever $P \subseteq A$. So, if the detective is in a doxastic state $\boldsymbol{D}_{i,w} \subseteq R$, but he observes a photograph of Madame Y painting the clearly fake falcon (call this observation P), this will still offer an answer to the detective as to whether the falcon is fake because $\boldsymbol{D}_{i,w}(P) \subseteq \neg R$. According to this much improved account, whether some proposition satisfies a questioner still depends on that questioner's state of knowledge but it is not so dependent as to cast aside all answers which are incompatible with that knowledge.

You might worry that this revised theory goes too far. If any $P \subseteq A$ such that $A \in \boldsymbol{W}/\boldsymbol{Q}$ brings agents to some doxastic state $\boldsymbol{D}_{i,w}(P)$ which offers an answer to $\boldsymbol{W}/\boldsymbol{Q}$, then how is it possible that the relevance of an answer depends on the beliefs of the questioner? When this condition holds, each semantic answer A automatically counts as an answer in the more restricted agent-relative sense. However, it is not generally the case that these semantic answers (or their subsets) are the *only* propositions which could be used to answer a question for some agent. It might be the case that an agent can find an answer to the question $\boldsymbol{W}/\boldsymbol{Q}$ by learning the truth of some proposition which isn't a semantic answer (or a subset of one) but one which, when combined with other aspects of their doxastic state, nonetheless takes them to a position where they believe one of the semantic answers to $\boldsymbol{W}/\boldsymbol{Q}$. Suppose that the detective is considering the question $\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$ depicted in the bottom-left square of Figure 2.1. They might be in a state where they know that either the Falcon is real and Mr. X stole it or it is fake, in which case Madame Y stole it. Let us call this scenario w and note that $Pr^i_w((X \cap R) \cup (Y \cap \neg R)) = 1$, where i is the detective. When he learns the falcon is fake, the detective *also* gets an answer to the question $\boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$, even though $\neg R$ is not the subset of any $A \in \boldsymbol{W}/\boldsymbol{Q} \sqcap \boldsymbol{W}/\boldsymbol{R}$. This is because $\boldsymbol{D}_{i,w}(\neg R)$ is a subset of $(Y \cap \neg R)$, which is an answer to the detective's question.

At this point, we have a three-way relation between propositions, doxastic states and question partitions which we will call the 'relation of answerhood'. A proposition P serves an an answer to the question $\boldsymbol{W}/\boldsymbol{Q}$ as it applies to $\boldsymbol{D}_{i,w}$ if and only if $\boldsymbol{D}_{i,w}(P) \subseteq A$, for some $A \in \boldsymbol{W}/\boldsymbol{Q}$. Where does this notion get us in terms of explicating the target relation of relevance between proofs and contexts? G&S (1984, p. 242) give the natural suggestion that a proposition P counts as relevant to a question $\boldsymbol{W}/\boldsymbol{Q}$ according to some speaker i at world w, if and only if learning P excludes at least one answer for that speaker—i.e. if $\boldsymbol{D}_{i,w}(P) \cap A = \varnothing$ for some $A \in \boldsymbol{W}/\boldsymbol{Q}$ such that $\boldsymbol{D}_{i,w} \cap A \neq \varnothing$. This falls in line with the view sketched at the beginning of this section, that a proposition is relevant if it makes some inroads into the program of inquiry established by a given question. However, given that we are working with a more fine-grained picture of conversations than G&S, which allows for probabilistic assertions as well as unqualified ones, I disagree that excluding an answer is the only kind of partial contribution a speech act can make to answering a question. In general, changing an answer's probability is enough to count as a relevant contribution to the solution of that question. This brings us full circle. Treating relevance this way combines the naïve Bayesian conception of relevant evidence with the question-based notion of relevance proposed by G&S (1984) and Cohen (2002).

Here is the full statement of this definition for the relevance of a proof (or any speech act) to a speaker in a given context:

Relevance If at time t a speech act is performed in context $\mathfrak{C}$ containing the question $\boldsymbol{W}/\boldsymbol{Q}$, then this speech act is relevant to the conversation for speaker i if and only if the resulting context $\mathfrak{C}[(q, \mathcal{X})]$ just after t is such that $Pr^i_@(A) \neq Pr^i_@[(q, \mathcal{X})](A)$ for some $A \in \boldsymbol{W}/\boldsymbol{Q}$.

Using this definition to characterise the irrelevance in cempedak, we see that the proof falls short because it changes the questioner's confidence in the proposition that the fruit is a cempedak *without* changing their confidence in either the proposition that it smells like garlic or it doesn't, which is the ostensive goal of their inquiry. When one of the friends asks whether the fruit smells of garlic, they add the question partition $\boldsymbol{W}/\boldsymbol{G} = \{G, \overline{G}\}$ to the conversational scoreboard. Given his limited background knowledge of exotic fruits, let this questioner i assign equal likelihood to both events in this question partition, such that $Pr^i_@(G) = Pr^i_@(\overline{G}) = 0.5$. When the other speaker j proves beyond all doubt that the fruit is a cempedak, i thereby replaces $Pr^i_@$ with a new probability measure $Pr^i_@[C]$, where C is the proposition that the fruit is a cempedak. Now, if the fruit being a cempedak has no bearing for i on whether it smells of garlic—that is, if $Pr^i_@[C](G) = Pr^i_@(G)$ and $Pr^i_@[C](\overline{G}) = Pr^i_@(\overline{G})$—then j's proof is irrelevant to whether G for the speaker i. Alternatively, if i knew that durians (and not cempedaks) smell of garlic, then j's proof would indeed be relevant to i in the given conversation. This will be the case when the selection function σ_i is such that $\sigma_i(w, C) \subseteq \overline{G}$, for all $w \in \boldsymbol{C}_i$.

Conclusion

I have defended a theory of proof and context which pays careful attention to the ease and difficulty with which we assert things. The dynamics for context revision given in Chapter 1 show how a diverse range of assertions are possible in any given conversation, whereas the requirements on proofs explored in Chapter 2—which determine how strong, relevant and informative a successful proof must be—show how hard it is for a speaker to ensure that what they say will be accommodated.

The benefit of adopting such a theory of conversations is that it makes way for a plausibly fallibilist view of our linguistic and epistemic endeavours. By taking contexts to be genuinely revisable, we can discharge an assumption underlying most formal theories of inquiry: that having a degree of belief in, asserting or deducing some proposition inexorably involves getting things right. Update procedures like Stalnaker's intersective rule and conditionalization only work when applied to context sets or probability measures which contain some solid foundation of true information which is then confirmed by all future assertions and observations. If something is said or seen which is incompatible with a context set or probability measure, then these two updating procedures fall silent. Those who unknowingly accept falsehoods are stuck in their tracks and cut off from the spoils of further inquiry.

By incorporating more flexible rules for the revision of probabilities and contexts, I have developed a model of inquiry which tolerates mistakes. Speakers are able to get many things wrong without being prevented from eventually learning the truth of a given matter. For every false presupposition garnered, there are also opportunities for speakers to withdraw them from the common ground. The theory of proof elaborated in Chapter 2 shows that not all procedures for overruling false presuppositions are on par. A proof only succeeds at bringing these erring speakers to the truth if it is able to meet them where they stand, eliminating the particular doubts which they have in a way which is acceptable according to the rules of their conversation.

Bibliography

[1] *Criminal Trial Courts Bench Book.* Judicial Commission of New South Wales, 2002.

[2] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[3] Ronald J. Allen. Debate: Legal probabilism–a qualified rejection: A response to hedden and colyvan. *Journal of Political Philosophy*, 28(1):117–128, 2020.

[4] Alan Ross Anderson and Nuel D. Belnap. *Entailment: The Logic of Relevance and Necessity, Vol. I.* Princeton University Press, Princeton, 1975.

[5] J. L. Austin. *How to Do Things with Words: The William James Lectures Delivered in Harvard University in 1955.* Oxford University Press UK, Oxford, England, 1962.

[6] Libor Běhounek, Petr Cintula, and Petr Hájek. Introduction to mathematical fuzzy logic. In Petr Cintula, Petr Hájek, and Carles Noguera, editors, *Handbook of Mathematical Fuzzy Logic*, volume 1. College Publications, 2011.

[7] Rudolf Carnap and Yehoshua Bar-Hillel. An outline of a theory of semantic information. Technical Report 247, Massachusetts Institute of Technology, Cambridge, MA, 1952.

[8] Laurence Jonathan Cohen. *The Probable and the Provable.* Clarendon Press, Oxford, 1977.

[9] L.J. Cohen. Some steps towards a general theory of relevance. In James Logue, editor, *Knowledge and Language*, pages 279–292. Kluwer Academic Publishers, 2002.

[10] Bruno de Finetti. *Theory of Probability.* Wiley, 2017[1970].

[11] Dylan Dodd. Belief and certainty. *Synthese*, 194(11):4597–4621, 2017.

[12] Fred Dretske. The pragmatic dimension of knowledge. *Philosophical Studies*, 40(3):363–378, 1981a.

[13] Fred I. Dretske. Epistemic operators. *Journal of Philosophy*, 67(24):1007–1023, 1970.

[14] Fred I. Dretske. *Knowledge and the Flow of Information.* MIT Press, Stanford, CA, 1981b.

[15] Jeffrey Dunn. Accuracy, verisimilitude, and scoring rules. *Australasian Journal of Philosophy*, 97(1):151–166, 2019.

[16] Kenny Easwaran. Regularity and hyperreal credences. *Philosophical Review*, 123(1):1–41, 2014.

[17] Benjamin Eva and Stephan Hartmann. The logic of partial supposition. *Analysis*, (2):215–224, 2021.

[18] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. *Reasoning About Knowledge.* MIT Press, Cambridge, Mass., 2003.

[19] Eduardo Fermé and Sven Ove Hansson. *Belief Change: Introduction and Overview.* Springer Verlag, 2018.

[20] Rohan French. A sequent calculus for urn logic. *Journal of Logic, Language and Information*, 24(2):131–147, 2015.

[21] Melissa Fusco. Jeffrey imaging revisited. *Analysis*, 83(3):447–464, 2023.

[22] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge and London, 1988.

[23] Peter Gärdenfors. Imaging and conditionalization. *Journal of Philosophy* 79(12):747-760, 1982.

[24] Leonard Goddard and Richard Routley. *The Logic of Significance and Context*. Scottish Academic Press, Edinburgh and London, 1973.

[25] Alvin I. Goldman. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(November):771–791, 1976.

[26] Daniel Greco. How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29(1):179–201, 2015.

[27] Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1989.

[28] J.AG. Groenendijk and M.J.B. Stokhof. *Studies on the semantics of questions and the pragmatics of answers*. PhD thesis, University of Amsterdam, 1984.

[29] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.

[30] Susan Haack. *Evidence Matters: Science, Proof, and Truth in the Law*. Cambridge University Press, New York, NY, 2014.

[31] Joshua Habgood-Coote. Group inquiry. *Erkenntnis*, 87(3):1099–1123, 2022.

[32] Alan Hájek. What conditional probability could not be. *Synthese*, 137(3):273–323, 2003.

[33] Harold W. Hake. The perception of frequency of occurrence and the development of "expectancy" in human experimental subjects. In Henry Quastler, editor, *Information Theory in Psychology*, pages 257–281. The Free Press, Publishers, Glencoe, Illinois, 1959.

[34] Brian Hedden and Mark Colyvan. Legal probabilism: A qualified defence. *Journal of Political Philosophy*, 27(4):448–468, 2019.

[35] J. D. Heydon. *Cross on evidence*. LexisNexis Australia, Fourteenth Australian edition, 2024.

[36] Jaakko Hintikka. *Knowledge and Belief: an introduction to the logic of both notions*. Cornell University Press, 1962.

[37] Jaakko Hintikka. *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic.* Clarendon Press, Oxford, 1973.

[38] Jaakko Hintikka. Answers to questions. In Henry Hiż, editor, *Questions.* Reidel, 1978.

[39] Jaakko Hintikka. Impossible possible worlds vindicated. In Esa Saarinen, editor, *Game-Theoretical Semantics*, pages pp. 367–379. D. Reidel, 1979.

[40] Paul Horwich. *Probability and Evidence.* Cambridge University Press, Cambridge, 1982.

[41] Elizabeth G. Jackson. The relationship between belief and credence. *Philosophy Compass*, 15(6):1–13, 2020.

[42] Richard Jeffrey. Dracula meets wolfman: Acceptance vs. partial belief. In Marshall Swain, editor, *Induction, acceptance, and rational belief*, pages 157–185. Reidel, 1970.

[43] Richard Jeffrey. Probabilism and induction. *Topoi*, 5(1):51–58, 1986.

[44] Richard C. Jeffrey. *The Logic of Decision.* University of Chicago Press, New York, NY, USA, 1965.

[45] James M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.

[46] James M. Joyce. *The Foundations of Causal Decision Theory.* Cambridge University Press, 1999.

[47] James M. Joyce. Accuracy, ratification and the scope of epistemic consequentialism. In Kristoffer Ahlstrom-Vij and Jeffrey Dunn, editors, *Epistemic Consequentialism*, pages 240–266. Oxford University Press, 2018.

[48] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[49] Henry Ely Kyburg. *Epistemology and Inference.* University of Minnesota Press, 1983.

[50] G.W. Leibniz. *New Essays on Human Understanding.* Cambridge University Press, 1996[1765].

[51] Isaac Levi. Possibility and probability. *Erkenntnis*, 31(2-3):365–86, 1989.

[52] Isaac Levi. *Mild Contraction: Evaluating Loss of Information Due to Loss of Belief.* Oxford University Press, New York, 2004.

[53] David Lewis. *Convention: A Philosophical Study.* Wiley-Blackwell, Cambridge, MA, USA, 1969.

[54] David Lewis. *Counterfactuals.* Blackwell, Malden, Mass., 1973.

[55] David Lewis. Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85(3):297–315, 1976.

[56] David Lewis. Truth in fiction. *American Philosophical Quarterly*, 15(1):37–46, 1978.

[57] David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359, 1979.

[58] David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.

[59] Hanti Lin. *The Open Handbook of Formal Epistemology*, chapter Belief Revision Theory. PhilPapers Foundation, 2019.

[60] Patrick Maher. *Betting on Theories.* Cambridge University Press, Cambridge, 1993.

[61] David Makinson. Propositional relevance through letter-sharing. *Journal of Applied Logic*, 7(4):377–387, 2009. Special Issue: Formal Models of Belief Change in Rational Agents.

[62] G.E. Moore. Proof of an external world. In Thomas Baldwin, editor, *Selected Writings*. Routledge, London, 1993.

[63] Sarah Moss. *Probabilistic Knowledge.* Oxford University Press, Oxford, United Kingdom, 2018.

[64] Sarah Moss. Knowledge and legal proof. In John Hawthorne Tamar Szabó Gendler and Julianne Chung, editors, *Oxford Studies in Epistemology*, volume 7. Oxford University Press, 2023.

[65] Michael Nielsen. A harder dilemma for partial subjunctive supposition. *Analysis*, 82(4):585–592, 2022.

[66] Graham Oddie. *Likeness to Truth*. Reidel, Dordrecht and Boston, 1986.

[67] Graham Oddie. What accuracy could not be. *British Journal for the Philosophy of Science*, 70(2):551–580, 2019.

[68] Graham Oddie and Gustavo Cevolani. Truthlikeness. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2022 edition, 2022.

[69] Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press UK, New York, NY., 2016.

[70] Karl Raimund Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge Classics. Routledge, 2002[1963].

[71] Graham Priest. Sylvan's box: A short story and ten morals. *Notre Dame Journal of Formal Logic*, 38(4):573–582, 1997.

[72] Graham Priest. Paraconsistent logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 6. Kluwer Academic Publishers, 2002.

[73] Frank P. Ramsey. *Philosophical Papers*. Cambridge University Press, 1990.

[74] Veikko Rantala. Urn models: A new kind of non-standard model for first-order logic. *Journal of Philosophical Logic*, 4(4):455–474, 1975.

[75] Nicholas Rescher. *Dialectics: a classical approach to inquiry*. Ontos Verlag, 2007.

[76] Jonathan Roorda. *Revenge of Wolfman: A probabilistic explication of full belief*. (Unpublished Manuscript), 1995.

[77] Daniel Rothschild and Seth Yalcin. On the dynamics of conversation. *Noûs*, 51(2):24–48, 2017.

[78] Deniz Rudin. Uncertainty and persistence: A Bayesian update semantics for probabilistic expressions. *Journal of Philosophical Logic*, 47(3):365–405, 2018.

[79] Stephen Schiffer. Vagueness and partial belief. *Philosophical Issues*, 10(1):220–257, 2000.

[80] Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, paperbound edition edition, 1963[1949].

[81] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.

[82] Brian Skyrms. *Signals: Evolution, Learning, and Information*. Oxford University Press, Oxford, GB, 2010.

[83] Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, England, 2008.

[84] Nicholas J. J. Smith. Degree of belief is expected truth value. In Sebastiano Moruzzi and Richard Dietz, editors, *Cuts and Clouds. Vaguenesss, its Nature and its Logic*, pages 491–506. Oxford University Press, 2010.

[85] Nicholas J. J. Smith. Vagueness, uncertainty and degrees of belief: Two kinds of indeterminacy–one kind of credence. *Erkenntnis*, 79(5):1027–44, 2014.

[86] Nicholas J. J. Smith. Respecting evidence: Belief functions not imprecise probabilities. *Synthese*, 200(475):1–30, 2022.

[87] Wolfgang Spohn. *The Laws of Belief: Ranking Theory and its Philosophical Applications*. Oxford University Press, Oxford, 2012.

[88] Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457, 1973.

[89] Robert Stalnaker. *Context and Content: Essays on Intentionality in Speech and Thought.* Oxford University Press UK, Oxford, GB, 1999.

[90] Robert C. Stalnaker. A theory of conditionals. In Nicholas Rescher, editor, *Studies in Logical Theory*, pages 98–112. Blackwell, 1968.

[91] Robert C. Stalnaker. *Inquiry.* MIT Press, Cambridge, MA, 1984.

[92] G. C. Stine. Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29(4):249–261, 1976.

[93] Eric Swanson. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(2):121–146, 2016.

[94] Pavel Tichý. Verisimilitude revisited. *Synthese*, 38(2):175–196, 1978.

[95] Rafal Urbaniak and Marcello Di Bello. Legal Probabilism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy.* Metaphysics Research Lab, Stanford University, Fall 2021 edition, 2021.

[96] Timothy Williamson. *Knowledge and its Limits.* Oxford University Press, New York, 2000.

[97] Seth Yalcin. Context probabilism. In M. Aloni, editor, *18th Amsterdam Colloquium*, pages 12–21. Springer, 2012.

[98] Snow Zhang. On the dilemma for partial subjunctive supposition. *Analysis*, 84(3):576–592, 2024.