Introduction

A possibility is a way the world could be, including the way it actually is. One way to understand this is that a possibility is a way the world could be (or is) at a certain level of description, abstracting away from irrelevant particulars, so that if I oversleep tomorrow morning, all of the multitudinous ways in which that could happen count as one possibility. Another way to understand the notion that a possibility is a way the world could be (or is) is that it is a possible state of the world in all its particulars. When we say in everyday terms that the conditional statement in (1) addresses the possibility that Alex misses the train, our use of the singular definite description, *the possibility*, invokes the first notion of possibility.

(1) If Alex misses the train, Sam will pick him up.

But we cannot say this under the second conception of possibilities sketched earlier. In a typical context, (1) would not be used to state a contingency plan for a single, fully specified or fully conceived state of the world in which Alex misses the train – say, one in which Alex leaves his hotel fifteen seconds after 8:31 am with plenty of time, but is late because his taxi to the station gets stuck in traffic caused by an overturned pizza truck at the corner of Walnut Street and 18th Avenue, causing Alex to arrive at the station exactly ten minutes and thirty-six seconds after the train departed.[1] Rather, (1) would normally be used to state a contingency plan for the realization of any one of a vast number of possibilities, some of which could be considered variants of the one in which the taxi is stuck in traffic, and some of which would be very different, but all of which could be contextually reasonable ways in which Alex misses the train. The number of such ways the world could be is vast, even if we limit consideration to ones that are realistic and within the bounds of normal expectations. There are a multitude of basic patterns or templates of ways in which Alex could miss the train, each of which could be expressed as a reason he missed the train. "Alex's taxi to the station got held up in traffic" describes one such

[1] Of course, much more detail would have to be added – perhaps infinitely more – to truly describe a unique way the world could be.

pattern. But it could be realized in thousands of different ways, or millions of ways, depending on how finely we count them. Another would be, "Alex forgot his briefcase and had to back-track to his office," which could also be realized in innumerable different ways. Clauses of natural language can serve as descriptions of possibilities, but as such, they characterize vast sets of possibilities, not individual ones.

The present work pursues the consequences of the second view of possibilities in order to elucidate and defend the theory of modal interpretation developed by Angelika Kratzer (1981, 1991), and a kindred theory of the interpretation of ascriptions of beliefs and other propositionally contentful mental states, called propositional attitudes. Possibilities on this view are "first-order entities," that is, each possibility is a singular discrete entity in our conceptual and linguistic ontology, what is often called a possible world. Kratzer's theory involves quantifying over such possibilities based on their accord with a contextually determined set of propositions, which she called an ordering source. In recent years, Kratzer's theory has been argued to be faulty on a number of grounds. Lassiter (2011), in particular, alleges a multitude of logical and empirical faults stemming from its reliance on quantification over possible worlds, and argues for the superiority of a theory that interprets modal statements as expressing relationships involving degrees on a scale. And von Fintel and Iatridou (2008) argue that Kratzer's theory suffers from internal inconsistencies that arise in the case when there exists a set of possibilities that are all maximally in accord with the propositions in the ordering source. (The assumption that this case holds is called the Limit Assumption in modal interpretation.) The present work begins by arguing that, under the conception of possibilities sketched (the second one), most of Lassiter's critiques (which seem to implicitly assume the first one) fail to gain purchase, and the Limit Assumption fails to have the consequences that von Fintel and Iatridou attribute to it. These issues are taken up in Chapters 1 and 2.

Hintikka (1969) proposed that ascriptions of belief and other propositional attitudes are interpreted by quantifying over possibilities consistent with everything the experiencer believes, and a significant literature has adopted that proposal (see especially Heim [1992]). But this approach has a fatal flaw – it breaks down completely in implementation when the experiencer of the propositional attitude holds inconsistent beliefs. The view of possibilities advanced here permits an implementation of Hintikka's idea, using Kratzer's modal quantificational apparatus, which does not break down when beliefs are ascribed to someone with inconsistent beliefs. The approach developed handles the effects of presupposition accommodation in cases of belief ascription discussed by Heim (1992) without Heim's recourse to dynamic semantics for this purpose. These issues are taken up in Chapters 3 and 4, completing Part I of the book.

The interpretation of attitude ascriptions given in Part I does not involve positing any sort of event or state that would correspond to the mental or emotional condition of the experiencer of the attitude. This is fortunate since, while many ascriptions with attitude predicates such as *believe*, *think*, *suspect*, *be sure*, and *expect*, among others, do seem to characterize such an event or state, many don't; they simply describe the doxastic situation of the experiencer – in effect, where in logical space their beliefs or expectations reside. Katz (2000, 2003, 2008) argues that stative sentences, in general, do not characterize goings on or conditions of the world that we could package and reify as events or states. Chapter 5 contends with this claim, arguing that Katz has not convincingly ruled out associated states of the world for all stative attitude ascriptions. Toward developing this argument, Chapter 5 first summarizes a view of the role of events and states in semantics.

Contrary to Katz, Hacquard (2006, 2010) argues that attitude ascriptions are descriptions of localized goings on or conditions of the world, which we can reify as eventualities (events or states), and that the assumption of such an eventuality is a standard part of the conventional semantic interpretation of an ascription. Chapter 6 argues that there are, in fact, two types of attitude ascription. One type is a quantificational condition on possible worlds with no associated eventuality; the other posits an associated eventuality and uses the propositional content, expressed in terms of quantification over possible worlds, to characterize it, in much the way Hacquard's theory does. The rest of Part II is largely devoted to substantiating the claim that stative attitude ascriptions are of two types, with consequences that we can detect in a number of different domains. For example, (2) reports on an event of Alex saying something, and this event is characterized in part by the content of what is said, so that if the content were different, it would not be "the same" saying event.

(2) Alex said that Bill stole the ring.
(3) a. Alex thought that Bill stole the ring.
b. Alex long thought that Bill stole the ring.
c. Alex guiltily thought that Bill stole the ring.

But (3a) is essentially ambiguous. It can report on an event of Alex having a certain opinion over a period of time, with certain causal consequences for Alex's state of mind, and this is the interpretation involved when the main clause is modified by a temporal adverbial as in (3b), or the manner adverb in (3c). In this case, (3a), like (3b,c), reports on an eventuality going on in the mind of Alex. But (3a) has a more minimal interpretation, as one interpretive option, in which it is unlike (2) and (3b,c) in not characterizing an associated eventuality. Chapter 7 argues that that is the interpretation on which negation in (4) can be interpreted with the lower clause (called "Neg-Raising" in the literature).

(4) Alex doesn't think that Bill stole the file.

Chapter 8 explores what might differentiate two types of modal statement, epistemic ones based on what the speaker and addressee know, and root forms of modality based on desires or preferences, abilities or other circumstances of the world, or what the law or convention requires. A proposal by Hacquard (2006, 2010) for syntactic, semantic, and contextual differentiation of the two types of modal statement is examined. This approach comports to a large extent with the views that animate the present work, but with some departures in detail, which will be discussed in that chapter and at various earlier points throughout the book.

Building on the work by Diessel (2004), Chapter 9 argues that the more minimal interpretation of (3a) is acquired during child language development before the eventuality-characterizing interpretation, which requires (and reflects) the development of the child's Theory of Mind. In view of ideas put forth by Nespoulous et al. (1998), Chapter 9 argues that in the clinical condition of fluent jargonaphasia, the minimal interpretation of (3a) survives, but the eventuality-characterizing interpretation is lost, as part of a more general degradation of the capacity to refer consistently to objects, events, and states not present in the immediate environment.

Part I

Modality and propositional attitudes

Intensional statements are about ways the world could or might be, however it may be in actuality. This is exemplified by modal statements and reports of propositional attitude. Consider modal statements first.

If Kim and Jules walk by a student dorm, and Kim, noticing that Bill's window is lit, utters *Bill might be working in his room*, Kim is not saying that Bill *is* working in his room, nor that Bill *is not* working in his room. On a widely adopted view, she is saying something about possible ways the world could be, namely, that among those in which Bill's window is lit, and things are otherwise relevantly similar to what Kim and Jules know about the world, there are some (at least one) in which Bill is working in his room. If Kim, in the same circumstances, and mindful of Bill's work habits and aversion to wasting electricity, utters *Bill must be working in his room*, she is saying something different, and stronger, about all possible ways the world could be, namely, that Bill is working in his room in all of those possible worlds in which Bill has something like his actual habits, and which are in other relevant respects realistic enough to be worth considering.

On this analysis, the interpretation of modal statements invokes a set of possible ways the world could be, that is, a set of possibilities, consistent with what Kim and Jules know. But what is known to Kim and Jules significantly underdetermines what possibilities are in play, leaving too many of them. Assuming that Kim and Jules haven't seen Bill in a couple of days and have had no confirmation that he is still alive (not that they have any reason to suppose otherwise), they could not rule out the possibility that Bill is dead. So the set of possibilities consistent with what they know includes the ones in which Bill died the day before. These should be somehow excluded from play, if, in fact, Bill *didn't* die the day before, to prevent them from spuriously falsifying Kim's statement that Bill must be in his room when she has the typical sort of evidence for truthfully making that claim. In addition to making reference to the set of all possibilities consistent with what Kim and Jules know, it would be useful to order those possibilities according to their proto-typicality for, or relevance to, the truth of Kim's statements about Bill working in his room. Possibilities consistent with what Kim and Jules know, but in which Bill

happened to die the day before, could ordinarily be ruled out of bounds since they involve something so beyond normal expectation that Kim's statement can reasonably be said not to have been about them.

Being indifferently educated, Kim and Jules know something about history, but not a lot. As a result, the set of possibilities consistent with what Kim and Jules know is truly vast. Suppose that all possibilities in which the history of the world unfolded as it actually did, and in which Bill's light is on, are possibilities in which he is working in his room. But suppose at the same time that there are other possibilities in which Bill's light is on, in which a battle during the Wars of the Roses unfolded rather differently than it did, while everything known to Kim and Jules are as they know them to be, but Bill is not working in his room. Suppose that in one of the alternate histories, there is a survivor of the battle who, in actuality, was killed in action before having children. Suppose further that, in this counterfactual world, a descendant of the survivor is responsible for Bill leaving his room with the light on just as Kim and Jules pass by. Then Kim's *must* statement is falsified by this possibility. It's too harsh to say that Kim must be counted wrong about the necessity of Bill working in his room if some such distant change in history would have sufficed to make it not so. And it works the other way as well: if Kim said that Bill might not be in his room, despite the light being on and Bill being frugal in habit, it would be too generous to Kim to say that she was right if the only possibilities in which Bill is not in his room are ones in which a battle during the Wars of the Roses, unbeknownst to Kim, unfolded slightly differently than it did (so that Kim was right only by coincidence, as it were).

For reasons such as this, Angelika Kratzer added a second component to modal interpretation.[1] In addition to delineating the set of possibilities consistent with what the discourse participants know (the modal base), there is a relation that serves to rank possibilities according to how suitable they are for the evaluation of the modal statement, one way or the other. Among possibilities consistent with what Kim and Jules know, those in which Bill is still alive, and the broader history of the world unfolded as it actually did, are "live" possibilities for the evaluation of Kim's statement. Others can be disregarded, so they do not impinge on the truth of Kim's statement in a way that could make it spuriously true or false.

Recognizing that the set of all possibilities consistent with what Kim and Jules know is generally too vast to serve as the correct basis for a modal statement, a ranking of possibilities with respect to their epistemic relevance is one

[1] Kratzer didn't put it this way. She was concerned to capture the notion of the accessibility of possible worlds, from the literature on modal logic, in a linguistically viable and natural way. But the concern is the same; we have merely chosen to enter the discussion here through some examples, rather than through the literature on modal logic.

way of taming that set. We would like to rank possibilities according to their suitability for the evaluation of the modal statement at hand. A natural idea at this point is that we could use the ranking to establish a certain threshold; then, disregarding possibilities that are outside that threshold (too far out to consider), we could evaluate *Bill might be working in his room* as true just in case the "prejacent," *Bill is working in his room*, is true in some possibility better than that threshold, and we could evaluate *Bill must be working in his room* as true just in case the prejacent is true in all possibilities better than the threshold.

But this idea cannot be implemented in such a simple way, since it isn't clear that we could always rank any two arbitrarily chosen possibilities according to whether one is better than the other for the evaluation of a given modal statement. Following Kratzer, define the epistemic modal base to be the set of propositions $\{p_i: i \in \mathrm{I}\}$, characterizing what the speaker knows, for an index set I (which might be a subset of the natural numbers, if the modal base is finite, or at most countably infinite).[2] Each of these propositions is a set of possibilities, specifically, the set of possibilities in which the given proposition is true. So the modal generalizes over possibilities satisfying each of these propositions, which is the set of possibilities $\cap\{p_i: i \in \mathrm{I}\}$. This set is called the set of epistemic alternatives for the epistemic modal statement. The ranking is derived from another set of propositions $\{q_j: j \in \mathrm{J}\}$, called an ordering source, each one a set of possibilities, indexed by an index set J (which typically could also be a subset of the natural numbers). If a possibility w_1 satisfies ordering propositions q_1, q_2, and q_3, and another possibility w_2 satisfies just q_1 and q_2, then clearly w_1 is better than w_2 in providing conditions for the evaluation of the modal statement. Following Kratzer, when w_1 is at least as good as w_2 for evaluation of the modal, we write $w_1 \leq w_2$, putting the better-or-equal possibility on the "less than or equal to" side of the inequality symbol. In general, the principle is that if w_1 satisfies ordering propositions $\{q_j: j \in \mathrm{J}_1\}$, for $\mathrm{J}_1 \subseteq \mathrm{J}$, and w_2 satisfies ordering propositions $\{q_j: j \in \mathrm{J}_2\}$, for $\mathrm{J}_2 \subseteq \mathrm{J}$, then $w_1 \leq \mathrm{w}_2$ just in case the set of ordering propositions for w_1 includes, as a subset, the set of ordering propositions for w_2, thus $\mathrm{J}_2 \subseteq \mathrm{J}_1$. But then if we pick any two possibilities, w_1, satisfying ordering propositions $\{q_j: j \in \mathrm{J}_1\}$, and w_2, satisfying ordering propositions $\{q_j: j \in \mathrm{J}_2\}$, there won't in general be any subset relation between the two sets of ordering propositions (e.g., $\{q_1, q_2, q_3\}$ and $\{q_3, q_4, q_5\}$). In that case, it is not true that either $w_1 \leq w_2$ or $w_2 \leq w_1$, so w_1 and w_2 are unranked with respect to one another. The relation $\leq$ defined in this way is not a complete ordering of any kind, since it is, in general, only defined on some pairs of possibilities.

And this means we can't adopt the idea mentioned earlier, where we simply take a candidate threshold of suitability, w_1, and say that *Bill must be working*

[2] For example, if $\mathrm{I} = \{1, 2, \ldots, 27\}$, then the epistemic modal base is $\{p_1, p_2, \ldots, p_{27}\}$.

in his room is true just in case the prejacent expressed by *Bill is working in his room* is true in all possibilities in the epistemic alternative set that are $\leq w_1$. For this wouldn't preclude the existence of another possibility in the alternative set, w_2, such that the prejacent is false on lots of possibilities $\leq w_2$, however good they get for the evaluation of the modal statement. Our outcome for evaluating the truth of the modal statement would depend on our starting point within the set of epistemic alternatives. Instead, we have to step back and say that we get this sort of behavior for any candidate starting point u that we pick, or if not at that point, then for some $v \leq u$. That is, if we pick candidate u at random, then there is some point v "inward of u" ($v \leq u$) that serves as a threshold in the simple sense we started with, namely, that the prejacent, *Bill is working in his room*, is uniformly true in possibilities $\leq v$. This more elaborate approach doesn't assume that any two possibilities in the alternative set are ordered by $\leq$, but it does assert that when we get "far enough down on" $\leq$, and a possibility is ordered $\leq$ any candidate threshold, then the prejacent is true. In effect, the prejacent is true at progressively more possibilities as the rank of possibilities improves according to $\leq$.

But the more elaborate approach raises delicate questions of its own, which go to the heart of the question of how dense is the set of possibilities in the epistemic alternative set with respect to the order $\leq$. As we take possibilities that get progressively better in rank according to $\leq$, will they converge to a set of best possibilities, or will they continue to get progressively better, without end? The former supposition is called the Limit Assumption – that there is a set of best possibilities, whereby possibilities converge to an ideal set, members of which satisfy all propositions in the ordering source $\{q_j: j \in \mathrm{J}\}$. The Limit Assumption would permit a fairly simple formulation of the truth conditions of modal statements.

One of the distinguishing features of the present work is that this simplifying assumption is not made. For reasons given in the overall introduction, it seems to reflect a substantive claim about possibilities that we don't generally have reason to believe is true. The set of propositions in the ordering source does not have to be finite. It could be countably infinite, as we numerate conditions that a possibility must satisfy in order to be "in play" for the evaluation of a modal statement. In certain mathematical contexts, it is even imaginable that there are nondenumerably many propositions in the ordering source. In such a case, possibilities in the epistemic alternative set could get progressively better, with the set of possibilities that satisfy all propositions in the ordering source getting smaller and smaller, without ever reaching a stable set whose members meet all propositions in the ordering source.

The idea at play here, and the needs it is designed to meet, may seem rather abstruse – that we use an elaborate quantificational apparatus to interpret modal statements in order to obtain correct interpretations, even in the special

case when the possibilities to consider keep getting progressively better for the evaluation of the modal statement. But this idea is the one we arrive at, following Kratzer, when we take seriously the idea that modal statements are really about all relevant possibilities. Furthermore, it turns out that this quantificational apparatus helps to meet several challenges that Lassiter (2011) and others pose to any theory of modal interpretation based on possible worlds. With these challenges met, we retain the apparatus of possible worlds to account for the sort of phenomena addressed by Kratzer, while also allowing, in a spirit of semantic pluralism, for Lassiter's preferred account of modal interpretation, based on mapping to scales and providing for a ready interface with numerical probability, to be retained for those purposes.

It also turns out that Kratzer's quantificational schemas, while accounting for what might seem to be abstruse problems in modal interpretation, play an even more critical role by rescuing the interpretation of propositional attitude ascriptions when they come up against real world conditions in which people have inconsistent beliefs, as will be previewed next.

It is common to analyze belief reports and other ascriptions of propositional attitudes as involving quantification over a set of possibilities. Consider the belief report, *Kim believes that Bill is working in his room*. On the view in question, this report makes a claim about all possibilities consistent with what Kim believes. Namely, it asserts that, in each of those possibilities, Bill is working in his room. The terminology used is parallel to that for modality. The set of all propositions characterizing what Kim believes is called Kim's doxastic base, and the set of all possibilities consistent with what Kim believes is called Kim's doxastic *context* or set of doxastic alternatives.

But just as in the evaluation of modal statements, we would want to restrict this quantification to accessible possibilities – those possibilities that are realistic enough to be considered "live alternatives," or to be considered "in play" for the evaluation of the ascription. It may be consistent with what Kim believes that aliens have abducted Bill so that Bill is no longer in his room (she may be somewhat skeptical of alien abduction, but not totally disbelieve in it), but ordinarily the existence of this far-fetched possibility wouldn't falsify the report that Kim believes Bill is working in his room. So the approach would naturally resort to an ordering source on doxastic alternatives to keep things in order.

The epistemic modal base is not the same as the doxastic base. Assuming that Kim is not infallible, she will have beliefs that are not among the things she knows. As a result, the epistemic modal base is a subset of the doxastic base: the propositions known are a subset of those believed. But given that the epistemic base is a subset of the doxastic base, there are fewer possibilities in the doxastic *context*, consistent with the greater number of propositions in the doxastic base, than there are possibilities in the epistemic alternative

set, consistent with what Kim knows.[3] Think of it this way: the possibilities consistent with what Kim believes have to pass a first test, being consistent with what Kim actually knows, and then a second test, being consistent with what Kim merely believes. Assuming some possibilities get thrown out in the second stage of this process, the set of possibilities in the doxastic context, used in the interpretation of the belief report, is a proper subset of the set of possibilities in the epistemic alternative set. So a belief report is interpreted much like a modal statement, with a set of possibilities satisfying propositions in a base, and an ordering source on those possibilities, but the doxastic context is different from the set of epistemic alternatives.

So the standard account of belief ascriptions in possible worlds semantics is to quantify over possibilities in a doxastic alternative set. But there are several problems with this that will require us to formulate it more abstractly. First, in reality, a person can hold inconsistent beliefs. If the set of all of Kim's beliefs happens to be inconsistent, then Kim's doxastic context will be empty, and the interpretive procedure sketched earlier would not have anything to work with. We will take seriously the idea that the experiencer of belief may have inconsistent beliefs, in which case the interpretation of belief reports should not involve quantification over the doxastic context; if Kim's beliefs are inconsistent, that set will be empty, and the quantification would be vacuously true (if it is universal quantification, as it is usually taken to be). Or it would be automatically false if by chance it were existential quantification.

Next, consider how the ordering source works on this account in comparison with the parallel account of modal interpretation. Interpreting modal statements such as *Bill must be working in his room*, we used the ordering source to cut back on possibilities that would make this necessity modal statement spuriously false, for example, one in which some details of the Wars of the Roses diverge from actual history, and Bill is not in his room, when Bill *is* in his room in any of the possibilities in which the Wars of the Roses unfolded as they did in actuality. The ordering source pinned down possibilities to make them accord with the truth (actual world) enough to give us reliable results for modal evaluation.

Now consider a comparable situation in the interpretation of the belief report, *Kim believes that Bill is working in his room*. Consider the set of all possibilities that are consistent with everything that Kim believes. If Kim believes that Bill is working in his room, then Bill is working in his room in all of these possibilities, regardless of whether or not those details about the Wars of the Roses that are beyond Kim's ken happen to accord with reality. Belief ascriptions are fundamentally different from epistemic modal statements in their lack

[3] All talk of fewer and greater number of possibilities here is actually proxy for proper subset relations. We will formulate this more precisely in Chapter 3.