What Did You Say?

A Philosophical Model of Communication and Indeterminacy

Inaugural-Dissertation
zur Erlangung des Doktorgrades der Philosophie
der Ludwig-Maximilians-Universität München

vorgelegt von
Conrad Friedrich
aus
Warburg
2025

Referent: Prof. DDr. Hannes Leitgeb
Korreferent: Prof. Dr. Christian List
Tag der mündlichen Prüfung: 07.02.2025

Contents

List of Figures

Acknowledgments

First and foremost, I would like to thank Hannes Leitgeb. Hannes is the reason I went to Munich to study, and his particular insight and clarity of thought never ceases to impress me. Hannes represents a regulative ideal to strive for not only academically; the amount of patience and goodwill he extends to everyone is exemplary as well. Above that, he has been a great supervisor and provided all the support I needed. I owe a great many thanks to Ursula Danninger, administrator at the MCMP. Ursula helped me navigate the life of a graduate student like no other and always had an ear for my troubles. Without her, it is doubtful whether I would've been able to complete this journey. I'd like to extend special thanks to my high school teacher in philosophy, Hein Besse of Hueffertgymnasium Warburg, who introduced me to the wonders of analytic philosophy. I am indebted to Thomas Grundmann, whose rigorous and convincing style of argumentation has had a great impact on me during my undergraduate studies. I'd further like to thank Mario Guenther, who went over many drafts of my papers and provided invaluable feedback. I'd also like to thank, for intellectual input, engaging discussion, or advice throughout my doctorate: Marianna Antonutti Marfori, Luis Bartolo Alegre, Javier Belastegui, Cordelia Berz, Sena Bozdag, Anne Deng, Maximilian Dore, João Ferrari, Martin Fischer, Andreas Frenzel, Sebastian Gäb, Ben Genta , Norbert Gratzl, Stefan Hartmann, Hosea von Hauff, Levin Hornischer, Elio la Rosa, Christian List, Alessandra Marra, Rodrigo Mena González, Will Nolan, Ragna Oeynhausen, Adrian Ommundsen, Leonardo Ortiz-Acuña, Hannah Pillin, Emily Redekop, Alexander Reutlinger, Tyko Schuff, Tom Sterkenburg, Diego Tajer, Masako Toyoda, Zhouwanyue Yang, Yingtong Ye, and Maximilian Zeuner. My thanks further go to everyone I had the pleasure of interacting with at the Munich Center of Mathematical Philosophy. The most heartfelt thanks I'd like to extend to my parents,

Susanne und Heiner, whose continuous support and unwavering faith in my abilities were absolutely crucial during my whole development and studies. Vielen lieben Dank euch beiden!

Deutschsprachige Zusammenfassung

Angenommen, Sie und Ihr Partner bereiten sich auf einen Spaziergang vor und sind sich nicht sicher, ob Sie den Regenschirm mitnehmen sollen. Ihr Partner ruft:

(1) Es regnet in Strömen!

Ohne lange zu überlegen, nehmen Sie den Regenschirm mit. Ihr Partner hat mit Ihnen kommuniziert, Sie haben verstanden, und die entsprechende Maßnahme ergriffen. So weit, so unscheinbar.

Aber warum war die Kommunikation erfolgreich? Was macht es aus, dass die verbale Äußerung Ihres Partners zur erfolgreichen Kommunikation gezählt werden kann? Die Standardantwort in der Philosophie lautet in etwa so: Ihr Partner hat etwas ausgedrückt, eine Bedeutung, indem er/sie die deutsche Sprache kompetent verwendet hat. Sie haben verstanden, was ausgedrückt wurde, denn auch Sie sind kompetent darin, Deutsch zu sprechen. Damit Sie Ihren Partner verstehen konnten, mussten Sie die Bedeutung—den Gedanken oder die Proposition—erkennen, die ausgedrückt wurde. Oft wird diese Bedeutung als Wahrheitsbedingungen erklärt: Sie haben verstanden, was ausgedrückt wurde, wenn Sie wissen, unter welchen Bedingungen der geäußerte Satz wahr wäre. Der Rest ist bloße praktische Schlussfolgerung. Es regnet, also nehmen wir den Regenschirm mit, um nicht nass zu werden. Ich nenne eine solche Sichtweise das *Klassische Modell der Kommunikation* und werde es in dieser Arbeit weiter entwickeln und seine philosophischen Ursprünge nachzeichnen.

Aber das klassische Modell steht vor einem Problem. Was sind eigentlich die Wahrheitsbedingungen für (1)? Bei näherer Betrachtung ist das nicht so offensichtlich. Zählt ein kurzer Nieselregen als *regnen in Strömen*? Zählt

ein Hagelsturm? Was ist, wenn Sie sich an der äußeren Grenze des Einflussbereichs einer Regenwolke befinden? Oder gibt es eine messbare Menge an Regen, die erforderlich ist, damit (1) wahr ist? Die Bedeutung der Äußerung scheint in mancher Weise *unbestimmt* zu sein.

Für den intuitiven Fall scheinen all diese Überlegungen keine Rolle zu spielen. Es regnet, und deshalb braucht man einen Regenschirm, Ende der Geschichte. Was ist daran kompliziert?

Das klassische Modell verlangt jedoch, dass sowohl der Sprecher als auch das Publikum in einer epistemisch qualifizierten Beziehung zu der *gleichen* Proposition stehen. Wenn die Proposition, die der Sprecher äußert, und die Proposition, die das Publikum aufgreift, voneinander abweichen, kann das klassische Modell den Erfolg der Kommunikation nicht erklären. Zugegeben, man könnte sagen, dass das klassische Modell eine Idealisierung ist und dass die meisten Fälle nicht so problematisch sind wie (1). Aber es gibt nichts wirklich Besonderes an (1). Wie in dieser Arbeit dargelegt wird, sind Fälle unbestimmen Inhalts allgegenwärtig und treten in vielen kommunikativen Situationen auf.

In dieser Arbeit wird das klassische Modell gegen die Herausforderung der Unbestimmtheit verteidigt. Diese Herausforderung wird an passender Stelle noch präzisiert werden. Zu diesem Zweck entwickle ich einen Ansatz, den ich als "Grobkörniges Modell der Kommunikation" bezeichne und der eine Erweiterung des klassischen Modells darstellt. Der Grundgedanke des Grobkörnigen Modells ist, dass die inhaltlichen Unterschiede zwischen Sprecher und Publikum für praktische Zwecke keine Rolle spielen. Für praktische Zwecke gibt es tatsächlich nur eine einzige Proposition, die kommuniziert wird. Das Grobkörnige Modell verteidigt also das klassische Modell, indem es das klassische Modell um eine Methode zur Bestimmung eindeutiger Wahrheitsbedingungen durch praktische Relevanz erweitert.

In dieser Arbeit wird also eine einfache Frage gestellt. Warum gelingt die Kommunikation zwischen zwei Menschen? Mit anderen Worten, was macht es möglich, dass eine Person einer anderen Person einen beliebigen Gedanken mitteilen kann und diese Person versteht, was die erste Person sagen wollte? Was sind also die Bedingungen für den Erfolg, und wann scheitert die Kommunikation? Oder, wie Wilson und Sperber es ausdrücken:

> Das Studium der Kommunikation wirft zwei wichtige Fragen auf: Erstens, was wird kommuniziert, und zweitens, wie wird Kom-

munikation erreicht? (Sperber & Wilson, 1986)

Die vorliegende Arbeit wird zu diesen Fragen Stellung nehmen und einen Rahmen für die systematische Erklärung des kommunikativen Erfolgs entwickeln. Ich werde relevante Darstellungen von Kommunikation in der Geschichte der Philosophie und zeitgenössische Darstellungen diskutieren und ihren Beitrag zum Klassischen Modell herausarbeiten. Das Klassische Modell beantwortet die zentralen Fragen der Dissertation in vielen Fällen überzeugend. Es sieht sich jedoch ernsthaften Herausforderungen durch eine Reihe von Fällen gegenüber, in denen es schwierig bis unmöglich erscheint, eine einzelne kommunizierte Proposition zu bestimmen. Ich werde verschiedene aktuelle Ansätze zur Bestimmung des Inhalts solcher Äußerungen betrachten und kritisch diskutieren, die allesamt den einen oder anderen Aspekt des klassischen Modells aufgeben. Diese Diskussion verlangt nach einer Lösung, und ich werde ein solches Modell im Detail entwickeln. Das Modell erweitert das Klassische Modell und ermöglicht es, in allen problematischen Fällen eine einzige Proposition als Inhalt der Äußerung zu bestimmen. Das Modell ist also maximal konservativ gegenüber dem klassischen Modell. Das Modell erklärt Kommunikation mit einer expliziten Artikulation des Klassischen Modells, erweitert um eine pragmatische Art der Bestimmung von Wahrheitsbedingungen für Äußerungen.

Plausiblerweise hat jede intuitiv ansprechende Darstellung von Kommunikation mit dem Begriff der *Bedeutung* zu tun. Wann immer wir in der mündlichen oder schriftlichen Kommunikation Ausdrücke verwenden, meinen wir mit unserem Ausdruck etwas. Bei den Ausdrücken kann es sich um Sätze, Äußerungen, Handgesten, Rufe und so weiter handeln. In der Philosophie ist dieses Etwas, das wir meinen, normalerweise von den Ausdrücken selbst verschieden und geht über sie hinaus, obwohl es Ansichten gibt, die diese weit verbreitete Annahme in Frage stellen (siehe z.B. Gauker, 2002). Etwas, das wir meinen, ist nicht „nur“ ein Symbol, in welcher Form auch immer, sondern etwas, wofür die im Ausdruck verwendeten Symbole stehen. Irgendwie überträgt der „Sender“ der Kommunikation, den ich im Folgenden „Sprecher“ nenne, obwohl Kommunikation natürlich nicht notwendigerweise verbal ist, etwas *zusätzlich* zum Ausdrucksmittel an den „Empfänger“, den ich „Publikum“ nenne. Die Kommunikation kann als erfolgreich bezeichnet werden, wenn das Publikum in einer epistemisch qualifizierten Beziehung zur Bedeutung des Ausdrucksmittels steht. Diese epistemische Relation ist

eine Erfolgsrelation—wie das Erkennen, das Wissen oder der Glaube in einer epistemisch qualifizierten Weise. Was genau erkennen? Betrachten wir eine verbale Äußerung. Erfolg könnte erfordern, dass man die Bedeutung erkennt, die die Äußerung *eigentlich* ausdrückt, falls es so etwas gibt. Oder der Erfolg erfordert das Erkennen der Bedeutung, die der Sprecher mit seiner Äußerung für das Publikum zu erkennen beabsichtigte, falls er dies tat. Es könnte genügen, eine dieser Bedeutungsarten nur teilweise zu erkennen. Oder der Erfolg könnte durch das Erkennen einer ganz anderen Bedeutung erreicht werden, die aber in einer Ähnlichkeitsbeziehung zur tatsächlichen oder beabsichtigten Bedeutung der Äußerung steht, die spezifiziert werden müsste. Dieser Gedankengang wirft auch die Frage auf, was für ein Ding eine Bedeutung ist, dass Menschen und Äußerungen in bestimmten Beziehungen zu ihr stehen können. Die Bedeutung kann physisch, mental, abstrakt usw. sein. Was auch immer die zu erkennende Bedeutung ist, wie kann das Publikum sie erkennen? Die Übertragung vom Sprecher zum Publikum könnte verzerrt und verrauscht sein, sodass es sein kann, dass das Publikum die Bedeutung des Ausdrucks nicht erkennen kann. Der Sprecher könnte die Wörter auch auf eine ganz andere Art und Weise verwenden als das Publikum. Ihr Sprachgebrauch könnte teilweise undurchsichtig, d. h. privat sein. Das Erkennen der beabsichtigten Bedeutung kann dann ernsthaft beeinträchtigt sein. Dennoch gelingt die Kommunikation oft, allen Widrigkeiten zum Trotz, könnte man meinen.

Diese kurzen prätheoretischen Überlegungen scheinen es dringend nötig zu haben, philosophisch organisiert und behandelt zu werden. Es überrascht nicht, dass Philosophen dies schon lange getan haben. Dies, so argumentiere ich, führte zum Klassischen Modell der Kommunikation, dessen genaue Formulierung ich aus MacFarlane (2020a, 2020b, 2020c) übernehme:

Klassische Pragmatik

(1) Der Inhalt einer Behauptung ist eine (einzelne) Proposition.

(2) Die Aufnahme besteht im Erkennen des behaupteten Satzes.

(3) Wenn die Behauptung akzeptiert wird, wird ihr Inhalt der gemeinsamen Gesprächsgrundlage hinzugefügt.

Klassischer Inhalt

Inhalte sind Weisen, wie die Welt sein könnte (Wahrheitsbedingungen).

Inwiefern *erklärt* das Klassische Modell den kommunikativen Erfolg? Es liefert Details darüber, *warum* Kommunikation erfolgreich ist: Etwas—irgendein Inhalt—wird zwischen den Personen übertragen; etwas potenziell extra-mentales. Beide Personen stehen in einer epistemischen Beziehung zu diesem Inhalt. Das Modell vermittelt auch eine Vorstellung davon, wie die Kommunikation gelingt: Inhalte werden durch kontextuelle Bedingungen beeinflusst, und man kann Inhalte modellieren, indem man Wahrheitsbedingungen angibt.

Im Teil I entwickle ich das Klassische Modell historisch, stelle das Problem der Unbestimmtheit dar und diskutiere zeitgenössische Ansätze, die allesamt das Klassische Modell nicht retten können.

Im Kapitel 2 gehe ich auf die historische Entwicklung der einzelnen Prinzipien des Klassischen Modells ein. Die Klassische Pragmatik (1) und (2), d.h. dass der Inhalt einer Behauptung eine Proposition ist und dass die Aufnahme im Erkennen der behaupteten Proposition besteht, lässt sich mindestens bis zu Aristoteles' *De Interpretatione* des *Organon* zurückverfolgen. Auch wenn Aristoteles' Position zu dem, was das Übermittelte in der Kommunikation ausmacht, von der heutigen Form des Klassischen Modells abweicht. Dennoch ist klar, dass er bereits einen Prozess beschrieben hat, der dem ähnelt, was Gauker (1992) als *Locksche Kommunikationstheorie* bezeichnet hat: Grob gesagt, die Verbindung von Klassischer Pragmatik (1) und (2). Lockes Sprachphilosophie ist dann tatsächlich auch ein Schwerpunkt des Kapitels 2. Die Behauptung, dass Inhalte Wahrheitsbedingungen sind, d.h. Klassischer Inhalt, wird hauptsächlich Gottlob Frege zugeschrieben, und die weitere Idee, Wahrheitsbedingungen als Mengen möglicher Welten zu modellieren, Rudolf Carnap. Obwohl viele Autoren im 20. Jahrhundert bedeutende Beiträge zu diesem Teil des klassischen Modells geleistet haben. An erster Stelle ist hier Kripke (1959, 1963a, 1963b) zu nennen.

Dass der Inhalt der Behauptung zur gemeinsamen Gesprächsgrundlage hinzugefügt wird, d.h. Klassische Pragmatik (3), ist eine neuere Entwicklung, die auf Paul Grices einflussreiche Arbeit über kontextabhängige Inhalte und dann auf Robert Stalnakers Arbeit über Kontext und gemeinsame Gesprächsgrundlage basiert.

Das klassische Modell geht also davon aus, dass der Inhalt einer Äußerung ein einzige Proposition ist. Dies funktioniert gut für Äußerungen, deren Wahrheitsbedingungen einigermaßen offensichtlich sind. Wenn der Sprecher

und die Zuhörer die Wahrheitsbedingungen einer Äußerung leicht erkennen können, ist die Kommunikation unproblematisch. Geringfügige Hindernisse wie die Bestimmung der Referenten von indexikalischen Begriffen oder einfache Zweideutigkeiten können überwunden werden. Probleme entstehen, wenn die ausgedrückte Bedeutung unbestimmt ist. Die Bedeutung kann in verschiedenen Weisen unbestimmt sein, und ich werde im Kapitel 3 diese Weisen disambiguieren und deutlich machen, welche davon ein Problem für ein Kommunikationsmodell darstellen. Die vielleicht offensichtlichste Weise lässt sich als *semantische* Unbestimmtheit zusammenfassen, die Fälle von vager Sprache einschließt. Wenn der semantische Inhalt des geäußerten Satzes keine genauen Wahrheitsbedingungen festlegt und der Kontext diese auch nicht zu bestimmen scheint, hat das klassische Modell Schwierigkeiten, den kommunikativen Erfolg zu erklären. In anderen Fällen ist der Inhalt einer Äußerung offensichtlich kontextabhängig, aber es ist keineswegs klar, wie die intuitiv kommunizierte Proposition bestimmt werden kann. Dies ist ein Fall von Unbestimmtheit in undurchsichtigen Kontexten. Wenn schließlich der Kontext das Publikum nicht zu verpflichten scheint, eine bestimmte Proposition aus einer ganzen Reihe möglicher Alternativen zu wählen, spreche ich von Unbestimmtheit aufgrund kontextueller Indifferenz.

Die Unbestimmtheit des Äußerungsinhalts, gleich welcher Art, führt zu einem erkenntnistheoretischen Problem, nämlich dem der Bestimmung des Inhalts eines bestimmten kommunikativen Akts. Es ist oft der Fall, dass die konventionelle, sprachliche Bedeutung eines Ausdrucks nicht ausreicht, um zu bestimmen, was mit einer Äußerung dieses Satzes gesagt wurde. Darüber hinaus hat nicht nur der Philosoph, der versucht, einen gelungenen kommunikativen Austausch zu erklären, Schwierigkeiten, mit Unbestimmtheit umzugehen, sondern auch das Publikum. Eine überzeugende Darstellung von Kommunikation muss sich meiner Ansicht nach mit der Frage befassen, wie das Publikum das epistemologische Rätsel lösen und die Botschaft des Sprechers erschließen, intuitiv erfassen oder auf andere Weise erkennen kann. Folglich stehen die Begriffe *was wird gesagt* und *was wird kommuniziert* im Mittelpunkt des in dieser Dissertation entwickelten Modells.

Die Herausforderungen für das klassische Modell sind vielfältig, und viele Autoren haben sich in ihren Ansätzen dazu entschlossen, einige Teile des klassischen Modells aufzugeben. Einige Ansätze streiten ab, dass Propo-

sitionen Wahrheitsbedingungen erfordern, andere behaupten, dass unbestimmte Äußerungen statt einer einzigen Proposition mehrere Propositionen ausdrücken, und wieder andere funktionieren aus technischen Gründen nicht. Ich werde mich im Kapitel 4 eingehend mit diesen Ansätzen befassen. Diese Ansätze schaffen es nicht, das Klassische Modell gegen den Einwand der Unbestimmtheit zu verteidigen.

Da ich mit den angebotenen Lösungen nicht zufrieden bin, stelle ich im Teil II die Alternative vor, die ich als „Grobkörniges Modell der Kommunikation" bezeichne. Zunächst werde ich die Begriffe präzisieren, die für die Formulierung des Grobkörnigen Modells notwendig sind. Der Schwerpunkt liegt dabei auf den Begriffen *Frage* und *Antwort*, die im Kapitel 5 erläutert werden. Da Fragen und ihre Antworten eine zentrale Rolle im Grobkörnigen Modell spielen, ist es sehr wichtig zu klären, was ich mit einer *Frage* und was ich mit einer *Antwort* auf diese Frage meine. Weitere Begriffe, die zur Entwicklung des Grobkörnigen Modells benötigt werden, werden im Kapitel 6 erläutert. Das Grobkörnige Modell wird im Kapitel 7 kurz und bündig dargestellt: Das Modell greift einige Elemente aus den zuvor diskutierten zeitgenössischen Darstellungen auf, stellt aber eine wichtige zusätzliche Behauptung auf. In vielen Fällen gibt es tatsächlich eine einzige Proposition, die als die Kommunizierte identifiziert werden kann. Diese Proposition ist die Erfüllung eines gemeinsamen Zwecks oder Ziels, das wiederum durch eine Frage repräsentiert wird, die Sprecher und Publikum zu beantworten suchen. Die Kommunikation ist also „gut genug" für die gegenwärtigen Zwecke des Gesprächs, auch wenn die verwendeten Ausdrücke unbestimmt sind. Dies kann auch dann erreicht werden, wenn Sprecher und Publikum die Ausdrücke leicht unterschiedlich interpretieren. Unter der Annahme, dass eine Frage zur Diskussion steht, die entweder explizit gestellt oder auf andere Weise implizit erzeugt wird, liefert das Modell ein Verfahren, um aus einem Äußerungskontext und einer Äußerung die kommunizierte Proposition zu bestimmen. Es ist klar, dass die Botschaft des Sprechers im Grobkörnigen Modell pragmatischer Natur und stark kontextabhängig ist. Das Ergebnis des Modells ist unterteilt in *was der Sprecher auszudrücken beabsichtigt* und *was das Publikum den Sprecher auszudrücken interpretiert*, die in ähnlicher Weise bestimmt werden, aber jeweils aus der Perspektive des Sprechers und des Publikums. Hier vorab eine These:

These *Was das Publikum den Sprecher auszudrücken interpretiert* mit einer

> wörtlichen assertorischen Äußerung in einem Kontext ist die Vereinigung aller Vergröberungen der Frage-Antwort-Interpretationen, die das Publikum für zulässig hält.

Diese These mag auf den ersten Blick schwer verständlich erscheinen und enthält einige idiosynkratische Ausdrücke. Ich werde hier eine intuitive Beschreibung der beteiligten Begriffe geben, um dem Leser eine Vorstellung von der späteren detaillierteren und technischen Beschreibung zu vermitteln. Eine *wörtliche assertorische Äußerung* ist eine ganz normale Aussage, die keine höheren sozio-kognitiven Anforderungen stellt wie z.B. Ironie es tut. Eine Interpretation ist eine determinierte Bedeutung des (möglicherweise) unbestimmten Ausdrucks. Enthält der geäußerte Satz beispielsweise einen vagen Ausdruck, dann ist eine Interpretation eine Möglichkeit, den Ausdruck genau zu deuten. Wir betrachten nur Interpretationen, die das Publikum für zulässig hält—oder die mit dem geäußerten Satz vereinbar sind (neben anderen Einschränkungen). Eine Interpretation ist eine Frage-Antwort-Interpretation, wenn sie eine Antwort auf die Frage zur Diskussion gibt. Mit anderen Worten: Eine Interpretation ist eine Frage-Antwort-Interpretation, wenn sie das Ziel des Gesprächs in irgendeiner Weise fördert. Eine Vergröberung einer Interpretation ist der Teil der Interpretation, der sich nur auf die Beantwortung der Frage zur Diskussion bezieht. Wir nehmen die Vereinigung all dieser Vergröberungen, um die Interpretation zu einer gemeinsamen Proposition zu vereinheitlichen. Die Vereinigung ist hier ein mengentheoretischer Begriff, da Proposition als Mengen möglicher Welten dargestellt werden. Intuitiv ausgedrückt ist also das, was das Publikum den Sprecher auszudrücken interpretiert, eine Antwort auf die Frage zur Diskussion, die mit allen Interpretationen übereinstimmt, die das Publikum mit dem geäußerten Satz für vereinbar hält. Der Prozess zur Bestimmung dessen, was der Sprecher auszudrücken beabsichtigt, ist ziemlich ähnlich, obwohl hier der Schwerpunkt auf den Interpretationen liegt, die der Sprecher für zulässig hält. Erfolgreiche Kommunikation ist dann gegeben, wenn sowohl das, was der Sprecher auszudrücken beabsichtigt als auch das, was das Publikum interpretiert, miteinander übereinstimmen. Das Ergebnis ist eine eindeutige Proposition, die kommuniziert wird und so den Erfolg der Kommunikation im klassischen Modell erklärt.

Im Kapitel 9 erörtere ich ausführlich eine Reihe von Einwänden gegen die Erklärung der Kommunikation mit dem Grobkörnigen Modell. Einer

dieser Einwände betrifft die Ansicht, dass das Grobkörnige Modell die Unbestimmtheit auflöst, indem es etwas Bestimmtes postuliert: Die zur Diskussion stehende Frage. Einfach ausgedrückt: Die Unbestimmtheit der Äußerung wird durch eine bestimmte Frage ersetzt, die zur Diskussion steht. Aber, so der Einwand, das ist doch nur ein Verschieben der Beweislast auf etwas, das auch nicht besser erklärt wird. Welche Rechtfertigung gibt es für die Annahme einer determinierten Frage, die zur Diskussion steht, könnte sie nicht ebenso unbestimmt sein? In gewissem Sinne handelt es sich hier um eine Art "Unbestimmtheit höherer Ordnung", ähnlich dem Begriff der Vagheit höherer Ordnung. Eine unbestimmte Frage, die zur Diskussion steht, garantiert nicht, dass sowohl der Sprecher als auch das Publikum ihre Bedeutung teilen. Mit anderen Worten: Publikum und Sprecher können die Frage unterschiedlich interpretieren, z. B. weil sie unbestimmte Ausdrücke enthalten kann.

Ein weiterer Einwand zweifelt die Annahme eines gemeinsamen Zwecks in Form einer direkten Frage zur Diskussion an. Es ist klar, dass der technische Begriff der Frage zur Diskussion einen großen Teil der Arbeit in der Darstellung leistet. Aber man könnte bezweifeln, dass eine solche Frage, die von den Diskursteilnehmern gegenseitig akzeptiert wird, in vielen Fällen plausibel ist. Was ist zum Beispiel, wenn jemand mit einer völlig neuen Information herausplatzt, die nichts mit dem vorherigen Thema der Diskussion zu tun hat? Es scheint keine vorher festgelegte Frage zu geben, auf die die Äußerung eine Antwort geben soll, und somit scheint die Erklärung nicht einmal anwendbar zu sein.

Ein wichtiger Grund, am klassischen Modell der Kommunikation festzuhalten, ist seine nahtlose Integration in die Semantik. So wie Propositionen bei der Erklärung der Bedeutung eines Satzes eine wichtige Rolle spielen, so spielen sie auch bei der Erklärung der Bedeutung einer Äußerung eine wichtige Rolle, die wiederum das zentrale Konzept des klassischen Modells ist. Die Idee ist also, dass linguistische Resultate über die Bedeutung eines Satzes wortwörtlich auch auf die Bedeutung einer Äußerung angewendet werden können. Das macht das klassische Modell zum Teil so attraktiv. Es funktioniert deshalb so gut, weil sich die Bedeutung eines Satzes normalerweise aus der Bedeutung seiner Teile zusammensetzt, die durch syntaktische Regeln kombiniert werden. Die Bedeutung eines Satzes—die Proposition—ist in diesem Sinne kompositionell. Aber die Bedeutung einer Äußerung, wie sie

nach oben genanntem Prozess bestimmt wird, ist bis zu einem gewissen Grad von der Bedeutung des geäußerten Satzes losgelöst, da sie auch durch die Frage zur Diskussion bestimmt wird. Die Bedeutung einer Äußerung ist dann nicht in gleichem Maße kompositionell wie die Bedeutung des geäußerten Satzes. Das Klassische Modell verliert so an Überzeugungskraft.

Das Modell scheint ferner eine bestimmte implizite Haltung zum metaphysischen Status des Gemeinten einzunehmen. Durch die Entkopplung der Interpretation einer Äußerung vom Kontext der Äußerung in Form einer möglichen Welt scheint das Modell an Aussagekraft zu verlieren. Denn in der Standard-2-D-Semantik ist es eine Tatsache, dass ein bestimmter Ausdruck eine bestimmte Intension in einem Äußerungskontext hat. Für eine gegebene Welt ist die Bedeutung dieses Ausdrucks festgelegt. Dies ermöglicht es, Aussagen *über* die Bedeutung von Ausdrücken als Sätze im gleichen Rahmen, d.h. als Mengen von Welten, darzustellen. Wenn aber z. B. in der realen Welt ein Ausdruck gültig in vielen verschiedenen Bedeutungen interpretiert werden kann, was soll man dann mit einer Proposition über die Bedeutung eines solchen Ausdrucks anfangen? Was sind seine Wahrheitsbedingungen?

Das Modell bestimmt die Bedeutung durch den logischen Ausschluss von Fragealternativen. Manchmal beantworten jedoch Proposition, die eine Frage intuitiv beantworten, diese Frage nicht in diesem eher technischen Sinne. Zum Beispiel:

(2) a. Frage. Wer ist wach?

b. Antwort. Kelly hat morgen einen großen Tag.

Hier liefert die Antwort nur nicht-deduktive Informationen in Bezug auf die Frage. Zu den nicht-deduktiven Informationen gehören induktive, ampliative, probabilistische und so weiter. Die Antwort liefert Evidenz zur Beantwortung der Frage, ohne die Frage im technischen Sinne zu beantworten. Das Modell ist jedoch auf die Beantwortung einer Frage zur Diskussion angewiesen, um der Äußerung überhaupt eine Bedeutung zuzuweisen. Für diese Art von Frage-Antwort-Paaren scheint das Modell also nicht anwendbar zu sein.

Dies sind nur einige der Einwände, die im Kapitel 9 diskutiert werden. Ich argumentiere, dass es für die meisten diskutierten Einwände plausible Strategien der Widerlegung für die Befürworter des Grobkörnigen Modells gibt.

Chapter 1

Introduction

Suppose you and your partner are getting ready to go for a walk, and are unsure about whether to take the umbrella. Your partner exclaims:

(3) It's pouring!

Without deliberating much, you'll take the umbrella with you. Your partner communicated with you, and you understood and took the appropriate action. So far, so unremarkable.

But why was the communication successful? What makes it so that your partner's verbal expression can count towards successful communication? The standard answer in philosophy goes something like this: your partner expressed something—a meaning—by using their command over the English language. You understood what was expressed, because you, too, are a competent English speaker. For you to understand your partner meant to recognize the meaning—the thought or proposition—expressed. Often, this meaning is cashed out in terms of truth conditions: you understood what was expressed if you know under which conditions the uttered sentence would be true. The rest is mere practical inference: sure, it's raining, so let's take the umbrella to not get wet. I'll call a view like this the *Classical Model of Communication*. I trace its philosophical origins and further develop the model in this thesis.

But the Classical Model faces a problem. What actually are the truth-conditions for (3)? Under closer inspection, that's not so obvious. Does a short drizzle count as pouring? Does a hailstorm? What if your place is just at the outer limit of the field of influence of a rain cloud? Or is there a measurable quantity of rain which is required for (3) to be true? The

meaning of the utterance seems to be in some way *indeterminate*.

Intuitively, all of these considerations don't seem to matter. It's pouring, and thus you need an umbrella, end of story. What's complicated about it?

But the Classical Model requires both speaker and audience to stand in an epistemically qualified relationship to the *same* proposition. If the proposition the speaker expresses and the proposition the audience takes up differ, the Classical Model cannot explain the communicative success. Granted, one might say that the Classical Model is an idealization and that most cases are not as problematic as (3). But it turns out there is nothing really special about (3). As will be argued in this thesis, forms of indeterminate content are ubiquitous and appear in many communicative situations.

This thesis defends the Classical Model against the challenge from indeterminacy. To this end, I develop a view which will be called the "Coarse-Grained Model of Communication", an extension of the Classical Model. This challenge will be made more precise in due course. The basic idea of the Coarse-Grained Model is that the differences in content between speaker and audience, like e.g. when exactly to count a situation as "pouring", do not matter for practical purposes. For practical purposes, there is in fact a single proposition communicated. The Coarse-Grained Model thus defends the Classical Model by extending it with a method to determine unique truth conditions by way of practical relevance.

This thesis thus asks a simple question. Why is it that communication between two people succeeds? In other words, what makes it so that one person can signify whatever thought they are entertaining to another person, and that person will understand what the first wanted to say? Consequently, what are the conditions for success, and when does communication fail? Or, as Wilson and Sperber put it:

> The study of communication raises two major questions: first, what is communicated, and second, how is communication achieved? (Sperber & Wilson, 1986)

This thesis will take a stance on these questions, and develop a framework for systematically explaining communicative success. I will discuss relevant accounts of communication in the history of philosophy and contemporary accounts, and their contribution to what amounts to the Classical Model. The Classical Model answers the central questions of the thesis convincingly

for many cases. But it faces serious challenges from a range of cases in which it seems difficult to impossible to determine a single proposition communicated. I will look at, and critically discuss, a variety of current accounts intended to determine the content of such utterances, all of which give up some aspect of the Classical Model. This discussion calls for a solution, and I will develop such a model in detail. The model extends the Classical Model and enables determining a single proposition for the utterance's content in all problematic cases. The thesis is thus maximally conservative w.r.t. the Classical Model. It explains communication with an explicit articulation of the Classical Model extended with a pragmatic way of determining truth conditions for utterances.

Any intuitively appealing account of communication plausibly will have to do with the concept of *meaning*. Whenever we use expressions in verbal or written or otherwise signifying communication, we somehow mean something with our expression. The expressions can be something like sentences, utterances, hand gestures, shouts, and so on and forth. Usually, in philosophy, this something we mean is taken to be distinct from and goes beyond the expressions themselves. There are views which challenge this widespread assumption (see, e.g. Gauker, 2002). This something we mean is not "just" symbols of whichever form, but something the symbols used in the expression stand in for. Somehow, the "sending agent" in the communication transmits something *in addition to* the expressive device to the "receiving agent". In the following I'll simply use "the speaker" for the sending agent, even though communication, of course, is not necessarily verbal. The receiving agent I'll just call "the audience". The communication can be said to succeed if the audience comes to stand in an epistemically qualified relation to the meaning of the expressive device. This epistemic relation is a success relation—like recognizing, knowing or otherwise believing in an epistemically qualified manner. Recognizing what, exactly? Let's look at a verbal utterance. Success could require recognizing the meaning the utterance *actually* expresses, if there is such a thing. Or success requires to recognize the meaning the speaker intended for the audience to recognize with their utterance, if they did so. It might be enough to just partially recognize one of these types of meaning. Or success could be achieved by recognizing an altogether different meaning, but which stands in some similarity relation (which would have to be specified) to the actual or intended

meaning of the utterance. This line of thought raises, too, the question what kind of thing a meaning is, that people and utterances can stand in certain relations to it. The meaning might be physical, mental, abstract, etc. Whatever the meaning to be recognized is, how can the audience recognize it? The transmission from speaker to audience might suffer distortion and noise, and the audience would be prevented from recognizing the expression's meaning. The speaker might also use words in a decidedly different manner from the way the audience does. Their use of language might be partially opaque, i.e. private. Recognizing the intended meaning might then be seriously impacted. Yet, communication often succeeds, against all odds, one might think.

These brief pre-theoretic considerations seem to be in dire need of being philosophically organized and addressed. Unsurprisingly, philosophers have long done so. This, I argue, leads to the Classical Model of Communication, the formulation of which I adopt from MacFarlane (2020a, 2020b, 2020c):

Classical Pragmatics

(1) The content of an assertion is a (single) proposition.

(2) Uptake consists in recognizing the proposition asserted.

(3) If the assertion is accepted, its content is added to the conversational common ground.

Classical Content

Contents are ways the world might be (truth conditions).

In what sense does the Classical Model *explain* communicative success? It provides details as to *why* communication succeeds: something—some content—is transmitted between subjects' minds; something potentially extramental, i.e. potentially part of neither mind. Both minds stand in an epistemological relation to this content. The model also provides an idea of *how* communication succeeds: Content is shaped by contextual constraints, and one can model content by specifying truth conditions.

In part I, I develop the Classical Model historically, pose the problem from indeterminacy, and discuss contemporary accounts, all of which fail to save the Classical Model.

In chapter 2, I detail the historical development of the individual principles of the Classical Model. Classical Pragmatics (1) and (2), i.e. that

the content of an assertion is a proposition and that uptake consists in recognizing the proposition asserted, can be traced back at least to Aristotle's *De Interpretatione* of the *Organon*. Even though Aristotle's position differs from the present form of the Classical Model. For Aristotle, something other than a proposition is transmitted in communication. Still it is clear that he already described a process akin to what Gauker (1992) termed the *Lockean Theory of Communication*: roughly, the conjunction of Classical Pragmatics (1) and (2). Locke's philosophy of language is then indeed also a focal point of chapter 2. The claim that contents are truth-conditional, i.e. Classical Content, will be attributed mostly to Gottlob Frege, and the further idea of modeling truth conditions as sets of possible worlds to Rudolf Carnap, although many authors in the 20th century made significant contributions to this part of the Classical Model. Foremost among these is Kripke (1959, 1963a, 1963b). Finally, that the content of the assertion is added to the common ground, i.e. Classical Pragmatics (3), is more a recent development, following Paul Grice's influential work on context-dependent content and then Robert Stalnaker's work on context and common ground.

The Classical Model assumes the content of an utterance to be a single truth-conditional proposition. This works well for utterances in which the truth-conditions are somewhat obvious. If the speaker and audience can easily know the truth-conditions of an utterance, communication is unproblematic. Slight obstacles like determining referents of indexical terms or straightforward ambiguities can be overcome. Problems arise when the meaning expressed is indeterminate. The meaning can be indeterminate in different ways. I will disambiguate these ways into types and make clear why they pose a problem for the Classical Model of communication in chapter 3.

The maybe most obvious type can be summarized as *semantic* indeterminacy, which includes cases of vague language. If the semantic content of the uttered sentence does not determine precise truth conditions, and the context does not seem to fill them in either, the Classical Model has trouble explaining communicative success. Consider a simple statement like "Fred is bald". What are the exact truth conditions of such an utterance? It's not obvious, if possible at all, to clearly delineate the number of hairs Fred has to have to not count as bald.

A second type is indeterminacy in opaque contexts. In many cases, the content of an utterance is apparently context-dependent, but it is far from

clear how the intuitively communicated proposition can be determined. Consider the statement "Biden is too old." Clearly, something in addition to the mere sentence is required to determine its truth conditions. Biden is too old for what? Knowledge of present political discussion will propose that Biden is too old to run again for office. But the context might be about the retirement age of skateboarders, or whatever else. Which feature of the context determines the truth conditions of the utterance? How do the truth conditions systematically depend on those features?

A third type is indeterminacy due to contextual indifference: this occurs if the context does not seem to require the audience to pick up any one particular proposition out of a whole range of possible alternatives. For example, consider the statement "I'm mad at you. You said that you're not a Ferrari fan. Everybody is a Ferrari fan. Even if they are not, they *are* a Ferrari fan." Who, exactly, is referred to by "everybody"? It seems that at least the person talked to is included. But surely not *everybody* in the whole world, even people who never heard of the brand. For the purposes of the conversation, it doesn't seem to matter much which exact extension one assigns to "everybody" to understand the utterance. The context is *indifferent* to the proposition expressed. But not arbitrarily so, this thesis will argue.

Indeterminacy in utterance-content, whatever its type, leads to an epistemological problem, namely that of determining the content for a given communicative act. It is often the case that a standing, conventional, linguistic meaning of a sentence token is not sufficient to determine what has been said by an utterance of that sentence. Additionally, not only the philosopher trying to explain a successful communicative exchange has trouble when dealing with indeterminacy, so too does the audience. A convincing account of communication, in my opinion, has to deal with the question how the audience can solve the epistemological puzzle and infer, intuit, or in some other way recognize the speaker's message. Consequently, the notions *what is said* and *what is communicated* by an utterance take center stage in the explanation developed in this dissertation.

The challenges for the Classical Model are manifold, and many authors have resolved in their account to give up some parts of the Classical Model. Some deny the truth-conditionality of propositions. Some state that instead of a single proposition such indeterminate utterances express multiple propo-

sitions. Other account simply do not work for technical reasons. I will spend some significant effort elaborating these accounts in chapter 4. As far as the Classical Model and the explanation of communication goes, the accounts are found lacking, but of course they might have merits elsewhere.

Not satisfied with the solutions offered, I present the alternative in what I term the "Coarse-Grained Model of Communication" in part II. First, I'll make precise the concepts necessary to state the Coarse-Grained Model. The focus is on what questions and answers are, detailed in chapter 5. Since questions and their answers play a central role in the Coarse-Grained Model, it's quite important to establish what I mean by "a question", and what I mean by "an answer" to that question. Further concepts needed to state the Coarse-Grained Model will be clarified in chapter 6. The Coarse-Grained Model will be stated succinctly in chapter 7. The model takes up some elements from the contemporary accounts discussed before, but makes an important additional claim. In many cases, there is indeed a unique proposition which can be identified as the one communicated. This proposition is the satisfaction of a common purpose or goal, which in turn is represented by a question speaker and audience are trying to answer. Communication is thus "good enough" for the present purposes of the conversation, even if the expressions used are indeterminate. This can be achieved even if speaker and audience interpret the expressions slightly differently. By assuming that there is a question under discussion, either explicitly stated or otherwise implicitly generated, the model delivers a procedure to determine, from a context of utterance and an utterance, the proposition communicated. Clearly, the speaker's message on the Coarse-Grained model is pragmatic in nature and vitally context-dependent. The result of the model is divided into *what the speaker intends to express* and *what the audience infers the speaker to express*, which are determined quite similarly to one another, but from the speaker's and the audience's perspective, respectively. Here is one claim to be developed:

Claim *What the audience infers the speaker to express* with a literal assertoric utterance in a context is the union of all coarsenings of the question-answering interpretations the audience deems permissible.

This claim might seem obtuse at first glance and contains some idiosyncracies. I'll here give an intuitive description of the concepts involved, to provide the reader with an idea of the more detailed and technical descrip-

tion later on. A *literal assertoric utterance* is just a normal statement which does not make higher socio-cognitive demands like, e.g. irony does. An interpretation is a determinate meaning of the (possibly) indeterminate expression. For example, if the sentence uttered contains a vague expression, then an interpretation is a way of interpreting the expression precisely. We only look at interpretations the audience thinks are permissible—or compatible with the uttered sentence (among other constraints). An interpretation is question-answering if it provides an answer to the question under discussion. In other words, an interpretation is question-answering if it furthers the goal of the conversation in some way. A coarsening of an interpretation is that part of the interpretation which only pertains to answering the question under discussion. We take the union of all these coarsenings to unify the interpretation into one common proposition. Taking the union here is a set-theoretic notion, since propositions are represented as sets of possible worlds. Intuitively put, then, what the audience infers the speaker to express is a way to answer the question under discussion which agrees with all interpretations the audience deems compatible with the uttered sentence. The process to determine *what the speaker intends to express* is fairly similar, although here the focus is on those interpretations the speaker deems permissible. Successful communication occurs when both *what the speaker intends to express* and *what the audience infers the speaker to express* coincide. The result is a unique proposition communicated, thus explaining success with the Classical Model.

This simple result might seem to be reached too easily. I discuss in detail a host of objections to explaining communication with the Coarse-Grained Model in chapter 9. One such objection concerns the view that the Coarse-Grained Model resolves the indeterminacy by postulating something determinate: the denotation of the question under discussion. Simply put, the indeterminacy of the utterance is replaced with a determinate question under discussion. But, the objection goes, that is just shifting the blame towards something not explained any better. What is the justification for assuming a determinate question under discussion, might it not be just as indeterminate? In a sense, this is a sort of "higher-order indeterminacy," similarly to the notion of higher-order vagueness. Relatedly, an indeterminate question under discussion does not guarantee that both speaker and audience share in its denotation. In other words, the audience and speaker might

interpret the question differently, for example, because it might involve indeterminate expressions. Thus, a prerequisite of the model isn't satisfied and it's not applicable. But this situation might be widespread and could thus constitute a serious objection to the model.

Another objection questions the assumption of a common purpose in form of a question under discussion directly. It's clear that the technical notion of a question under discussion does a lot of the heavy lifting in the account. But it might be doubted that such a question mutually accepted by the discourse participants is a plausible part of many cases. For example, what if someone blurts out some entirely new information, having nothing to do with the previous topic of discussion? There does not seem to be a previously established question under discussion the utterance is intended to answer, and thus the account does not seem to even be applicable.

An important *reason* to stick to the Classical Model of communication is its seamless integration with semantics. Just as propositions feature prominently in explaining the meaning of a sentence, so do they feature in explaining the meaning of an utterance, which in turn is the central concept used in the Classical Model. The idea is that results in linguistics about the meaning of a sentence can verbatim be also applied to the meaning of an utterance. This, in part, makes the Classical Model so attractive. This works so well because usually, the meaning of a sentence is composed of the meaning of its parts combined by syntactical rules. The meaning of a sentence—the proposition—is compositional is this sense. But the meaning of an utterance as determined under this account is detached to some extent from the meaning of the uttered sentence, as it is determined also by the question under discussion. The meaning of an utterance is then not to the same degree compositional as the meaning of the uttered sentence, if at all, which takes away a reason to adopt the Coarse-Grained Model.

The model seems to take a particular implicit stance on the metaphysical status of what is meant. By decoupling the interpretation of an utterance from the context of utterance in form of a possible world, the model seems to lose some expressiveness. For in standard 2-D semantics, it is a fact of the matter that a certain expression has a certain intension at a context of utterance. For a given world, the meaning of that expression is fixed. This enables one to represent statements *about* the meaning of expressions as propositions in the same framework, i.e. as sets of worlds. But when

at, e.g., the actual world an expression can be validly interpreted as having many different intensions, what then should we make of a statement about the meaning of such an expression? What are its truth conditions?

The model determines meaning by excluding question-alternatives. Sometimes, though, statements which answer a question intuitively do not answer that question in this more technical sense. For example:

(4) a. Question. Who is awake?

b. Response. Kelly has a big day tomorrow.

Here, the statement only provides non-deductive information w.r.t. to the question. Non-deductive information is inductive, ampliative, or probabilistic. The statement provides evidence towards answering the question, without answering the question in the technical sense. But the model relies on answering a question under discussion to even assign meaning to the utterance. Thus for these types of question-statement pairs, the model does not seem to apply.

These are just some of the more prevalent objections discussed in chapter 9. I'll argue that there are plausible strategies of rebuttal for proponents of the Coarse-Grained Model for most objections discussed.

Finally, before getting into the main part of the thesis, some clarifications and hedges are in order. The term *model* here refers to the theoretical representation usually employed in philosophy of a more quotidian, worldly phenomenon. The term *model* here is parasitic on the term *model* used in scientific explanations and can be seen as a philosophical counterpart. The philosophical theory underpinning models in philosophy might lack in comparison to philosophical theory underpinning models in science. But nothing much hangs on the term model, it's just used to express the following similarities to scientific models: the philosopher develops a model which exhibits some properties deemed important to the phenomenon to be explained, explicated, analyzed, and so forth. This model necessarily makes simplifying assumptions about the thing it explains. In other words, the model idealizes to some extent.

Some premises will be assumed in this thesis. The most central one is that communication succeeds only if the speaker expresses something like a meaning, and the audience entertains or grasps the meaning so expressed. In other words, there is something like a meaning transmitted from speaker

to audience if communication is to succeed. This has historically been and still is controversial. Yet the thesis does not provide extensive argument for this premise, and instead operates under its assumption.

The cases under consideration will be tokens of direct literal verbal communication between two individuals, as the problems which amass in explaining this phenomenon are already plenty enough to fill a thesis. I will completely ignore the phonetics of verbal communication, and treat the issue as if sentences were spoken directly, disregarding any issues pertaining to voice and sound.

In the literature, there are a lot of expressions associated with communicated content. On the speaker's side, there are *what is expressed*, *what is said*, *what is implicated*, *what is intended*, etc. On the audience's side, there are *what is entertained*, *what is taken up*, *what is understood*, etc. I do not make fine distinctions about these different types of content in this thesis, and in particular use *what is said* and *what is expressed* interchangeably. With one exception: what is said literally and what is implicated can potentially come apart to such an extent as to be relevant for the purposes of this thesis. I talk more about this in chapter 9.

In communication, there are many different types of speech acts, as pertains to speech act theory. I will not engage with any more involved case here, and what I'm looking at are simple assertive illocutionary acts.

Lastly, a brief note on the way I visually represent content in this thesis. Contents here usually are propositions, represented as sets of possible worlds. For example, the propositions R, J, and M can be represented with a Venn-diagram like in Fig. 1.1. The propositions potentially intersect, and each possible world is either element of a proposition or its complement. That is, at each possible world, either the proposition, e.g., R holds or its negation $\bar{R}$. Thus, the propositions partition logical space. This can be represented in a Venn-diagram like in Fig. 1.2. Equivalently, this partition can be represented slightly more economically, like in Fig. 1.3. Propositions are represented as sets of worlds, or, equivalently, functions from worlds into truth values T, F. For some purposes, it will be helpful to make use of representing a proposition as a function. Then I'll display the propositions in a table, like in Fig. 1.4. Here, each world w_1 to w_8 is an element of a different partition cell. With these preliminaries, let's head into the history of the Classical Model in the following chapter.

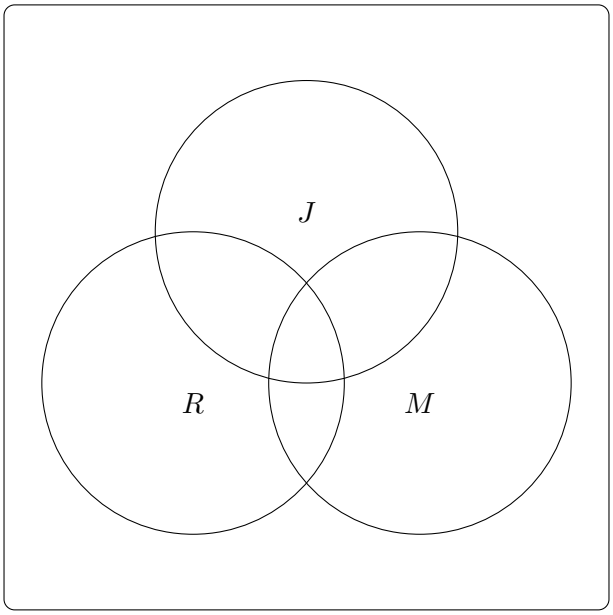

Figure 1.1: Propositions in a Venn-diagram.

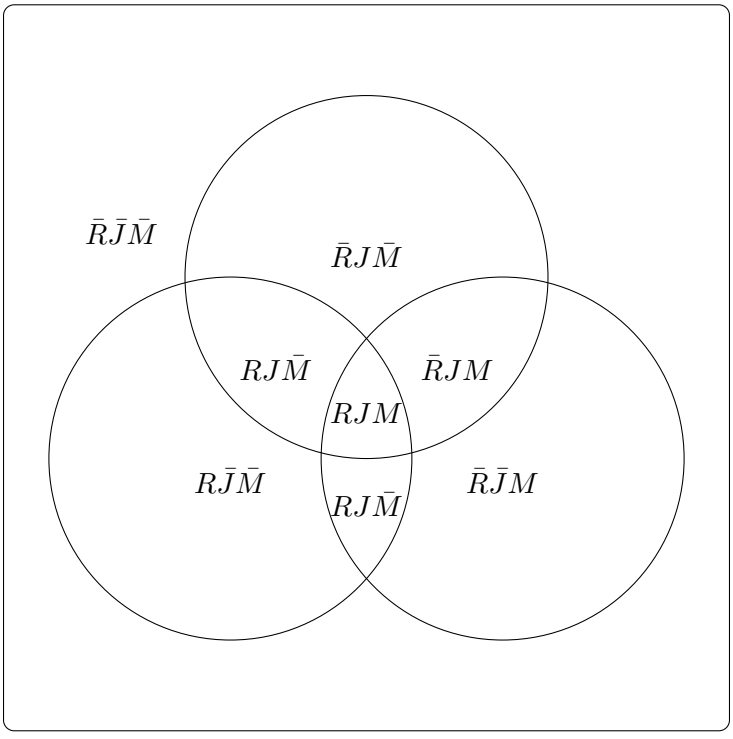

Figure 1.2: Propositions in a Venn-diagram partition logical space.

RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$
$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$

Figure 1.3: Alternative illustration of how propositions partition logical space.

	RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$	$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$
	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
R	T	T	T	F	T	F	F	F
J	T	T	F	T	F	T	F	F
M	T	F	T	T	F	F	T	F

Figure 1.4: How propositions partition logical space, as a table.

Part I

Models of Communication and Indeterminacy

Chapter 2

The Classical Model of Communication

2.1 Introduction

This chapter presents and develops the Classical Model of Communication. Since the model will be defended in this thesis, a whole chapter is devoted to its historical origins. This, of course, is not an exact science, but I think it instructive to relate the central model of the thesis to some historical ideas.

A modern formulation of the Classical Model, which I adopt for the purposes of this thesis, is given by MacFarlane (2020a, 2020b, 2020c).

Classical Pragmatics

(1) The content of an assertion is a (single) proposition.

(2) Uptake consists in recognizing the proposition asserted.

(3) If the assertion is accepted, its content is added to the conversational common ground.

Classical Content

Contents are ways the world might be (truth conditions).

A few clarificatory remarks are in order. The situations explained by the Classical Model are, to make matters simple, one-to-one verbal communications between a speaker and an audience. The terms "speaker" and "audience" denote *roles* relative to a particular utterance. Two people conversing might repeatedly switch roles. The history of the conversation up to

the particular utterance under consideration is important, as the context of the utterance is affected by it. But the Classical Model is not an account of the complex unfolding dynamics of a conversation, like, e.g., Groenendijk and Stokhof (1991), Heim et al. (1983), Kamp (1981b), and Roberts (2006) offer. In the formulation of the Classical Model, the speaker and the audience are left implicit in Classical Pragmatics (1) and (2). "Uptake", for example, denotes uptake by the audience.

What does this Classical Model provide, exactly? First, it explains what's going on in the case of literal communication in terms of the transmission of *content* from a speaker to their audience. This content is truth-conditional and context-dependent. Second, it provides a way to integrate the dynamics of a conversation by way of specifying the effect of an utterance on the contextual situation. Third, the model gives a criterion of successful communication. It is necessary for successful communication that the audience grasp the content expressed by the speaker. What does the model say about sufficient conditions? It seems that Classical Pragmatics 1-3 together with Classical Content suffice for success. This can be debated, of course, and has much to do with how, exactly, one spells out the technical terms involved (e.g. "assertion"). For the purposes of this thesis, it is important that the Classical Model is to impose necessary conditions on successful communication. That is, for every case of (intuitively) successful communication, the Classical Model shall be satisfied. This entails that there should be a single proposition expressed by the speaker which is taken up by the audience. It is not required that for every such case, the Classical Model provides such a proposition. Yet, we'll encounter cases in this thesis in which authors claim that there *cannot be* a single proposition expressed and taken up. This thesis argues that even in these cases there in fact is a single proposition communicated, and provides some extended means of determining this proposition. The thesis thus provides an argument for the Classical Model by way of constructing the single proposition required.

In the following, Classical Pragmatics (1) and (2) is found to have roots reaching as far back as Aristotle reaching as far as Locke. Classical Content, i.e. truth-conditionality, is attributed to Frege and following authors making up the then nascent analytic philosophy. Treatment of the contextual impact on communication in Classical Pragmatics (3), finally, is attributed chiefly to Grice and Stalnaker.

2.2 The Content of an Assertion is a Proposition

Gauker (1992) describes what he calls the "Lockean theory of communication": communication succeeds if the audience grasps something expressed by the speaker, yet distinct from the speaker's words. Under the assumption that this something can be described as a proposition, Gauker's Lockean theory comes close in content to the conjunctions of Classical Pragmatics (1) and (2): The content of an assertion is a (single) proposition, and uptake consists in recognizing the proposition asserted. This section traces the Lockean Theory through philosophical history.

In *De Interpretatione*, part of the *Organon*, Aristotle expounds that "spoken words are the symbols of mental experience and written words are the symbols of spoken words," (16a4-5) according to the translation by E.M. Edghill (Aristotle, 2001). More familiar to the reader, perhaps, is the translation by J.L. Akrill of the same sentence, from whom I will adopt all following citations: "now spoken sounds are symbols of affections in the soul, and written marks symbols of spoken sounds" (Aristotle, 1985). In the same paragraph, Aristotle can be interpreted to give an explanation of successful communication: since affections of the soul (or mental experiences) and the things they are likenesses or experiences of are *the same* for everyone, successfully communicating is then "just" a matter of using the right symbols to express one's affections.

Kretzmann (1974, p. 3) values this passage as the "most influential text in the history of semantics." Affections of the soul are later called "thoughts", even if both terms might not denote the exact same concept (Kretzmann, 1974, footnote 12). Aristotle further states that in communicating, "the speaker arrests his thought and the audience pauses," (De Interpretatione, 16b20) which Kasabova and Marinov (2016, footnote 6) take to clearly indicate a transmission of thought from speaker to audience, encoded into spoken word by the speaker and decoded by the audience. How do spoken words, symbols and signs come to be associated with thoughts, for Aristotle? What determines the "code" according to which content relates to expressions? For Aristotle, that is convention (De Interpretatione, 16a19, 16a27). Linguistic expressions are significant by convention, whereas affections of the soul are like actual things without conventional variability, they instead signify naturally. For Aristotle and the later scholastic tradition, (spoken)

words (and sounds) were either nonsignificative or significative. Significative words were further divided in those which are significative naturally, like laughs and groans, and conventionally significative, like words in a language (Ashworth, 1981, p. 305). Two affections of the soul or mental impressions, as Kretzmann calls them, instantiated in different individuals, have a single Aristotelian form: they are two tokens of the same type (Kretzmann, 1974, p. 11). This is the condition for successful communication, one might infer: when two individuals associate mental impressions of the same type with the expression, the speaker successfully transmitted their thoughts to the audience. An interpretation of what makes communication successful, for Aristotle, might then be the following. The speaker expresses a thought (or mental experience, affection of the soul, mental impression) with their spoken words. The audience associates the spoken words with that thought by convention. Since thoughts are 'the same for all', the audience takes up the thought the speaker expressed (and thus communication succeeds).

For example, the name "Nikomachos" is not a natural sign of the person Nikomachos, but comes to be symbolic—and thus to be a name at all—by convention. What does the name symbolize? A mental impression which is a likeness of Nikomachos. Since mental impressions are "the same for all," the speaker can make themselves understood if the audience associates the same mental impression as the speaker does. If additionally both employ the same conventions, communication may succeed.

Kretzmann (1974) gives the following summary of relations:

(A1) Spoken sounds are symbols of mental impressions.

(A2) Spoken sounds are (in the first place) signs of mental impressions.

(A3) Mental impressions are likenesses of actual things.

For Aristotle, then, words are primarily signs—natural symbols—of mental impressions. They might still signify actual things secondarily. Some authors have later noted that these types of relationships are not dissimilar to something like a division into (primary) meaning and (secondary) reference (Ashworth, 1981).

The following are crucial premises in Aristotle's account:

(A4) Spoken sounds are not the same for all.

(A5) Mental impressions are the same for all.

(A6) Actual things are the same for all.

The justification for (A5) is problematic. On what grounds should one adopt the premise? Aristotle states that mental impressions are the likenesses of actual things, actual things are the same for all (A6). Let's grant the latter claim (A6). What makes two likenesses of an actual thing identical? Clearly, any rational reconstruction of a likeness or similarity relation will fail to guarantee sameness of two things similar to a third. (Kretzmann, 1974, p. 11) gives an illustrative example, compare Fig. 2.1.

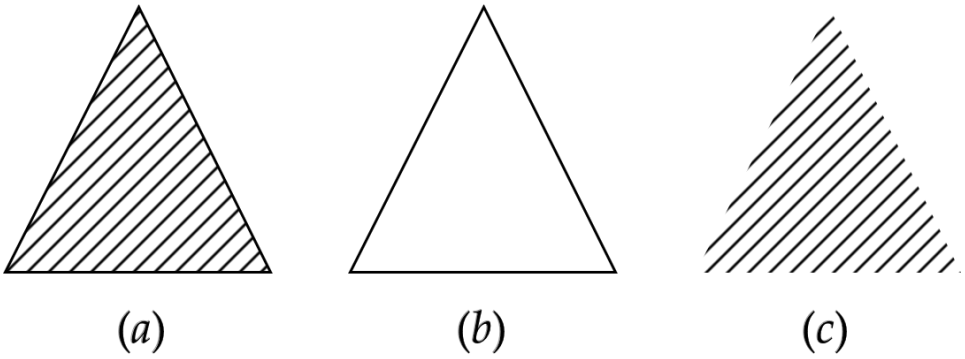

Figure 2.1: A similarity relation is not an equivalence relation.

One thing (Fig. 2.1a) is similar to another (Fig. 2.1b) and to a third (Fig. 2.1c), but the latter two are not similar, and surely not identical. Consequently, the sameness of actual things does not guarantee the sameness of their likenesses, and thus not the sameness of mental impressions. For Kretzmann, this indicates that Aristotle was indeed *not* intending the beginning of De Interpretatione as a theory of communication. Another move is to adopt a radical externalist position. According to the radical externalist, communication succeeds *in every case* only because both speaker and audience stand in the right relation to a mind-independent, external object they are both referring to. Then the sameness of mental impressions would not matter for Aristotle's explanation, since it's the sameness of external referent which facilitates communication.

For Boethius, who provided commentary of Aristotle's work in the earliest middle ages, the "whole arrangement" concerning verbal communication consists in "things, thoughts and spoken sounds" (Commentaries on On Interpretation, 20,11-16), in the translation of Andrew Smith (Smith, 2014). *Things* are ontologically fundamental. Thoughts are conceived from,

or about, things, and thus secondary. Spoken words signify thoughts. Mediated by thoughts, then, words are about things as well. Boethius makes the following claims: "Those who use the same spoken sounds have the same thoughts in their mental conception [...] But this does not hold conversely. For those who have the same thoughts do not automatically have the same spoken sounds [...]" (21,11-13). For the latter claim Boethius provides an example. A Roman and a Greek might see the same horse, and have the same thought that it is a horse, but use different words to talk about it. The first claim, on the other hand, strikes one as rather peculiar from a modern perspective. Why should not two people use the same words, but have different thoughts they want to express, e.g. two Romans talking about different horses? It seems that the individuation of thoughts is the issue here. Perhaps one can construe it such that both Romans have the same thought, e.g., with an indexical component. "That, which is in front me, is a horse" would then be the same thought, but the actual horse referred to is not constitutive part of the thought. But since thoughts are parasitic on (or "preceded by") the actual things thought about, and "thoughts are always of things" (21,3)[1], it would be hard to argue with Boethius that the horse referred to is not part of what the thought is about. It was stipulated, though, that these horses are different. Apart from this, Boethius explanation of communication follows Aristotle's exposition very closely (as one might expect), and does not seem to add any new perspective relevant to this thesis.

The following millenium of medieval scholars shall not be part of this exposition; certainly there would have been one or the other subtle and potentially interesting theoretical point to be made. But see Eco (1984) for an overview. The next step in the historical journey is the early modern treatment of philosophy of language by John Locke.

In his important "An Essay Concerning Human Understanding" (from here on "ECHU"), Locke (1694) presents his theory of meaning and communication. For Locke, "words, in their primary and immediate signification,

[1] Although just below, Boethius, in the Andrew Smith translation at least, contradicts this and states that "Nor is there always a subject for a thought, for there are thoughts which have no corresponding thing, e.g. the centaurs or chimaeras invented by poets" (22,2-6). I suspect some hidden contextual equivocation, such that this latter use of "thing" stands for something like "real or existing thing." The original latin uses "res" for both potential senses of "things," though: "rerum enis semper intellectus sunt [...]" (21,2-3) and "sunt enim intellectus sine re ulla subiecta [...]" (22,3-4) (Boethius & Meiser, 1987).

stand for nothing but *the ideas in the mind of him that uses them*, how imperfectly soever or carelessly those ideas are collected from the things which they are supposed to represent" (ECHU, III.2.2, original italics). An *idea* is the primary signification of a term, and something private: "[...] words [...] can properly and immediately signify nothing but the ideas that are in the mind of the speaker [...]" (ECHU, III.2.4). Ideas are what thoughts are made of, and what is made known to others in communication (ECHU, III.2.1). As Kretzmann (1968) notes, Locke is somewhat liberal with the terminology to describe the relation of words and ideas. Words "are signs of, or are marks of, or are names of, or signify, or mark, or correspond to, or are annexed to ideas" (Kretzmann, 1968, p. 179). In the following Kretzmann proposes to use a neutral "signify" for all these different senses, with Locke's intended meaning to be developed. For Kretzmann, Locke's statement that words signify ideas is the main *thesis* of the semantic theory. The argument for the thesis is "from the uses of words", i.e. from their role in communication. Locke is intending to explain the meaning of words via their function in communicating, and what is communicated are the ideas of the speaker. This way of putting things invites immediate objection, like famously John Stuart Mill (1843) formulated in A System of Logic (ASoL):

> When I say, 'the sun is the cause of day,' I do not mean that my idea of the sun causes or excites in me the idea of day; but that the physical object, the sun itself, is the cause from which the outward phenomenon, day, follows as an effect. (ASoL, I.2.1)

For Mill, names are of things themselves, and not of the ideas of things. Locke, on the other hand, at least in the context of communication, deems this use of words 'perverted':

> [...] it is a perverting the use of words, and brings unavoidable obscurity and confusion into their signification, whenever we make them stand for anything but those ideas we have in our own minds. (ECHU, III.2.2)

This seems to indicate that for Locke, words do *not* stand for things themselves other than ideas at all. But, as I will argue, the Lockean theory has more in common with Aristotle's than it might seem at first glance.

Mill's reading of Locke's theory is not as charitable as it could be. If words signify ideas [and nothing but ideas], then Locke's theory is only

about a very special type of communication. Namely when the speaker reports what they are thinking, or what they are feeling, or what they are experiencing. Thoughts, feelings, and experiences are the subjects of this type of communication. This talk is *about* ideas. Yet most normal talk is not about ideas, but about the world around us, at least on any common-sense account of communication. Locke would be guilty of an "incredible oversight" (Kretzmann, 1968, p. 183). Attributing such a blunder to Locke is not as charitable as possible. And it ignores an important part of the thesis quoted above: words *properly and immediately* signify nothing but ideas. It is rather plausible that "immediately signify" here is to be understood as "signify directly, without intermediary". That words can improperly or mediately signify things other than ideas is thus compatible with Locke's thesis. And Locke seems to intend as much, as Kretzmann elaborates.

> For our Names of Substances being not put barely for our *Ideas*, but being made use of ultimately to represent Things, and so are put in their place, their signification must agree with the Truth of Things, as well as with Men's *Ideas* [...] (III,11,24)

So at least for Lockean substances, words signify not only ideas, but also mind-independent things. Roughly, Locke distinguishes between complex ideas of "substances" and of "modes", where substances are supposed to have independent existence, and modes exist only derivatively (Uzgalis, 2022). Without going too much into Locke's metaphysical account, substances include ordinary objects which form the substrate of complex ideas, that in which ideas "subsist" (II,23,1). Modes, on the other hand, include relational properties of substances and other modes. It's enough for present purposes that indeed, talk about ordinary objects is used, ultimately, to signify things. So it is not the case that words (talking about ordinary objects) signify nothing but ideas, only that words signify nothing but ideas *immediately*. Note that there are dissenting voices. For example, Landesman (1976) argues that signifying and signifying immediately are not two subtly different relations, but instead identical. That is, whenever something is signified it is immediately signified and *vice versa*. Additionally, Lockean ideas are not mind-dependent, private objects, but have a referential component as well and are "outside the mind". Landesman thus rejects the distinction between signifying immediately and mediately.

What, then, are the proper connections between words, ideas, and things which facilitate communication, in Locke's theory? For Ashworth (1981), it's important to study Locke's writing in the historical context. Students at Oxford were "still reading scholastic texts in the mid-seventeenth century [...] written as university text books by Roman Catholic philosophers, predominantly Jesuit, who consciously placed themselves within the tradition of medieval philosophy and theology while at the same time making use of sixteenth-century developments in Aristotelian studies." Ashworth argues that Locke's theory of language is heavily influenced by late scholastic works. The notion that words signify ideas is Aristotelian in spirit, claims Ashworth. Because neither takes the historical context into account, Ashworth deems incorrect both Kretzmann's and Landesman's analysis of Locke's theory of meaning, although Kretzmann is "nearer the truth" than Landesman. Ashworth (1981, p. 311) notes that the scholastic discussions make "[...] much better sense if one thinks of 'signify' as being replaceable by 'make known', 'reveal', 'manifest', or 'express' than if one tries to equate 'signify' with 'mean'." The reason Ashworth gives is that apparently, "Do words mean ideas or things?" is not a sensible question, whereas "Do words make known ideas or things?" is. Asking for a meaning instead is asking for a translation or definition, so Ashworth.

To address the question whether words make known things or ideas, Ashworth discusses some relevant commentators in the scholastic tradition. Most seem to agree that both things and ideas are signified. They differ in the details: For Smiglecius (1618), words immediately signify things and only mediately ideas. For Couto (1606), a word has two kinds of immediate signification, namely both ideas and things. There is a duality in response upon hearing a word uttered, perceiving actual things as well as ideas of the speaker. Finally, Burgersdijk (1632), Masius (1617) and of St. Thomas (1930) state that words immediately signify ideas in the mind of the speaker, and only mediately objective things in the world. Locke seems to follow the third group of authors as is clearly indicated in the passages from ECHU quoted above. Yet Ashworth makes an additional point: since Locke was aware of the writings of the authors mentioned and saw himself in their tradition, he was not putting forth a radical new thesis about semantic meaning. Locke's claim is not about the meaning of words, but about making known, just as in the writings of scholastic authors. The meaning of a word is in-

stead a definition produced to answer the question "What does this word mean?", so Smiglecius claims (Ashworth, 1981, pp. 318). And since "Nearly everything [Locke] says about language in the first two chapters of Book III [of ECHU] is closely parallel to the scholastic texts which were read at Oxford when Locke was an undergraduate and a tutor there" (Ashworth, 1981, pp. 325), Ashworth attributes this position also to Locke. Ashworth also provides direct evidence: "the defining of words [...] is nothing but declaring their signification" (III,3,10). However, Ashworth is open to the idea that definition consists in complex ideas.

For the purposes of the current thesis, either interpretation will do: whether words *mean* ideas immediately or only *make known* ideas immediately does not crucially alter the explanation of communication we can attribute to Locke. For it seems that in order to communicate, a speaker ought to make their ideas known to the audience. What is made known might be, but does not have to be, fully determined by the expression's meaning.

Words immediately signify private ideas, and objective and intersubjective things mediately. Locke argues *from* communication *to* his notion of meaning (or making known, if one follows Ashworth), and does not seem to provide a criterion for successful communication. But it's certainly appropriate to interpret Locke as claiming that if the speaker is indeed successful in making their ideas known to the audience, communication succeeds. How, though, can the private ideas be known by the audience, if all they have to go by is the spoken word? One premise adopted in this thesis is that successful communication requires some thing, some meaning, to be transmitted, i.e. expressed by the speaker and correctly grasped by the audience. What is this thing, for Locke? If ideas are private, we are lacking any criterion for when the private idea the speaker expresses is the one the audience comes to entertain. After all, the audience might come to entertain a slightly or entirely different idea. But since the idea is private, the audience does not have a method for determining whether they entertain the idea the speaker expressed. And neither does the theorist trying to give an account of successful communication. For Aristotle, as described above, this is "solved" by requiring the relation between the mental entity, the affection of the soul, to be a natural sign of the thing it represents. Natural signs are not arbitrary and thus not private, such that the problem does not arise in the same way.

While the simple critique of Locke's theory of language, as argued above, doesn't hold up, it is in explaining successful communication that Locke's theory is insufficient.

2.3 Contents are Ways the World Might Be

And Gottlob Frege indeed objects to something like Locke's theory of communication, argue Miller (1998) and McCulloch (1995). To see this, first some background on Frege's theory of language. For Frege, it was of utmost importance to clearly distinguish between the psychological on the one hand and the logical on the other, or the subjective and the objective (Frege, 1953, p. xxii). With a normal meaningful linguistic entity, Frege associates a sense (*Sinn*) and a reference (*Bedeutung*). A sense, in Frege's terminology, "contains" the mode of presentation, a way in which the referents are presented. Referents are the objects denoted, e.g. individuals, places, instants, or stretches of time. Frege's gives an example: three lines a, b, and c which meet in the same point. Then "the point of intersection of a and b" and "the point of intersection of b and c" have the same referent—the point at which all three lines meet" (Frege, 1948, p. 210).[2] Yet they have different senses. Sense is used to explain why these two expression have different cognitive value, i.e. provide different information, even though they refer to the same point. Similarly, identity statements of the form "$a = a$" and "$a = b$" clearly differ in cognitive value, even under the assumption that a and b refer the same object. To each linguistic expression in a precise formal language there should correspond a definite sense and in turn a definite referent, even if Frege acknowledges that natural languages often exhibit context-dependency. Other exceptions are possible: an expression can have a sense but lack a referent. For example, "the largest integer" appears to have a sense to be grasped but does not refer to any number. So far, Frege's account seems to be in lockstep with Locke's. Can't we simply assume that senses are Lockean ideas? After all, senses are about cognitive value, which appeal to human cognition. And doesn't everyone understand expressions in their own, private way? It is a crucial point for Frege that senses are *not* to be equated with private ideas. Instead, senses are intersubjective, "common

[2]Most terminology for Frege's notions is taken from Max Black's translation (Frege, 1948) of (Frege, 1892). For translator's notes, see Black (1948); see Church (1948) for commentary on the translation.

to many," (Frege, 1953, p. 212) and "[...] objective, inasmuch as it can be used by several [...]" (Frege, 1953, p. 213). That is, senses are not private and not part of an individual mind. Lockean ideas resemble much more closely what Frege calls a conception (*Vorstellung*). These are private and include such things as mental images. "A painter, a horseman, and a zoologist will probably connect different conceptions with the name 'Bucephalus' [the horse of Alexander the Great]" (Frege, 1953, p. 212). Senses are something "in between" reference and conception:

> The referent of a proper name is the object itself which we designate by its means; the conception, which we thereby have, is wholly subjective; in between lies the sense, which is indeed no longer subjective like the conception, but is yet not the object itself. (Frege, 1953, p. 213)

Complete, literal indicative sentences without subordinate clauses express a *thought* (*Gedanke*) as their sense, and they refer to a truth value as their denotation (Frege, 1948). It bears repeating that although "thought" seems to indicate something psychological and relative to an individual, thoughts—and senses—for Frege are intersubjective and thus shared between individuals. Interestingly, Frege goes one step further and states that two people *cannot* have the same conception, even though they may grasp the same sense. This seems irritating at first, since it's intuitively possible that two different people have the same conception of a thing, even if that conception doesn't figure in the theory of language. But this irritation is quickly resolved, since Frege adds indices to conceptions, at least a person and a time. On the supposition that two different people cannot be the same person under any circumstance, two people cannot have the same conception in this trivial sense. Proper names, i.e. names denoting individuals, refer to those individuals. For example, "Aristotle" refers to Aristotle. Their sense is supposed primitive and unexplained. The referent of a predicate is a function, and which kind of function depends on the type of predicate. The referent of a monadic, i.e. 1-ary predicate of the form $F(x)$, maps an individual onto a truth value. For example, the referent of "is a philosopher" maps the referent of "Aristotle" onto the truth value *true*, or "the true", as Frege calls it. These types of functions are also called "concept" (Frege, 1952). The referent of the sentence "Aristotle is a philosopher" is then the true. Zalta (2024) ar-

gues that the sense of sentences is similarly compositionally established. The sense of "$a = a$" would thus be a composition of the senses of "a", "=", and "a" again, whereas the sense of "$a = b$" would be composed of the senses of "a", "=", and "b". This distinction yields the grounds for an explanation that both sentences do not have the same cognitive significance, although they refer to the same truth value (again, on the supposition that both a and b refer to the same object.) Frege's work on the theory of language doesn't end here, of course. For example, much of (Frege, 1948) deals with indirect speech or what we would today call intensional contexts. Additionally, Frege acknowledges the pragmatic *force* of different types of speech, like assertion, thus predating speech act theory, or so Dummett (1996, chp. 10) argues.

What is Frege's criticism towards Locke's theory of language? Mainly, it's that communication becomes impossible if we are to account for the meanings of an expression in terms of a private item like Lockean ideas. That is, at least, what Miller (1998) attributes to Frege, following McCulloch (1995). A similar critique is mounted against the Lockean explanation of communication just above. This can be read as a reductio in favor of Frege's theory of language, like so:

(P1) A major task of any theory of language is to explain successful communication between individuals.

(P2) Communication has to be explained by giving an account of meaning of expressions which are transmitted from one individual to another.

(P3) Private Ideas cannot be transmitted from one individual to another.

Thus. Private ideas cannot explain successful communication.

Hence. Private ideas cannot be used to fulfill a major task of a theory of language.

Does the argument have merit? Even though the version presented here is sloppy and does not strictly yield a logically valid argument, it is easy to see that one can produce such an argument on pains of legibility. So let's treat the argument as valid. The premise (P3) is probably the least controversial, since it is analytic, dependent on the definition of "private" one uses. It can be argued that a theory of language is mostly about semantic properties of sentences and their parts, and communication does not fall under its purview,

thus rejecting (P1). Instead, it would be the task for pragmatic theories of communication. However, this can hardly be a position we attribute to Frege. Premise (P2) is a statement which can be counted as part of the so-called "code model" of communication, after Shannon and Weaver (1949), see section 2.4 below. This thesis accepts (P2) and takes it as the basis for further theorizing, as described in the introduction. Of course, there are authors rejecting the code model, for example Sperber and Wilson (1986) and Gauker (2002), whose accounts will not be discussed further in this thesis.

What does Frege's account offer, so as not to fall prey to the argument, too? As is apparent from the discussion above, Frege explains the meaning of expressions as senses. Senses, for sentences, are thoughts, and thoughts are intersubjective and not private. It seems that the argument presented even *requires* Frege to adopt this position. What kind of thing is a thought, that two individuals may grasp it? Thoughts aren't mental objects, since they are intersubjective, abstract objects. The process of grasping a thought is, however, psychological and can differ of person to person. Frege "infamously" never specified what the components of thoughts, i.e. modes of presentation of objects, are (Heck & May, 2006, p. 25). Kripke (1972), interpreting Frege, claimed that the sense of a proper name is given by a definite description. This seems to be in line with an example Frege provides: the sense of "Aristotle", could be taken to be "the pupil of Plato and the teacher of Alexander the Great" (Frege, 1948, p. 210). However, according to Heck and May (2006) it is doubtful whether Frege held such a view at all; they cite Dummett (1973, chp. 5) as a prominent critic, among others. In any case, it's clear that for Frege sense determines reference. So for a sentence, the thought determines the truth value of the sentence. Thoughts are the objects of propositional attitudes. In the case of successful communication, the speaker expresses a thought with their utterance, which the audience then grasps. This might still fail, however, and in particular due to the inexactness of natural languages. The audience might come to associate a different thought with the sentence. But for two individuals to grasp the same thought is for them to be in the same psychological state (w.r.t. the thought). This also fixes the truth value. In the case of proper names, this assumption was famously challenged by Putnam (1975), who argued "with a little help from science fiction" that two individuals can associate the same

sense with a name while their associated extensions differ.

The Ludwig Wittgenstein of the Tractatus (Wittgenstein et al., 1990) very much agrees with Frege's call for an idealized language. So much so that Baldwin (2006) ascribes to Dummett the attribution of the early Wittgenstein's theory of meaning to Frege, although, of course, the details differ: Wittgenstein deems ordinary language to already have ideal logical structure, even though that structure might not be obvious: "Alle Sätze unserer Umgangssprache sind tatsächlich, so wie sie sind, logisch vollkommen geordnet" (Wittgenstein et al., 1990, prop. 5.5563). And instead of an abstract object, thoughts are the "logical picture of facts": "Das Logische Bild der Tatsachen ist der Gedanke" (Wittgenstein et al., 1990, prop. 3). Whether Russell (1905) agrees with Wittgenstein that thoughts are representations or what is sometimes called a Russellian proposition—a complex composed of objects, properties and relations—shall not matter for the present purposes.

It is Carnap (1947) who introduces an explication of the notion of meaning which is closely related to the more recent model of propositions as sets of possible worlds, and which Church (1951) used to explicate Frege's notions of sense and reference (Leitgeb & Carus, 2023). Rudolf Carnap assigns to sentences an extension as well as an intension. Intensions, most generally, are functions from state descriptions to extensions. State descriptions contain "for each atomic sentence either this sentence or its negation" (Carnap, 1947, §2), i.e. for each propositional variable p, either p is member of the state description or not-p is. In this way state descriptions are similar, but not identical to, possible worlds (Kripke, 1963a; Kripke, 1959; Kripke, 1963b). For Carnap, there is no variability of intension between individuals:

> The term 'property' is to be understood in an objective, physical sense, not in a subjective, mental sense; the same holds for terms like 'concept', 'intension', etc. (Carnap, 1947, §4)

For a sentence, the intension is a proposition which are "objective, nonmental, and extra-linguistic entities" (Carnap, 1947, §6). Putnam complains that *grasping* an intension is left entirely unexplained (Putnam, 1975, pp. 263), and then hypothesizes grasping to be a subjective psychological process. For Carnap, grasping is a "logical or semantical" operation, "in its technical form it is based on the semantical rules concerning the given expression" (Carnap, 1947, pp. 202). This leaves open the possibility that different individuals

apply different semantic operations, thus associating a different proposition with an uttered sentence. Communication for Carnap, however, would seem to require that a speaker expresses an intension and the audience grasps that intension, and in addition, that the communicated sentence has *true* as its extension. For this it seems paramount to agree on a conceptual framework, which is a "construed object language with a hierarchy of metalanguages in which to define and explore truth, analyticity, synonymity, designation and other semantic resources of the object language" (Leitgeb & Carus, 2023). Of course, Carnap was mostly concerned with frameworks for scientific language.

Not a philosophical, but a mathematical model of communication is presented by Shannon and Weaver (1949), which Carnap (Carnap & Bar-Hillel, 1954) develops into a semantic theory and in particular influenced Dretske (1981). In the communicative situation according to the model by Shannon and Weaver, a source is sending a message to a destination. In order to do so, the message is encoded into a signal and transmitted via a channel, which is subject to interfering noise, potentially corrupting the signal. The signal is received and then decoded, such that the original message can arrive at the destination. The focus of Shannon and Weaver is on the probabilistic properties of the information contained in a signal. The encoding and decoding processes are taken as given. Sperber and Wilson (1986) term this and similar approaches "the Code Model." They translate the situation the model describes onto verbal communication between speaker and audience. The speaker's central thought processes produce a thought which is then linguistically encoded into an acoustic signal. The signal is transmitted via the air and subject to interfering noise. The audience decodes the acoustic signal into a thought, again, which is entertained by their central thought processes. The focus in this thesis is what makes communication possible in the entire absence of noise and does not make the same assumption that speaker and audience command perfectly symmetrical encoding and decoding processes. Thus the Code Model is applied to separate issues from the ones discussed in this thesis. It is helpful, though, to conceptualize the situation of verbal communication between speaker and audience as it is in the Code Model, even if the assumption about symmetrical encoding and decoding is dropped. For with the additional assumption that the channel is noiseless, it brings into focus what is at issue here: what is it that is

transmitted between speaker and audience, and how may the audience grasp it?

W.V.O. Quine holds that what a speaker means with their utterance is not a matter of objective truth. There are in each case multiple competing hypotheses equally valid, each relative to their respective scheme of translation. The available empirical evidence with regards to speaker meaning is "radically indeterminate" (Quine, 1960, chp. 2). Science, therefore, has no place for meaning and propositional attitudes, and the latter should be disregarded in favor of empirical study of behavior (Baldwin, 2006, p. 78). Quine would thus not assent at all to the project of this dissertation, rejecting the central premise that to explain communication is to explain what is transmitted from speaker to audience. This thesis is not the place to defend against Quine's challenge from underdetermination.

Two further prominent authors disagree with the idea of transmitting a meaning from speaker and audience to facilitate communication: Donald Davidson and the later Wittgenstein. Davidson (1967) outright rejects that there is some entity assigned as the meaning of a word, and provides Fregean functions for unsaturated predicates as an example. For Davidson, giving the truth conditions of a sentence is a way of giving the meaning of that sentence. These truth conditions are not to be confused with intensions as described above, however. Truth is analyzed along broadly Tarskian marks, i.e. giving a metalinguistic translation of the form "s is T if and only if p", where s is a sentence of the object language and, T a truth-predicate and p a sentence expressing the meaning of s in the meta language. However, in explicating meaning, Davidson cannot use on pains of circularity the notion of synonymy—sameness of meaning—to explain the relation between s and p. What, then, renders a sentence like (S) not part of a theory that "accounts for the workings of natural language" (Davidson, 1967, pp. 311)?

(S) 'Snow is white' is true if and only if grass is green.

After all, the biconditional (S) is in fact true. Davidson thus requires theories of meaning to be *holistic*, in the sense that sentences in a languages are meaningful only in relation to a theory of meaning for all sentences of that language. To achieve such a theory is to align it with empirical evidence of a speaker's utterances. Davidson makes use of Quine's theory of radical translation in what he terms "radical interpretation" (Davidson, 1973).

The focus in analyzing meaning—and communication—is then not on some meaning-entity which is transmitted from speaker to audience, but instead of general patterns in the behavior of the speaker (Malpas, 2023). Successful communication consequently does not require any sameness of content between speaker and audience, as Davidson recognizes the potential failure of convention leading to incommensurable idiolects (Davidson, 1986). Instead, communicative success is a chiefly pragmatic matter (Baldwin, 2006, p. 85).

The Wittgenstein of the *Philosophische Untersuchungen* disregards treating meaning as a represented entity as useful for the philosophy of language entirely, as it cannot do justice to the richness and variety of ordinary language (Wittgenstein, 1953, I §1 pp.). Instead, he holds that "die Bedeutung eines Wortes ist sein Gebrauch in der Sprache"—the meaning of a word is its use in the language (Wittgenstein, 1953, I §43). This is meant, of course, as a slogan for a more involved view. Wittgenstein's "language games" are self-contained systems establishing their own idiolect of a language. This language can be characterized by a set of rules. However, any statement of these rules is compatible with any conduct, as the rules first have to be interpreted (Wittgenstein, 1953, I §201 pp.).[3] Wittgenstein gives the example of a hypothetical student who he taught to continue series of numbers (Wittgenstein, 1953, I §185 pp.). The student is asked to continue a series "+2" after 1000, and the student answers "$1004, 1008, 1012, \ldots$", for the student believes "+2" to involve the rules "add 2 up to 1000, add 4 up to 2000, and so forth". Wittgenstein argues that interpreting to follow a rule correctly is "inescapably indeterminate" (Baldwin, 2006, pp. 88). Wittgenstein proposes to determine rules by investigating tokens of common practices of the participants in the language game. The meanings of expressions thus emerge through careful study of ordinary language use. For a theory of communication this has immediate effects. These find expression in what is often reconstructed as Wittgenstein's famous "private language argument," which has been the subject of much debate.[4] Rule-following, and therefore speaking meaningfully, is not possible "in private" (Wittgenstein, 1953, I §202), but only as a fundamental human activity, viz. as part of a community (Beaney, 2006, p. 54). Successfully communicating it seems, then, is taking part in

[3] It might be important to note that Wittgenstein's notion of "interpretation" differs from the standard model-theoretic one used throughout this thesis.

[4] See, for example, Kripke (1982), McDowell (1989), and Stroud (2000).

the same language game following *public* rules. But if rules are constituted by actual practice, just as meaning is actual use, doesn't this entail that communication cannot fail? For it seems that whatever action performed, there can be a set of rules which the action satisfy. Here again Wittgenstein defers to the "practices of the speakers engaged in the language-game in question, which will include the cases which are picked out as paradigms, the types of evidence taken to be relevant, the authority of different speakers," and so forth (Baldwin, 2006, pp. 88). We only have rules where there is agreement (Wittgenstein, 1953, I §224), so it seems that successful communication at least requires some agreement on the rules of the language game.

2.4 The Content is Added to the Common Ground

One of the most important figures which can justifiedly be counted as contributing to the Classical Model in the 20th century is Paul Grice. Grice (1957) gave an answer to the attack on "traditional" philosophy of language, in which meaning and truth conditions take center stage,[5] mounted by philosophers prioritizing ordinary language as the focus for philosophy.[6] The conflict can be framed thus: first, distinguish between *sentence-meaning* and *speaker-meaning*. Sentence-meaning is the literal meaning of a sentence and can often be characterized by a sentence similar to the following: "Snow is white" is true iff snow is white. Sentence-meaning is closely linked to a sentence's truth conditions. "Speaker-meaning," on the other hand, is a term to describe meaning that isn't truth-conditional, at least not obviously so. Prime examples of utterances in which speaker-meaning seems to be what explains that the sentence uttered is meaningful are questions, commands, expressions of approval or disapproval, and so on. For example, an expressivist position in metaethics denies that moral judgments are subject to truth conditional analysis and instead liken them to expressions of the form "Φ? Buargh!". Still, this utterance is meaningful—it seems to communicate something and can be understood. The conflict stems from deciding which of the two types of meaning is *explanatorily fundamental* (Miller, 1998, pp. 318). Grice contributed to this debate in a substantial way, and it seemed to be his

[5]Proponents include Carnap (1947, 1963), Frege (1948), Kaplan (1964), Kripke (1963b), Lewis (1970), Montague (1960, 1968, 1970), Scott (1970), and Tarski (1935).

[6]Proponents include Austin (1961, 1962), Ryle (1949, 1953), Strawson (1950, 1952), and Wittgenstein (1953, 1991).

goal to re-establish truth-conditional semantics as emerging from pragmatic analysis. He is firmly in the second camp:

> [...] the causal theory ignores the fact that the meaning (in general) of a sign needs to be explained in terms of what users of the sign do (or should) mean by it on particular occasions; and so the latter notion, which is unexplained by the causal theory, is in fact the fundamental.[7]

Put differently, he was to reconstruct sentence meaning *from* speaker meaning. Whether one judges Grice's efforts successful, the analysis of speaker meaning alone is quite influential. One formulation, given by Strawson (1964), is the following:

> [A person] S [...] means something by an utterance x if S intends
>
> (i1) to produce by uttering x a certain response r in an audience A and intends
>
> (i2) that A shall recognize S's intention (i1) and intends
>
> (i3) that this recognition on the part of A of S's intention (i1) shall function as A's reason, or a part of his reason, for his response r.

This is a set of conditions which are jointly sufficient for S to mean *something* by their utterance. That is, if S intends all (i1)-(i3), they mean something by their utterance. Why are all three necessary, and not only (i1)? Grice (1957) gives some examples. Consider: S might leave B's handkerchief at the scene of the crime intending to induce the detective to believe that B was the murderer. But Grice does not want to say that S *meant* by leaving the handkerchief that B was the murderer. The case is violating an additional condition, which Grice specifies in (i2). S does not intend to for the detective to recognize their intention of type (i1)—that would undermine the deceitful plan. But even (i1) and (i2) together do not cover all cases.

[7] "Causal theory" here refers to a causal theory of utterance meaning in the style of Stevenson (1944), according to which "for x to mean$_{\text{NN}}$ something, x must have (roughly) a tendency to produce in an audience some attitude (cognitive or otherwise) and a tendency, in the case of a speaker, to be produced by that attitude, these tendencies being dependent on an elaborate process of conditioning attending the use of the sign in communication." (Grice, 1957, pp. 379). This is distinct from what's commonly called the "causal theory of reference" developed later by Geach (1969), Donnellan (1970) and Kripke (1972).

Because Grice wants to distinguish between cases of "deliberately and openly letting someone know" or "getting someone to think" on the one hand and "telling" on the other. These three examples are given to elucidate.

(i) Herod presents Salome with the head of St. John the Baptist on a charger.

(ii) Feeling faint, a child lets its mother see how pale it is (hoping that she may draw her own conclusions and help).

(iii) I leave the china my daughter has broken lying around for my wife to see.

Herod intends to produce in Salome the belief that John the Baptist has been killed, and intends that Salome recognizes Herod's intention for her to believe so. Yet, Grice argues, by presenting the decapitated head to Salome Herod does not *mean* that John the Baptist was killed. Instead, Herod only lets her know or gets her to think that John the Baptist was killed, but does not tell her so. Neither Herod's intention of type (i1) nor (i2) are part of the reason for Salome to form her belief. This claim is presented without argument by Grice, and seems on first glance to lack justification. For a very good reason to form a belief on any account, really, is confirming evidence. But Herod presenting the decapitated head seems to be extraordinarily good evidence for Salome to believe that John the Baptist was killed. Miller (1998, pp. 322) tries to make things clearer by interpreting "part of the reason" to instead say "play a part in the explanation". He cashes out "play a part in the explanation" counterfactually. If Herod wouldn't have had the intentions of type (i1) and (i2), would Salome still believe that John the Baptist had been killed? For Miller then the simple possibility that Salome could have come to believe the news in some other way is enough to conclude that Herod presenting the head does not play a part in the explanation of Salome coming to believe the news. Wharton (2008) gives a more detailed analysis of what Grice might have had in mind. In cases like (i) above, there are two "layers" or types of information available to the audience. The first layer is basic and consists of the information being presented to the audience. In this case, the decapitated head itself. The second layer concerns intentions and consists of the intention of the speaker that the information of the first layer is presented to the audience. In the problematic cases, the intention to inform is not crucial to the audience being able to derive the intended

response, as the first—basic layer— is already sufficient. We therefore cannot ascribe to Herod the intention that Salome form her belief because of his intentions—the second layer—since she might just as well derive the intended belief from the decapitated head—the first layer—alone. For Grice, then, the head itself might mean something, but in an altogether different fashion than the meaning analyzed by Grice's conditions (i1)-(i3). Indeed, Grice picks up Aristoteles' distinction mentioned in the beginning of this chapter between natural and conventional meaning, where the latter is dubbed "non-natural" meaning. Grice uses these considerations to develop a notion of sentence-meaning. In a nutshell, he requires there to be some regularity in the association of a sentence and speaker-meaning. This regularity is to appear for different occasions of use for the same speaker, resulting in the "timeless meaning for an individual" (Grice, 1968, pp. 231). The regularity should also appear across different speakers within a group of individuals (e.g. a language-community), thus addressing the conventional meaning of an utterance-type.

Some ordinary language philosophers also held, besides claiming that speaker-meaning is the fundamental notion, that it is a methodological mistake to analyze many speech acts with truth-conditional methods. Grice defused this tension by providing a comprehensive account of what he calls, and has come to be known since, "conversational implicature." What is implicated (e.g. implied, suggested, indicated etc.) is to be distinguished from what is said by an utterance. An implicature goes beyond the literal, truth-conditional meaning of the sentence uttered. This notion makes it possible to systematically describe non-literal uses of language which appear frequently in day-to-day communication. Consider Grice's example:

> I am reporting on a pupil at Collections. All I say is "Jones has beautiful handwriting and his English is grammatical." We might perhaps agree that there would here be a strong, even overwhelming, implication that Jones is no good at philosophy. (Grice & White, 1961)

Grice doesn't *say* that Jones is not a good philosopher, but merely implicates it. He *says* that Jones has beautiful handwriting and that his English is grammatical. Grice provides a simple test: implications of this sort are cancelable. That is, it is felicitous to negate the candidate implicature, for

example by saying that "Jones has beautiful handwriting and his English is grammatical, and indeed he is also good at philosophy." This is crucially different from other types of implied content, for example in the case of suppositions. "Jones' thesis was decent philosophy" supposes that Jones wrote a thesis. "Jones' thesis was decent philosophy, but in fact he never wrote one" seems infelicitous, failing the cancelability test.

Grice distinguishes between *conventional* and *non-conventional* implicature. Conventional implicature holds in virtue of the meaning of the words, as Grice puts it. One example is "He is an Englishman and therefore brave" (Grice, 1975, p. 44). The implicature, here, is that being brave is a consequence from being an Englishman. Grice stresses that this is not part of what is said, but what is implicated: "But while I have said that he is an Englishman, and that he is brave, I do not want to say that I have *said* (in the favoured sense) that is follows from his being an Englishman that he is brave, though I have certainly indicated, and so implicated, that this is so" (Grice, 1975, pp. 44). Conventional implicatures aren't usually cancelable. One central species of non-conventional implicature is *conversational* implicature. These can be particularized, i.e. their intended meaning is dependent on features of the particular context of utterance. The example about Jones' philosophical skills is of this type. To infer the implicature, the relevant context is needed: this is a reference letter from one professor to another, and is supposed to highlight Jones' philosophical skills. On the other hand, *generalized* conversational implicatures *normally* differ from the compositional meaning of the uttered words. Grice's example is "X went into a house yesterday and found a tortoise inside the front door" (Grice, 1975, p. 56). Without any further information about the context, the audience assumes that X went into a house *other* than their own house. But the sole meaning of the uttered words does not yield any information about whose house it was. It is merely information about the normal use of a phrase like this which licenses the inference that a house other than X's is referred to here. Intuitively, this is because the speaker would've described the house as X's own house if that were the case. Grice gives a detailed and influential analysis of this intuition. He grounds this analysis in a postulate, called the *Cooperative Principle*:[8]

> Make your conversational contribution such as is required, at the

[8]Grice's maxims and the Cooperative Principle are quoted from (Grice, 1975).

> stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

Participants of a discussion are implicitly expected to observe this principle, Grice claims. By this expectation, the audience is put in a position to infer the meaning-intention of the speaker. What is the purpose or direction of the exchange? Grice lists a *question for discussion* initially proposed as one possibility. The account developed in this thesis will circle back to ideas of this kind and make heavy use of it. For Grice, other types of purposes are also valid, however lenient, with the requirement that *some* contributions to the exchange are to be regarded as non-purpose-driven. This means that the purpose is defined operatively and only insofar as at least some contributions do not further the exchange towards the—however indefinite—goal of the conversation. Which contributions count as valid under the Cooperative Principle? Grice provides four general categories of more specific maxims, named *Quality*, *Quantity*, *Relation*, and *Manner*. I'll discuss each in turn as they are to play a role in justifying claims of the present thesis. The category named *Quantity* comprises the following maxims.

> (Qn1) Make your contribution as informative as is required (for the current purposes of the exchange).
>
> (Qn2) Do not make your contribution more informative than is required.

Utterances—conversational contributions—are supposed to provide the right amount of information. The notion of information is not further specified, but it is safe to assume that at least considerations of the following sort hold: a conjunction generally is at least as informative as either of its conjuncts. For example, "Alice goes to a classical concert tomorrow" is more informative than "Alice goes to a concert tomorrow." If the purpose of the exchange is represented, for example, by the question *for* discussion "What does Alice do tomorrow" and, let's say, it is contextually salient that Alice's options are given by the cultural offering of Vienna, either utterance can arguably be said to satisfy (Qn1) and (Qn2). If, however, the purpose is rather "What kind of concert does Alice attend tomorrow?", then "Alice goes to a concert tomorrow" fails (Qn1)—it simply doesn't provide any new information. Grice comments that (Qn2) might not be necessary, as it might be subsumed under (Re), as will be discussed below. The category *Quality*

comprises the following maxims.

(Ql1) Do not say what you believe to be false.

(Ql2) Do not say that for which you lack adequate evidence.

This maxim seems entirely reasonable. The audience should expect the speaker, in a cooperative situation, to be sincere and truthful. Truthful in addition not by accident, i.e. not by arbitrarily guessing, but instead by basing their beliefs and corresponding utterances on sufficient evidence. It is, as one might guess, not part of Grice's theory to provide explications of evidence and evidential support, but it seems any run-of-the-mill epistemological theory will be compatible. The most explanatorily load-bearing maxim is the one of *Relation*, simply

(Re) Be relevant.

Relevant to what, exactly? Grice has in mind that contributions should be relevant to the purpose of the conversation. This is expressed, for one, in (Qn2), following which the conversational contribution shall not be more informative than is required by the purpose of the conversation. The maxim requires for the contributions to be directly related to the purpose of the conversation. Expecting the speaker to observe maxim (Re) is a powerful tool: many conversational contributions can be made sense of with the help of (Re). Consider:

A: Smith doesn't seem to have a girlfriend these days.

B: He has been paying a lot of visits to New York lately.

Grice explains: "The speaker implicates that which he must be assumed to believe in order to preserve the assumption that he is observing the maxim of relation" (Grice, 1975, p. 51). In this case, it's the indication that Smith may have a girlfriend in New York. B's contribution would seem to be irrelevant to A's contribution if not for assuming the implicated content. A's contribution must, accordingly, be understood to set the direction of the conversation. In a sense, A poses an open question whether Smith has a girlfriend. I'll have much to say about this kind of process in the course of this thesis. Of course, A's contribution could be interpreted differently, for example as a fervent statement. But without additional contextual clue, the canonical interpretation Grice gives is most plausible. Grice adds a

fourth element to the list of categories making up the Cooperative Principle, *Manner*, under which he subsumes the following maxims:

(Ma1) Avoid obscurity of expression.

(Ma2) Avoid ambiguity.

(Ma3) Be brief (avoid unnecessary prolixity).

(Ma4) Be orderly.

These are but examples, says Grice, and summarized by "be perspicuous", i.e. easy to understand. That this is a set of maxims not often valued as highly in practice as the others is immediately obvious to anyone who ever set foot in an institute of philosophy. Grice states as much himself: "[...] the observance of some of these maxims is a matter of less urgency than is the observance of others" (Grice, 1975, p. 46).

Now, what's with all these maxims? How do they help determine what was said or implicated? Grice takes great care to list many examples, which, by mere common-sense reasoning, deliver an intuition about what is being implicated. In some cases, certain maxims are violated intentionally (or "flouted"), leading to apparent failure of communication. But, to Grice, these just show the strength of his account. For example, consider a case of irony, adapted from Grice's. X, a friend of Y, has betrayed Y in some fashion. Y and their audience both know this. Y says "X is a *fine* friend." The description of the situation here includes that it's obvious to both Y and their audience that Y's utterance isn't what Y believes. Y is saying something false, flouting (Ql1). But instead of the communication just falling flat, Grice argues, the mutual assumption of rationality and thus the assumption of following the Cooperative Principle suggests that Y implicated something true. He states: "This must be some obviously related proposition; the most obviously related proposition is the contradictory of the one [Y] purports to be putting forward" (Grice, 1975, p. 53). So, Y ostensibly flouts maxim (Ql1). But the audience is capable of determining the implicature by assuming that they must mean something in accordance with the maxims.

As a side note, there seems to be some leeway of interpreting Grice's theory on way or the other, and his reasoning *in this original paper from 1975* is not completely obvious to me. Grice divides the cases explained roughly into three groups: Group A, in which no maxim is violated, Group B, in which there is some maxim violated, but which is to be explained in terms of

a clash of two competing maxims, and Group C, in which a maxim is flouted demonstratively. The case of Smith's girlfriend belongs to Group A, the case of irony to Group C. My confusion has to do with whether *what is said* has to satisfy the maxims or *what is implicated*. If *what is said*, or the direct, literal sentence-meaning has to follow the maxims, then B's answer in the Smith case arguably violates maxim (Re). Without its implicature, B's answer doesn't relate to the purpose of the conversation. But this is an example of Group A, so there is no violation of a maxim. Then seemingly *what is implicated* should satisfy the maxims. Yet, take a look again at the friend case: here, *what is implicated* is the negation of Y's utterance's content. But the negation does not violate any maxim; the literal meaning of the original utterance does. Its classification, though, is into Group C, in which maxims are flouted, that is, intentionally and demonstratively violated. I propose for the purposes of this thesis to sideline the issue of categorization of cases into groups as outlined. For the present purposes, it's enough to establish the basic outline: what has been said literally seems to violate some of the maxims. By assuming that the speaker is yet following the maxims, the audience infers what the speaker actually meant—implicated—by a sort of rational reconstruction of their intention.

Grice has been seen as contributing substantially to the discussion on the divide between semantics and pragmatics. The distinction between syntax, semantics, and pragmatics (commonly traced back to Morris, 1938) is a recurring theme in the philosophy of language. A general proclivity to sort phenomena and methodologies into these categories does not necessitate their metaphysical reality, and discussions surrounding related issues are not the focus of this thesis. That these categories can generally be understood and that certain considerations have a general tendency to be assigned to either of these categories is, however, assumed. Roughly, syntax has to do with structural properties of signs, semantics has to do with relations between signs and signified entities, and pragmatics has to do with the use of signs and the users of signs (Bezuidenhout, 2006). Although, of course, much more elaborate distinctions are possible and have been made. Grice proposed a way to explain sentence-meaning on the basis of speaker-meaning. Speaker-meaning has to do with the use of language by the speaker and as such is a prime candidate to be categorized as pragmatic. Some even take this to be a defining difference between semantics and pragmatics (Gauker,

2012). That there is a sharp dividing line, however, is doubted by many recent authors (e.g. Carston, 2002; Recanati, 2003). The discussion where to draw this line is quite orthogonal to the aims of this thesis, so I'll sideline this issue as well.

One key takeaway is that to explain communication and establish what it is that is communicated, the mere literal meaning of a sentence, however defined, is often inadequate to explain a whole range of phenomena. Some of these have been highlighted by Grice. This thesis will ultimately argue that pragmatic considerations may have an effect on the truth-conditional content of an utterance in circumstances which will be delineated amply; roughly, those in which the literal meaning of an utterance—its sentence-meaning—does not provide adequate means, even in context, to determine the truth-conditional content. A phenomenon to this effect has been termed "pragmatic intrusion" (Levinson, 1995, 2000). This phenomenon has some parallels to the "encroachment" of pragmatic considerations in questions of knowledge ascriptions (Kvanvig, 2011). Whether these pragmatic intrusion are a matter of *what is said* or *what is implicated* has been the subject of some debate as well (Carston, 2013; Recanati, 1989).

This still leaves open the question with which tools to analyze the pragmatic component of communication, insofar as it has anything to do with explaining successful communication. Stalnaker developed, in a series of papers, a method to model conversational contexts and conversational contributions using broadly the tools of truth-conditional analysis (Stalnaker, 1974, 1975, 1978, 1986, 1988, 2002, 2004). Other such accounts drawing "both on the ideas of the anti-formalist Grice-Austin tradition and on the resources of the kind of formal semantics that those philosophers were reacting to" have been developed, which Stalnaker lists duly:

> There is situation semantics (Barwise & Perry, 1983), discourse representation theory and file change semantics (Heim, 1982; Kamp & Reyle, 1993), update semantics, dynamic predicate logic, dynamic Montague grammar (Groenendijk & Stokhof, 1990, 1991), among others. (Stalnaker, 1998, pp. 3)

Stalnaker claims that his approach is foundational with respects to describing the dynamics of discourse. Contexts play two roles in the description of discourse: on the one hand, the context has an effect on the contents of

an utterance—or more generally, on a speech act. On the other hand, the context is "acted upon" by the speech act, i.e. what has been contributed to the conversation usually influences the context against which the next contribution is to be evaluated. Stalnaker then proposes to "[...] identify a context (at a particular point in a discourse) with the body of information that is presumed, at that point, to be common to the participants in the discourse" (Stalnaker, 1998, p. 5). The body of information is represented with a "context set" of possible worlds compatible with the information. The body of information is also called the "common ground." What this body of information includes is, of course, context-dependent. It is information that discourse participants take for granted. For a discussion about Formula 1, for example, it may include the information that Lewis Hamilton signed for Ferrari in 2025. Stalnaker adds: "But the information presumed to be common will also include facts about the discourse that is currently taking place, since when one is engaged in a conversation, one can normally take for granted that speakers and audiences are aware that the conversation is taking place, that speakers are saying what they are saying in the way they are saying it" (Stalnaker, 1998, pp. 5). It is assumed that everyone in the discussion knows these pieces information, and that everyone knows that everyone knows and so forth. There is mutual presupposition: if I take it to be common ground that ϕ, then I also take it to be common ground that it is common ground that ϕ, and so does my conversation partner. Stalnaker also speaks of common ground as presupposition: "To presuppose something is to take it for granted, or at least to act as if one takes it for granted, as background information—as common ground among the participants in the conversation" (Stalnaker, 2002).

Lewis (1975) gives an influential analysis of accommodating presuppositions: "Say something that requires a missing presupposition [to be acceptable], and straightway that presupposition springs into existence, making what you said acceptable after all [...]" Accommodation, according to Lewis, describes that the "conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play," at least under the assumption of co-operative intent. The conversational score meanwhile keeps track of commonly presupposed propositions and other contextual factors. It's clear, then, that central attributes of the context are to be explained with reference to the participants' propositional

attitudes. The common ground can be modified through moves in the conversation, i.e. speech acts, and foremost among these is the act of *assertion* (Stalnaker, 1978, 2004). If an assertion is accepted by the participants, its content is added to the common ground. For example, if Phoebe utters "I saw an interesting movie last night", that *Phoebe* uttered the sentence is part of the common ground. Thus the indexical "I" can be resolved to Phoebe, and similarly the reference to a place in time. Note though that Stalnaker's account does not provide specific mechanisms to resolve context-dependent expressions. Instead, it provides the grounds to license the audience to infer the intended content. The utterance can be assigned a definite content in the present context. If Phoebe's utterance is accepted, the content is added to the common ground. Possible worlds in which Phoebe watched a boring movie or no movie at all are discarded from the context set. This way, the context affects the content of the utterance, and in turn is affected by the utterance. I'll necessarily gloss over many issues that deal with details of Stalnaker's account; for the time being, it suffices to present the central notions of common ground and context set.

Chapter 3

The Problem of Indeterminacy

3.1 Introduction

Problems for the Classical Model appear when the proposition expressed proves difficult to determine. This happens in cases of what I call *utterance-indeterminacy*. Utterance-indeterminacy occurs if the linguistic meaning of the uttered sentence type together with the context of the utterance does not suffice to determine a single proposition expressed. It will be assumed that the notions "linguistic meaning" and "context" are reasonably well-defined to work in this definition; in any case, I have more to say about these notions as the thesis develops. I argue that the cases I'll discuss are cases of *apparent* utterance-indeterminacy. I argue this by showing that the linguistic meaning together with the context, suitably spelled out, does in fact suffice to determine a single proposition expressed. Hence, the cases I'll discuss are cases of utterance-indeterminacy as just defined only apparently—at first glance. For reasons of readability, I will leave out the "apparent" in most places when discussing utterance-indeterminacy. MacFarlane (2020c) calls such cases in which communication still succeeds in spite of utterance-indeterminacy *felicitous underspecification*. Note that this is not to suggest that, e.g. vagueness does not constitute genuine semantic indeterminacy. Utterance-indeterminacy also does not, qua its definition, require any stance towards judging the cases w.r.t. metaphysical indeterminacy. Metaphysical indeterminacy is understood as that for some proposition p, there is no fact of the matter whether p (Taylor & Burgess, 2015). Utterance-indeterminacy and metaphysical indeterminacy, so understood, are logically independent.

The notion of utterance-indeterminacy goes further than Carston's (2002, p. 19) thesis of *underdeterminacy*. Carston's underdeterminacy thesis states that "linguistic meaning underdetermines what is said". Linguistic meaning, for Carston, is "the information encoded in the particular lexical-syntactic form employed", and what is said is "the thought or proposition which [the linguistic meaning] is being used to express" (Carston, 2002, p. 17). But usually, for Carston, the context provides enough information to determine what is said. Carston also considers a pragmatic form of indeterminacy much like utterance-indeterminacy, but associates it with Grice's implicature (Grice, 1975, p. 40) and Sperber and Wilson's weak implicature (Sperber & Wilson, 1986, pp. 195). Utterance-indeterminacy can arise in virtue of different phenomena. Broadly, I'll distinguish three categories of utterance-indeterminacy in the sense above.

3.2 Semantic Indeterminacy

First, utterance-indeterminacy can arise due to semantic indeterminacy, e.g. due to vague or imprecise expressions. I'll focus on vagueness in this discussion. Some similar considerations can also apply to imprecision, generality, or to ambiguity. For some authors, these phenomena even exhibit the same cognitive features (Dunbar, 2001).

Consider:

(5) Richard is tall.

Which truth conditions does an utterance of (5) express? One can infer with some certainty that Richard is not small, at least. But "small" exhibits similar semantic properties. There are two obvious ways in which an utterance of (5) introduces utterance-indeterminacy. The utterance might be indeterminate because the context is lacking a standard to evaluate Richard's height against. Is Richard tall for a basketball-player or for a horse racing jockey? To distinguish this type of case from the ones discussed below, I'll assume that such a standard has been contextually established or, where there isn't, that the issue is otherwise unproblematic. The other way in which an utterance of (5) introduces utterance-indeterminacy is if its truth is unclear, even though, say, Richard's exact height is common knowledge. No matter which standard for "tall" the context established, there are conditions under which the truth of (5) seems indeterminate. Richard in this

scenario is what's in the literature usually called a "borderline case". That it is indeterminate whether Richard is tall or not seems to be a semantic feature of the predicate. Note that ambiguity doesn't usually present the same problem for the explanation of communication. For in many cases, ambiguous expressions—with multiple distinct meanings—are used to communicate successfully by disambiguation. The context provides enough clues for the audience to decide between the multiple meanings of the ambiguous expressions, thus determining a proposition inferred to be the content of the utterance. In cases where the audience cannot discern such, usually the communicative exchange is said to fail.

Now, what's the problem for the Classical Model, exactly? The Classical Model requires determinate truth conditions expressed by the speaker to be taken up by the audience. Knowing the truth conditions of a sentence is knowing under which conditions the sentence is true. If Richard height is h meters, and to be h meters tall is to be a borderline case of the predicate "tall" (by the standard set by the context), then there seems to be no fact of the matter whether "Richard is tall" is true (note that this is different from there being no fact of the matter whether Richard is tall). Thus, for the conditions under which Richard is h meters tall, the audience does not know whether "Richard is tall" is true and hence does not know its truth conditions. Is there a way for the audience to infer the sentence's truth conditions? Suppose that the speaker actually intends to express determinate truth conditions, i.e. has in mind, or the ability to decide, for every possible height of Richard whether "Richard is tall" is true. Since many such "precisifications" of truth conditions are potentially expressible with (5), the audience has to decide between them if they are to take up determinate truth conditions. By assumption, neither the literal meaning nor the context provides enough clues to decide between these precisifications. Thus choosing the precisification the speaker expresses is a matter of luck for the audience. But communicative success of utterances like (5) do not seem to be a matter of luck; their success seems to be systematic. What gives? Standard accounts of vagueness do not seem to offer a solution. As an aside, this thesis does not take a stance on whether any of the standard accounts of vagueness are correct. I'll merely argue here that taken at face value, the most straightforward explanations of communication resulting from either standard account don't seem to square with the Classical Model.

Kit Fine (1975) proposed to "not only consider the truth-values that sentences actually receive but also the truth-values that they might receive under different ways of making it completely precise." Vague predicates are indeterminate: some cases do not receive a truth value. Thus, there are ways of making the predicate precise by assigning a determinate truth-value to every case. Fine develops a logic based on the notion of *super-truth*. A sentence is super-true if it is true under all ways of making it precise. Since this thesis will pick up concepts and themes from Fine's account to model indeterminacy, it is instructive to take a closer look.

Fine assumes an "intuitively understood, but possibly vague" first order language L. To deal with vague expressions, *partial specifications* assigning truth values to atomic sentences of L may also assign a truth value I for *indefinite*, besides T and F for *true* and *false*. Specifications are *appropriate* if the meanings of the predicates are intuitively understood. Fine's example: "Yul Brunner is bald" is assigned *true*, "Mick Jagger is bald" assigned *false* and "Herbert is bald" assigned *indefinite* in case Herbert is a borderline case of "bald". The central question which Fine addresses is how the semantics of such a language might work. In what way are truth values assigned to formulas, given a specification of truth values (*true*, *false*, and *indefinite*) for all atomic sentences? In particular, the issue is what happens for composites like $P \wedge Q$ if either part has indefinite truth-value. Fine briefly considers truth-functional assignments. On such an approach, the truth-value of a compound is entirely determined by the truth-values of the individual formulas. One way of spelling this out is a "maximization" approach:

$f_\wedge$	T	I	F
T	T	I	F
I	I	I	F
F	F	F	F

The function $f_\wedge$ assigns truth values for a conjunction. For example, $P \wedge Q$ is assigned I if (the truth-value of) P is T and Q is I. But $P \wedge Q$ is assigned F if P is F and Q is I. Here, then, definite truth values are assigned just in case the truth value does not depend on how the indefinite truth value is made definite. In a sense, indefinite truth-values "dither". Opposite sits the "minimization" approach:

$f'_\wedge$	T	I	F
T	T	I	F
I	I	I	I
F	F	I	F

On this approach, indefinite truth-values "dominate". Whenever an individual formula is indefinite, so is the compound. Logics on the basis of these types of three-valued assignments have been studied extensively, perhaps most notably by Kleene (1952) and Lukasiewicz (1930). But, Fine argues, any truth-functional approach does not account for what he calls *penumbral connections* between indefinite sentences. These are (potentially obvious) logical connections between sentences which get obscured when they are assigned an indefinite truth-value. For example, an atomic sentence P might be indefinite. However, we would expect $P \wedge \neg P$ to be false, not indefinite. But on the most plausible truth-functional approach to negation, either approach to conjunction above renders $P \wedge \neg P$ indefinite:

$f_\neg$	
T	F
I	I
F	T

Fine then instead presents his influential supervaluationalist approach. Here, similar to a Kripkean possible world model for intuitionistic or modal logic, Fine gives up truth-functionality. Simplifying slightly, a specification as defined above can be *extended* by another specification. The extension preserves all definite truth values. This means that a specifications u extends a specification u, written $t \leq u$, if and only if

$$\text{If } t \models A \text{ then } u \models A, \text{ and}$$
$$\text{If } t \dashv A \text{ then } u \dashv A.$$

Here, "$t \models A$" says "t assigns *true* to A" and "$t \dashv A$" says "t assigns *false* to A". If a specification assigns only definite truth-values, its called *complete*. Complete specifications are classical, and any specification can be extended into a complete specification. Here is the central statement, then:

$$t \models A \text{ iff for all complete } u \text{ with } t \leq u : u \models A, \text{ and}$$
$$t \dashv A \text{ iff for all complete } u \text{ with } t \leq u : u \dashv A.$$

In other words, a sentence is super-true (or super-false) at a specification if and only if it is true (or false) at all complete extensions.

What is the effect of this semantics? Very straightforward. Definite atomic sentences at a specification keep their truth-values at all extensions. Indefinite atomic sentences at a specification are extended into different truth values. For example, an indefinite atomic sentence without any additional constraints is true at one extension and false at another. The interesting effect is for compounds like $P \wedge \neg P$. If P is indefinite at a specification t, then there are, let's assume, complete extensions t' and t'' at which P is true and false, respectively. But since complete extensions are classical, $P \wedge \neg P$ is false at both t' and t'', and so at all other complete extensions. The sentence $P \wedge \neg P$ is hence false at t, too, even though P is indefinite at t. Similar arguments apply for sentences like $P \vee \neg P$, $P \supset \neg P$, etc. This, then, works for all classical logical truths. Merely contingent sentences with vague predicates, however, aren't normally super-true: there is, in most cases, a precisification at which the sentence is true (or false), even if at other precisifications the sentence is false (or true).

How would the supervaluationist propose to explain communication? First, we'd have to extend Fine's extensional account into an account of truth *conditions*, that is, an intensional one. I assume that this is straightforwardly possible, as does Fine (Fine, 1975, p. 274). Fine uses partial specifications to assign indefinite truth-values to atomic sentences. In analogy, define *partial interpretations* as a function assigning partial specifications to possible worlds. To make this case in the simplest possible way, let's assume the language consists of just one atomic sentence "Mick Jagger is bald". The extension at a possible world can then be further simplified to just a single truth-value: *true*, *false*, or *indefinite*. Interpretations thus can extend one another in the intuitive way in analogy to specifications. Also, call an interpretation *complete* if it assigns only the definite truth-values *true* and *false*. Suppose we order (classes of) possible worlds by the number of hairs on Mick Jagger's head, s.t. at world w_1, Mick Jagger has a certain low number of hairs, while at w_2 he has more hair, at w_3 even more and so on. Then while at w_1, Mick Jagger is definitely bald, whereas at w_8, say, he definitely isn't. A partial interpretation i thus assigns *true* for the definite case of baldness, *false* for the definite case of non-baldness, and *indefinite* for others. See Fig. (3.1) for an illustration.

	...	w_1	w_2	w_3	w_4	w_5	...
i		T	I	I	I	F	

Figure 3.1: Assignments of truth values by partial interpretation i to "Mick Jagger is bald".

This partial interpretation can be extended into complete interpretations in different ways, depending on where to set the boundary for being bald. Some complete interpretations c_1 to c_4 extending i are illustrated in Fig. (3.2). Assuming that, for each interpretation, a bald Mick Jagger can't have more hair than a non-bald Mick Jagger.

	...	w_1	w_2	w_3	w_4	w_5	...
c_1		T	F	F	F	F	
c_2		T	T	F	F	F	
c_3		T	T	T	F	F	
c_4		T	T	T	T	F	

Figure 3.2: Assignments of truth values by complete extensions c_1 to c_4 of i to "Mick Jagger is bald".

There are two straightforward options for the supervaluationist to explain what is communicated with the help of truth conditions. First, what's communicated could be the partial interpretation i. This seems to be a good candidate, at first glance: only definite information is communicated, and the speaker does not take a stance at all towards borderline cases—the speaker remains indifferent. The audience is then free to take up any complete interpretation extending the partial interpretation. But it might be doubted that this describes correctly what intuitively is communicated. I'll have more to say on this in later chapters. The crucial point for current purposes is that a partial interpretation only determines partial truth conditions. Partial truth conditions do not work in the Classical Model—at least not *as is*. Instead, the Classical Model requires definite truth conditions.

Second, what is expressed could be a *set* of complete extensions. That is, the speaker would communicate multiple propositions, not just one. The audience might take up one or more of these propositions. In some sense the communication could be said to be successful if *some* of what the speaker expressed is taken up by the audience. This approach is similar in structure to the *Multiple Propositions View* (or *Cloudy Picture*), which will be discussed in section 4.2.2. In any case, this approach does not work in the Classical Model either—as multiple, mutually incompatible truth conditions do not uniquely determine truth conditions communicated. To be fair, a proponent of this approach could develop a process how to determine unique truth conditions from the different complete extensions.

Both the first and second option don't satisfy the Classical Model for another reason. They fail to satisfy *Classical Pragmatics (3)*, repeated here:

Classical Pragmatics

(3) If the assertion is accepted, its content is added to the conversational common ground.

The speaker is not in a position to know which of the possible truth-conditions the audience takes up. What happens then is that the update of the common ground with the newly asserted information is only common to both speaker and audience by luck or accident, if at all. The Classical Model requires unique, bivalent truth conditions in part to set up the dynamics of the conversation (MacFarlane, 2020a).

For epistemicists like Schiffer (1999), Sorensen (2001), and Williamson (1994), whether Herbert is bald might be a case of ignorance. In such a case, we do not know that Herbert is bald and we do not know that Herbert is not bald. We are ignorant of something: either "Herbert is bald" expresses an unknown truth or "Herbert is not bald" does (Williamson, 1994, p. 187). For Williamson, "Herbert is bald" is either true or false—vague utterances are bivalent. But in these borderline cases, even after we established all the ways of measuring and deciding on Herbert's baldness, we are in a state of ignorance about whether Herbert is bald. And this, moreover, *necessarily so*, Williamson argues.

I agree in this thesis with the epistemicists that neither the speaker nor the audience *know* where the precise boundary between baldness and non-baldness lies, if there is any. Nor does the speaker have to decide, for them-

selves, on a precise boundary and neither does the audience. On the Coarse-Grained Model, communication succeeds because the precise boundary *does not matter* for the practical purpose of the conversation. The semantic content of the vague sentence only constrains, but does not determine, what is communicated in a particular exchange.

However, this thesis does not stipulate the unknowability of such a boundary, even though it'd be compatible with the Coarse-Grained Model developed. The problem marrying the epistemicist account of vagueness to the Classical Model becomes apparent when looking at what, exactly, is unknowable, for Williamson. It is not the content of a declarative sentence containing vague elements which is unknowable, but what is said. Utterances themselves are the bivalent truth-bearers (Williamson, 1994, p. 187). But if what has been said is in their determinate truth-conditions unknowable in principle, it is unknowable in particular for the audience. And if what is said is taken to be the only thing that is communicated, then the audience *cannot* take up the truth-conditions expressed. Thus the account fails to satisfy Classical Pragmatics (2), and hence the Classical Model.

A broad family of views on vagueness hold that the content of a vague expression—and the truth values of sentences containing the expression—is context-dependent (e.g. Graff, 2000; Kamp, 1981a; Raffman, 1994; Shapiro, 2006). In some obvious sense, many people will agree to this, as described above: whether "Richard is tall" is true depends on whether Richard is tall for a basketball player or a for horse jockey. The context can help disambiguate the different senses of "tall". But contextualist theses go further than that. For example, for Soames (1998), the speaker is at discretion to use a vague term to apply to borderline cases warrantedly. The individual might be an absolute borderline case, such that no semantic rules determine whether the vague predicate applies. The speaker has the pragmatic agency, as it were, to apply the predicate nevertheless. If accepted by the audience, this adds up to a proposal to, in this particular context, treat the individual as a definite case from here on. To put it simply, when applying the predicate to a case which might semantically be a borderline case, the speaker "shifts the border" *in this context* to somewhere else. Most work on vagueness is concerned with solutions to the sorites paradox rather than explaining communication. As far as I can surmise, contextualists are not claiming that *in this context*, the speaker uses the predicate precisely. The border is shifted, but the predicate

is not made precise. Thus, the truth *value* of the sentence might contextually be made salient, but the problem of determinate truth *conditions* has merely been moved. On the contextualist account, it is consistent with the speaker's use of the vague expression that the expression still has borderline cases. So Richard could still be a borderline case of tall if he had a different height than he actually does, even after I asserted "Richard is tall". It is understandably not the aim of the contextualist to determine a proposition expressed by the sentence uttered containing vague expressions. But this then leaves us with no recourse to repatriate contextualism into the Classical Model. What's expressed is not a proposition with determinate truth conditions.

3.3 Indeterminacy in Opaque Contexts

The second source of utterance-indeterminacy I'll consider can be called *indeterminacy in opaque contexts*. This occurs in cases where there may well be, intuitively and including contextual information, a unique proposition communicated. But where this proposition is context-dependent in ways which seem to lie beyond the scope of systematic explanation. How to determine that unique proposition can't seemingly be straightforwardly explained. Consider the following examples (taken or adapted from Bach, 1994):

(6) a. Steel is strong enough. [to be used in the space shuttle]
b. That lamp is cheap. [relative to other lamps]
c. Mutual knowledge is relevant. [to communication]
d. Biden is too old. [to run again for office]
e. Gregor was merely a bookkeeper. [as opposed to an accountant]
f. The princess is late. [for the party]
g. Tipper is ready. [to dance]
h. The king has arrived. [at the palace]
i. Al has finished. [speaking]

The list of examples shall illustrate the point that, without contextual completion, the utterances can be considered incomplete. But they are perfectly fine to assert: they are grammatically well-formed and completely adequate to utter, depending on the context. The suggestion of a completion in brackets is one of the ways in which a context could contribute. But

for any particular sentence, many such completions are possible. Consider variations on (6g):

(7) a. Tipper is ready. [to dance]
b. Tipper is ready. [to go home]
c. Tipper is ready. [to receive priestly ordination]

These completions clearly express different truth-conditions. The different completions, it seems, depend in their adequacy on the context of the utterance. Yet it is not clear how the correct completion, if any, depends *systematically* on the context. The range of examples is wide, and each example seems to require careful consideration to explain its context-dependency.

This type of indeterminacy appears to be not semantic, but rather pragmatic in character. That is, if one agrees that those sentences on their own express *some* proposition—but not the one said. This is the view of semantic minimalists like Cappelen and Lepore (2006), more on which in section 4.2.1. Bach, however, does not think so. For Bach, what is expressed by these sentences is a partial proposition ("propositional radical"). Bach claims that these examples are semantically underdeterminate: "When a sentence is in this way semantically underdeterminate, understanding its utterance requires a process of completion to produce a full proposition." Thus on this view, utterance-indeterminacy in opaque contexts is not *due* to the opaque context, since the indeterminacy (or underdeterminacy) stems from semantic properties of the sentence. Bach makes a further distinction between completion and expansion, the former semantic and the latter pragmatic, but which will be smoothed over in the present discussion. What matters is that the examples all exhibit a type of utterance-indeterminacy as described above.

The examples above are missing an unarticulated constituent, claim Perry and Blackburn (1986). Bach (1994) calls the examples in (6) cases of *constituent underdetermination*: they "require the insertion of additional conceptual material". These are to be distinguished from cases of *structural underdetermination*, which "require the articulation of structural relations among existing material". Consider the following examples:

(8) a. Willie almost robbed a bank. [he tried and nearly succeeded/he barely refrained from robbing a band/intend on robbing *something*, he robbed something else instead]

b. I love you too. [I, besides others/love, besides hate/you, besides others]

c. The number of planets [nine/whatever it is] may be even.

d. In 1996 the [now/then] president of the US will be a Republican.

e. Gyro believes that the inventor of the YO-YO [whoever that may be/Donald Duncan] is rich.

Bach (1994) struggles to pinpoint exactly what cases like (8a) are lacking. It is something like a contextually supplied scope for *almost*, supines Bach, but not a structural ambiguity on the syntactic level. Cases of structural underdetermination, just as cases of constituent underdetermination, act as sources of utterance-indeterminacy. Without the ability to explain how the unique proposition is determined, the Classical Model does not explain the communicative success: it is merely stipulated that both speaker and audience are able to infer the relevant proposition from contextual cues. This is different for perfectly determinate sentences. Here, knowledge of the linguistic meaning of the uttered sentence determines how to interpret the contextual clues (for example, to fill in indexicals). Bach develops a *sort of* answer (Bach & Harnish, 1979) of communication as a coordination game. But Bach himself admits:

> Now to describe the general character of communication is not to explain how it works. Sperber and Wilson (1986, pp. 20, 69-70) rightly point out that we 'pragmatists' have not supplied much in the way of psychological detail and about how the process of understanding utterances works. (Bach, 1994, p. 155)

Sperber and Wilson provide such a theory, but reject every aspect of the Classical Model—and reject Gricean Maxims, too. The Coarse-Grained Model developed in this thesis will not supply psychological detail, either. The task taken up here is to instead *rationally reconstruct* the inference process the audience might undertake to determine a unique proposition.

Bezuidenhout (2002) develops a contextualist account or "truth-conditional pragmatics" for types of apparent context-dependence. Bezuidenhuit adds to the list of examples. Here is one of them in full:

> A movie director is trying to determine whether everything is ready for the next scene he'll be filming. In this scene a cat must

> seem to fly up into the air and out through a skylight in the roof. So a cat has been fitted with a special harness. The cat will be suspended just above a mat, with its paws not quite touching the ground. When the director gives the signal, the cat will be whisked through the air on an invisible line and out through the skylight. The director, wanting to know whether everything is in place, asks his assistant where the cat is and she replies:

(9) It is on the mat.

> In this case what she says is true provided that the cat is suspended by wires over the mat with its paws not quite making contact with the mat.

It's clear here that the constituent expressions are neither ambiguous nor vague. The sentence is also not incomplete, nor does it lack structural determination. Still, the intuitively communicated content differs from what the lexical meaning and syntactic arrangement of the constituent words would suggest. Arguably, the expressed content is context-dependent, even for a trivial-seeming sentence like (9). Another example from Bezuidenhout: a son returns inside the house after playing with a baseball bat and a ball with his dad. He exclaims:

(10) I was playing baseball!

Now, the way he was playing resembled the rules of an actual baseball match only remotely, yet he intuitively doesn't speak falsely with (10). It seems that even the meaning of the term "baseball" can vary with context to such a significant extend to change truth value. The contextualists claim that this is the rule and not the exception. Cases of this sort are a type of indeterminacy *sui generis* and Bezuidenhout labels the phenomena "occasion-sensitivity".

If cases of occasion-sensitivity are anywhere as widespread as the contextualists would have it, the Classical Model is challenged to provide an explanation how communication still succeeds. The Coarse-Grained Model developed in this thesis provides such an explanation.

3.4 Indeterminacy due to Contextual Indifference

The third type of utterance-indeterminacy I distinguish presents an additional challenge to the Classical Model. Here, even after all relevant details about the context are fixed, there does not seem to be a single proposition either expressed by the speaker nor taken up by the audience. These are cases in which the speaker seems to be *indifferent* with regards to which of the multiple propositions compatible with the sentence uttered in the given context the audience takes up. Here's a case by Buchanan (2010):

> While preparing for their first party at their new off campus apartment, Chet and Tim go out to buy provisions for the night. After a long and heartfelt discussion, Chet convinces Tim that "sophisticated" partygoers, like the charming ladies next door, do not like to drink beer from a keg—'especially if it is domestic, bro'. To cater to the sophisticates that they hope will show up later that night, they decide to go to a local corner store to pick up several cases of imported bottled beer which they will serve from a giant ice-filled plastic bucket, decorated in a pirate motif, which is to be located in their back yard. An hour before the party is to begin, Tim asks Chet 'Are we ready to rage?' 'So bro', Chet responds, 'We are totally ready. The living room totally looks like a pirate ship. The strobe lights are up. Every beer is in the bucket. I just need to find an eye patch to wear with this pirate hat.'

Consider just (11):

(11) Every beer is in the bucket.

What did Chet communicate to Tim? Intuitively, they communicated successfully. Chet updates Tim on the progress of preparing the party. The cases of beers are placed in the ice-bucket, ready to be served at the party. It should be straightforward to determine the unique proposition communicated, then. But Buchanan notices a problem: it is far from clear which proposition Chet expresses. And this is so even if the scope of the quantified phrase "every beer" and the reference of "the bucket" is extensionally fixed. That is, even if it is clear which beers and buckets are referred to, there is

still the problem of determining the proposition expressed. This becomes obvious when looking at possible extensions of the sentence (Buchanan, 2012, p. 349):

(12) a. Every beer *we bought at the bodega* is in the bucket *in the backyard*.

b. Every beer *we will serve at the party* is in the bucket *decorated in pirate motif*.

c. Every beer *for our guests* is in the bucket *filled with ice*.

d. Every beer *at the apartment* is in the bucket *next to the hot tub*.

Arguably, this list of extensions can be continued indefinitely. Each of (12a) to (12d) (and beyond) is a candidate for what has been said by Chet. Each is perfectly compatible with the literal meaning and the contextual clues provided by the situation. It seems that Chet is *indifferent* with regards to the proposition Tim takes up. What, then, should Tim take Chet to say? What is the proposition communicated? Suppose that Chet expressed the content of (12b) by uttering (11). On the Classical Model, for communication to succeed Tim has to take up (12b). But nothing in this scenario enables Tim to make an informed decision on which proposition to take up. Either proposition (12a) to (12d) (and beyond) seems an equally good candidate. The success of the communication would be a matter of mere luck on the Classical Model. Buchanan's solution is to introduce a technical notion: a *restricted proposition type* is a rule for how to complete a proposition in a given context. Then instead of a proposition, it's a proposition type which is communicated. I'll discuss Buchanan's proposal in more detail in section 4.2.2.

The next chapter looks at and critically discusses contemporary accounts of indeterminate communication. The accounts address one or more of the types of indeterminacy just defined.

Chapter 4

Contemporary Accounts of Indeterminacy in Communication

4.1 Introduction

This chapter presents contemporary accounts which deal with indeterminacy in communication. None of the accounts both retain the Classical Model and successfully address indeterminacy in communication. They fail to do so for different reasons.

Cappelen and Lepore (2005) develop *Semantic Minimalism*, an account concerning the semantic content of utterances. What is communicated, on their account, is due to many pragmatic features, and not adequately captured in a single proposition. Thus, they reject Classical Pragmatics (1).

MacFarlane (2020b) groups together a series of accounts under the label the *Cloudy Picture* or, less prosaically, the *Multiple Proposition View*. These have in common that they reject some principle of Classical Pragmatics, like Cappelen and Lepore, typically by rejecting that a single proposition is communicated. Instead, they provide constraints on which more or less well-defined set of propositions is communicated. I discuss accounts by Buchanan (2010), Bowker (2019) and Hodgson (2018).

MacFarlane (2020c) also proposes an original account, called *Plan Expressivism*. This account takes serious the idea of pragmatic content, and rejects simple truth-conditionality (and thus Classical Content). Instead, what

is expressed with an utterance is a set of world-hyperplan pairs. Hyperplans represent plans about, among others, linguistic decisions, for example to set thresholds for vague predicates.

Abreu Zavaleta (2021) takes an idea from MacFarlane (2014) in the tradition of Barwise and Perry (1981, 1983) and rejects truth-conditionality as an assignment of truth-values to possible worlds. Instead, the content of an utterance is an assignment of truth-values to mere situations or (incomplete) states of the world. Abreu Zavaleta thus also rejects Classical Content.

The account developed by Schoubye and Stokke (2016), which I call the *Minimal Content Constraint*, does indeed try to recover a unique proposition expressed. The account also is compatible with Classical Pragmatics. The account though, I will argue, does not succeed in determining a unique proposition for mostly technical reasons, and thus fails to defend the Classical Model.

Conversational Exculpature developed by Hoek (2018) does not explain indeterminacy in communication, but rather cases of pragmatic weakening. The account Hoek develops is sufficiently similar to the Coarse-Grained Model to warrant detailed discussion, but as such does not defend the Classical Model either.

This exposition establishes that none of the accounts in the literature upholds the principles of the Classical Model. In part II, directly following this chapter, I develop the Coarse-Grained Model. The Coarse-Grained Model defends the Classical Model against the challenge from indeterminacy. All the while the Coarse-Grained Model upholds Classical Pragmatics and Classical Content.

4.2 Rejecting Classical Pragmatics

4.2.1 Semantic Minimalism

Hermann Cappelen and Ernest Lepore (2005) firmly reject that the semantic truth-conditional meaning of an utterance is context-dependent in any major way: “the most salient feature of Semantic Minimalism is that it recognizes few context sensitive expressions, and, hence, acknowledges a very limited effect of the context of utterance on the semantic content of an utterance”

(Cappelen & Lepore, 2005, p. 2).[1] The effect of the context of utterance on the semantic content of an utterance is limited to what they call the *Basic Set of Context Sensitive Expressions*. This set includes indexicals, after Kaplan (1989): personal pronouns like "I", "you", etc., demonstrative pronouns like "this", "that", etc., adverbs like "here", "now", "yesterday", etc., and adjectives like "actual", "present", etc. The set further includes so-called *contextuals*, but Cappelen and Lepore doubt their validity as truly context-dependent expressions (Cappelen & Lepore, 2005, p. 1). Among these are relational nouns like "enemy", "outsider", "native", etc., and associated adjectives like "foreign", "local", "imported", etc. They leave up to the reader to decide for themselves whether these expressions should count as genuinely context-dependent and refer to the literature (Partee, 1973; Vallée, 2003). For their central argument it seems crucial that they don't count, as will be discussed below. They introduce some intuitive tests for context-dependence.

For example, consider the so-called *Inter-Contextual Disquotational Indirect Report* test for context-dependence. If the occurrence of an expression in such a report makes the report false, then that's evidence that the expression is context-dependent. In such a case, let's say that the expression passes the test for context-dependence. An Inter-Contextual Disquotational Indirect Report has the form "A said that S". Take a decidedly context-dependent sentence. Suppose Jesus utters to Peter on Holy Friday:

(13) I will return in three days.

Peter then reports to the other disciples the day after the crucifixion: "Jesus said that I will return in three days". Peter's report is false, as "I" in Peter's report, which is a different context to the context of Jesus' utterance, now refers to Peter instead of Jesus, and "in three days" refers to a different date. Hence, this sentence passes the test for Inter-Contextual Disquotational Indirect Reports. Cappelen and Lepore argue that this test only marks expressions as context-dependent which are in the Basic Set of Context Sensitive Expressions. They offer the example (14) (Cappelen & Lepore, 2005, p. 90):

(14) John is ready.

[1] Cappellen and Lepore consistently use the term 'context-sensitivity' instead of 'context-dependence'. It does not seem, however, that they make use of MacFarlane's (2007, 2009, 2014) distinction between context-dependence and context-sensitivity. I will use the term 'context-dependence' throughout.

The reader is asked to imagine several contexts. In context $C1$, the discussion is about exam preparation. Nina utters (14). In context $C2$, people are getting dressed to go out into the rain. Nina utters (14). Then there is a third context called $5stC$, in which Cappelen and Lepore drink iced tea in a cafe on 5th street in New York City. In $5stC$, they report on Nina's utterance in $C1$ with (15a) and on Nina's utterance in $C2$ with (15b).

(15) a. Nina said that John is ready.

b. Nina said that John is ready.

Both reports are intuitively true, they claim. But they shouldn't be: context $5stC$ is entirely different from either $C1$ or $C2$. They conjecture: reports would be true for any utterance of (14) (Cappelen & Lepore, 2005, p. 91). Thus, this example fails the Inter-Contextual Disquotational Indirect Report Test for context-dependence, indicating that the Contextualists are on the wrong track when claiming that examples like (14) are context-dependent.

Charitably reading their argument, I have to conceit that I probably have not fully grasped their intent. For it seems to me that a sentence like (14) may indeed pass an Inter-Contextual Disquotational Indirect Report Test for context-dependence, especially if the test is developed in analogy to (13) above. Suppose we're in context $C2$, that is, getting ready to go out into the rain. But instead of Nina uttering (14), Felicitas utters (16), referencing the earlier context $C1$, in which Nina was talking about exam preparation.

(16) Nina said that John is ready.

But it seems that in Felicitas' context $C2$, when everybody is getting dressed to go out, uttering (16) is misleading. The audience will take Felicitas to say that Nina said that John is ready to go out into the rain, and not, as Nina intended in Context $C1$, that John is ready for the exam. Intuitions here are not so clear, then, or even point into the direction that (14) does indeed pass the Inter-Contextual Disquotational Indirect Report test. If these intuitions are right and (14) passes the test, then (14) contains context-dependent expressions. Cappelen and Lepore might defend against this by insisting that these intuitions are simply misguided: it doesn't matter at all for the semantic content of (16) whether the utterance is in any way appropriate or expected by the audience. It just matters whether the utterance is true. But this seems to put the cart before the horse, in other words, assume Semantic Minimalism.

The second test they propose is the *Collective Descriptions* test: "Suppose we know there are two contexts in which 'Yesterday John left' and 'Yesterday Bill left' are true respectively (though we don't know the times of these contexts). It doesn't follow that there is a context in which 'Yesterday John and Bill left' is true" (Cappelen & Lepore, 2005, p. 100). This gives evidence that 'yesterday' is context-dependent, which it of course indeed is. Let's say that in such a case the expression passes the Collective Descriptions test for context-dependence. For context-independent expressions, one should be able to immediately infer such a conjunctive sentence. More precisely, if 'F' is content-independent, and there are some true utterances of the form 'a is F' and 'b is F', then one can truly utter 'a and b are F' (Cappelen & Lepore, 2005, p. 100). Cappelen and Lepore claim that this is the case for the following scenario: suppose in one circumstance, astronaut Smith steps on a scale in full astronaut gear, but on an unexplored planet. The scale shows 80kg. Suppose further that in another circumstance, Jones steps on a scale in the morning. It shows 80kg, too.

Consider (17a) and (17b).

(17) a. Smith weighs 80 kg.

b. Jones weighs 80 kg.

Both (17a) in the first circumstance and (17b) in the second circumstance are "natural" utterances. And it is also true that (18), Cappelen and Lepore claim.

(18) Both Smith and Jones weigh 80 kg.

Thus, it would seem, "weighs 80 kg" fails the Collective Description test. To me, this conclusion does not follow. For if "Both Smith and Jones weigh 80kg" is true, then it is also true that they weigh the same. And by common sense, this means that when both stand on the same scale, the scale shows the same value. But nothing in the circumstances described earlier ensures that. The gravitational acceleration on the unexplored planet might be any actual value, and we know nothing about Smith's weight on earth (except, perhaps, that he likely fulfills weight requirements for astronauts). If the scale ends up showing the same value, then this is entirely accidental. Thus if we know "Both Smith and Jones weight 80kg" is true, we wouldn't be licensed to infer that they actually weigh the same. An epistemologically problematic position, in my opinion.

In any case, Cappelen and Lepore would like to limit the range of context-dependent expressions to the Basic Set of Context Sensitive Expressions. Semantic context-dependence is a purely grammatical matter. Any utterance of a sentence expresses the semantic content of the sentence, after semantic values for the expressions in the Basic Set of Context Sensitive Expressions are fixed. They explicitly summarize: "Our view is that sentences are context sensitive just in case they contain an expression from what we call the Basic Set of Context Sensitive Expressions" (Cappelen & Lepore, 2005, p. 17).

They give a list of examples like (19):

(19) Rudolf is happy.

The sentence (19) expresses the proposition *that Rudolf is happy*, and is true just in case Rudolf is happy. And any utterance of (19) expresses that same proposition, since (19) does not contain any expressions in the Basic Set of Context-Sensitive Expressions. They add, polemically: "if you find it surprising that someone would write a book defending conclusions so obvious, we have a great deal of sympathy" (Cappelen & Lepore, 2005, p. 3).

If one is following Cappelen and Lepore's argument up to this point, one might wonder where the problem actually lies that the Contextualists and this thesis stress so much. For example, recall Bezuidenhout's example about the movie set where a cat is suspended on cables hovering above a mat. The cat is about to be whizzed into the air. Answering a question about where the cat is, the production assistant replies:

(9) It is on the mat.

For Cappelen and Lepore, the situation is clear: The pronoun 'it' is context-dependent and refers to the cat contextually supplied. The utterance (9) then simply expresses the proposition *that the cat is on the mat*, case closed. How could anyone want to defend something so obvious? It is apparent that there is some tension here. How can one square both intuitions? On the one hand that obviously, "It is on the mat" expresses *that the cat is on the mat.* On the other hand that obviously, in this example the production assistant didn't mean that the cat is sitting on the mat like any ordinary cat, but instead is slightly suspended, hovering above the mat.

Cappelen and Lepore's attempt to resolve this tension by divorcing the semantic content of an utterance from what is said with that utterance. They introduce *Speech Act Pluralism*:

> No one thing is said (or asserted, or claimed, or ...) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated. What is said (asserted, claimed, etc.) depends on a wide range of facts other than the proposition semantically expressed. It depends on a potentially indefinite number of features of the context of utterance and of the context of those who report on (or think about) what was said by the utterance (Cappelen & Lepore, 2005, p. 4).

What is said, asserted, etc. by an utterance (or *speech act content*) does not have to be strongly correlated to the semantic content of the utterance. A proposition asserted can even be contradictory to the semantic content. Even all features of the context together with the semantic content are not enough to fix speech act content. Instead, even the context of those who think about the utterance (e.g. the philosophers theorizing about it) in part determines speech act content (Cappelen & Lepore, 1997).

Cappelen and Lepore further describe their account: "The semantic content of a sentence S is the content that all utterances of S share. It is the content that all utterances of S express no matter how different their contexts of utterance are" (Cappelen & Lepore, 2005, p. 142). They tout as one of the account's *virtues* to be explanatorily powerful w.r.t. communication. This might be surprising given the foregoing discussion, since it seems that divorcing semantic content from speech act content makes it harder for the semantic content to explain communication. But the reasoning behind their claim becomes clearer when one notices that for Cappelen and Lepore, the problem of explaining communication is different one from the one described in this thesis so far. Everywhere in this thesis I assumed that the problem of explaining communication can be described like the following: given a context, a speaker and an audience, what makes it the case that the speaker and the audience can communicate successfully? But for Cappelen and Lepore, what they deem much more challenging is a problem described in this way: given a speaker and a context of utterance, and an audience and a context in which the audience learns of the speaker's utterance, what makes it the case that the speaker and the audience can communicate successfully (Cappelen & Lepore, 2005, pp. 152)? In other words, what explains successful communication *between contexts*? They paraphrase an argument by Frege (1977):

> Frege in this passage is, at least in part, trying to show that if thoughts were psychological states, then it would be hard to see how individuals could communicate. The analogy is this: If communicated contents are restricted to (or, essentially tied to) specific contexts of utterance, then it is hard to envision how speakers who find themselves in different contexts can communicate, i.e., under such circumstances communication between contexts is thrown into doubt. (Cappelen & Lepore, 2005, p. 153)

In this thesis, the problem of explaining inter-contextual communicative success will be sidelined.

What is the semantic content of an utterance? It's a so-called *minimal proposition*. As one would expect, the minimal proposition can be determined by determining the lexical meaning of the constituent expressions and their semantic composition (i.e. via syntax). Additionally, ambiguous expressions should be disambiguated, vague expressions made precise, and the semantic values of context dependent expressions should be fixed. All context dependent expressions are present in the Basic Set of Context Sensitive Expressions. A question immediately arises—in which way should the expressions be disambiguated and made precise? That a unique precisification can be determined which is *the* semantic content of a sentence containing a vague expression might be doubted, as argued in Chapter 3. But that issue aside, Cappelen and Lepore claim that there is a unique minimal proposition for an utterance like (20).

(20) Tipper is ready.

The minimal proposition for (20) is *that Tipper is ready*. Truth-conditionality of *that Tipper is ready* would require that the predicate "is ready" has a fixed extension. But which one? In chapter 3 example (20) appeared as (7).

(7) a. Tipper is ready. [to dance]
b. Tipper is ready. [to go home]
c. Tipper is ready. [to receive priestly ordination]

Each of these context-dependent potential completions refer to a different extension. Schoubye and Stokke (2016) help themselves to the sentences' existential closure in cases like this:

(21) Tipper is ready *for something or other*.

The sentence (21) is very general and weak: it would be true whenever Tipper is ready for anything imaginable, and arguably be false in outlandish scenarios. The content of (21) could thus provide the minimal proposition for (20). Recanati (1989) provides additional examples:

(22) a. It will take us some time to get there.

b. I have had breakfast.

For Recanati, the minimal proposition expressed by (22a) is quite weak: "The proposition we get at this point is the truistic proposition that there is a lapse of time (of some length or other) between our departure, or some other point of reference, and our arrival at a certain place" (Recanati, 1989, p. 303). And sentence (22b) expresses "the proposition that the speaker has had breakfast at least once before the time of utterance" (Recanati, 1989, p. 303). Whereas (22a) already is overtly existentially quantified, it seems that (22b) is covertly so: there is a moment in time at which the speaker has had breakfast. That, at least, is one way of interpreting the minimal proposition according to Recanati. The point being, there remains work for Semantic Minimalists to spell out the structure of minimal propositions and how to determine them.

In addition to Semantic Minimalism, Cappelen and Lepore offer Speech Act Pluralism to explain inner-contextual communicative success, a kind of agnosticism about speech act content. By presenting a complex example of speech act content, they conclude that there just is no one right way to describe *what the utterance said*. A long list of contextual factors are to be considered, like facts about the speaker's intentions and beliefs, facts about the conversational context of utterance, other facts about the world, logical relations, etc. One of the many propositions said by an utterance is the semantic content of that utterance. There might not be any systematic theory from which all of what is expressed can be derived (Cappelen & Lepore, 2005, Chp. 13). How then, on this account, do people successfully communicate?

Cappelen and Lepore develop their theory in order to provide an alternative to Radical Contextualism. Radical Contextualism is a family of positions held together by the view that every expression is context-dependent. This means that without contextual supplementation, a sentence by itself does

not express a proposition. Notable proponents of this view are Bezuidenhout (2002), Moravcsik (1998), Recanati (2001), Searle (1978), Sperber and Wilson (1986), and Travis (1996, 2000), claim Cappelen and Lepore (Cappelen & Lepore, 2005, p. 43). However, many authors hold a more moderate contextualist view. In fact, as Cappelen and Lepore stress, many "sensible philosophers" would even diagnose a *reductio ad absurdum* if some argument leads to Radical Contextualism (Cappelen & Lepore, 2005, p. 53). Moderate Contextualists expand the set of context-dependent expressions, and develop theories how these get their semantic value in context. But they do not hold, contra the Radical Contextualists, that *every* sentence has context-dependent semantic value. Cappelen and Lepore foretell disaster for the Moderates, though. One of the central arguments in their (2005) diagnoses a slippery slope from Moderate to Radical Contextualism. They claim that any Moderate Contextualist has to set an arbitrary boundary for context-dependence of expressions. But arbitrary boundaries betray inconsistency in the evaluation of evidence: the arguments supporting Moderate Contextualism over Semantic Minimalism and Speech Act Pluralism also support Radical Contextualism over Moderate Contextualism (Cappelen & Lepore, 2005, chps. 3-5). These arguments are grouped into *incompleteness* and *context-shifting* arguments. Thus a rational philosopher who prefers a Moderate Contextualist view over a Semantic Minimalist view in virtue of incompleteness and context-shifting arguments, should also, they argue, prefer a Radical Contextualist over a Moderate Contextualist view. They establish this by first describing how these types of arguments are constructed by the Moderate Contextualist for the Semantic Minimalist. The arguments show that the Basic Set of Context Sensitive Expressions fails to include seemingly context-dependent expressions. The Moderate Contextualist then argues for extending the Basic Set of Context Sensitive Expressions. But then Cappelen and Lepore claim that for any such limited set, additional incompleteness and context-shifting arguments can be constructed. They argue that this process does not terminate until the Moderate Contextualist is forced into a Radical Contextualist position. In other words, Moderate Contextualism suffers from instability. Accordingly, they reject the validity of these arguments altogether. The examples used in incompleteness arguments Cappelen and Lepore discuss are structurally very similar to the examples for utterance-indeterminacy presented in Chapter 3. To get an idea

of how they argue against Moderate Contextualism, let's look at an example of a context-shifting argument. Cappelen and Lepore discuss quantifier domain restriction (Cappelen & Lepore, 2005, pp. 41). Consider (23):

(23) There are no French girls.

How does the Moderate Contextualist argue for the context-dependence of (23)? By describing two different scenarios. First, imagine a classroom where it is shared knowledge that the aim is to find a French girl. Then (23) is uttered. Intuitively, the speaker said something true as long as there are no French girls in the classroom. The quantifier is restricted by the context of utterance. This indicates a context-dependent semantic value for the quantifier. Second, imagine that the same speaker "is scurrying about New York City desperately seeking French girls". The speaker then utters (23). Intuitively, this time the speaker said something true as long as there are no French girls in New York City. It seems that intuitions about truth conditions of utterances of (23) vary with context. Thus the Moderate Contextualist makes the case to enter quantified expressions into the Basic Set of Context Sensitive Expressions. Cappelen and Lepore want to show: if anyone is convinced by that argument, then they should be equally convinced by similar arguments showing that *any* English sentence is context-dependent. This would entail Radical Contextualism.

They boldly state: take any arbitrary sentence! Consider (24):

(24) John went to the gym.

For any sentence, they claim, some imagination can produce scenarios in which uttering (24) yields different truth conditions. In other words, for any sentence one can construct a context-shifting argument. Cappelen and Lepore offer for (24), among others: First, suppose John likes to take walks at night, where the gym is closed. Uttering (24) intuitively says that John went somewhere close to the gym. Second, suppose John's exercise diligence is at issue, and whether he exercised yesterday. Uttering (24) seems to be truthful only if John went into the gym to exercise. The utterance is not true if John only went somewhere close to the gym. The sentence (24) does not contain any context-dependent expressions, it seems. Yet the intuitive truth value of its utterance changes, depending on the context. The Moderate Contextualist is now challenged to explain this context-dependence as

a case of a particular type of context-dependence, for example, with hidden or surprising indexicals or unarticulated constituents. Yet Cappelen and Lepore chose this particular sentence because "there are no local fixes for them" (Cappelen & Lepore, 2005, p. 50). And if there is no systematic theory of context-dependence available to the Moderate Contextualist to distinguish (24) from other arbitrary sentences, the Moderate Contextualist is at danger to slide into Radical Contextualism. Cappelen and Lepore claim the Moderate Contextualist established an arbitrary border between context-dependence and context-independence and go on to undermine its validity by turning incompleteness and context-shifting arguments against the Moderate Contextualist position.

They offer more examples; like (25):

(25) Justine is a philosopher.

In one context, Justine's character is at issue, as they are "a rather unreflective person with virtually no commitments to any position one way or another" (Cappelen & Lepore, 2005, p. 47), such that (25) might seem false when uttered. In another context, talking about professions, (25) seems true when uttered, as Justine is indeed a professional philosopher. Examples of this kind don't seem to be systematically context-dependent, at least in no obvious way. Yet the intuitions about changing truth-values with context seem stable enough. Cappelen and Lepore's claim that Moderate Contextualism inevitably slides into Radical Contextualism depends heavily on the validity of their generalization: that for any arbitrary sentence, such context shifting arguments can be constructed. But it seems quite apparent that this is an extremely strong position. Even some counterexamples of common parlance which don't display the effects of context-dependence would threaten their claim.

Some comments. It seems that already their example (25) is not as clear-cut as it might appear at first glance. In the first context, why would uttering (25) be deemed uttering something false? Uttering (25) rather seems inappropriate, eliciting a response of the sort "Yes, I know Justine is a philosopher, but she still has an noncommittal character". This seems so, I think, because (25) *implicates* something about Justine's character, because of stereotypical associations people might have about what character it takes to become a professional philosopher. What is said (literally) by

(25) has nothing to do with Justine's character. My point is that it is far from clear that what is said by *any* sentence is subject to a context-shifting argument. This can be seen at (26).

(26) Paris is the capital of France.

The sentence (26) is an "institutional" fact, a fact by human agreement (Searle, 1995), which is true iff Paris is indeed the capital of France. Both Paris and France are individual names, and being-the-capital-of is a well-defined relation. This sentence implicitly is context-dependent on the *time* of utterance, granted. Uttering the sentence in revolutionary times where the existence of a French nation-state is controversial might have an effect on its truth-value. But over and above that, there is no room for context-dependence, I maintain. If that is right, then the Moderate Contextualist is *not* forced, by their own rationale, to adopt a Radical Contextualist view. In certain respects, the Coarse-Grained Model developed in this thesis adopts a form of Moderate Contextualism. What is said by an utterance is context-dependent in ways that go beyond utilizing the Basic Set of Context Sensitive Expressions. But the context-dependence will be explained to be entirely systematic—and systematically limited.

What can be said for Cappelen and Lepore's account with respect to the Classic Model of Communication? On the one hand, it seems that their approach is quite conservative, due to their insistence that the semantic content of an utterance be truth-conditional. The semantic content of an utterance is simply the minimal content of the sentence uttered. This seems in line with the Classical Model, as it requires that contents are truth-conditional. But it is not the semantic content which is at issue in the Classical Model, but the content communicated. And, as established above, Cappelen and Lepore divorce the content communicated by an utterance from the semantic content of the uttered sentence. What Cappelen and Lepore claim explicitly is that there *is no* single proposition communicated. Instead, they offer their pragmatic or speech act pluralism as outlined above. Here, different, even innumerably many, propositions are communicated. This is of course in stark opposition to the Classical Model. What's the merit of their alternative explanation of communicative success? An important question, but which I won't attempt to answer in this thesis.

4.2.2 Multiple Proposition View

A different solution to the problems of determining the content of utterances with indeterminate expressions is to assign not a single, but multiple propositions to an utterance, as some authors have proposed. MacFarlane (2020c) subsumes strategies implementing this idea under the label *the Cloudy Picture*: multiple propositions form something like a cloud, metaphorically speaking. He argues that all of these accounts reject some aspect of Classical Pragmatics, without providing a suitable replacement. MacFarlane discusses explicitly the accounts of Braun and Sider (2007), Buchanan (2010), and von Fintel and Gillies (2011), but mentions many others developing a similarly cloudy picture (Caie, 2018; Khoo & Knobe, 2016; King, 2014; Suikkanen, 2019). There additionally are some other recent versions of this strategy (Armstrong, 2023; Bowker, 2019; Hodgson, 2018).

Let's look at the account by Buchanan (2010) first. Buchanan discusses an example already mentioned in section 3.4, repeated here. Tim and Chet are planning a party, and Tim asks Chet whether everything is set up. Chet responds with:

(11) Every beer is in the bucket.

What is communicated? Clearly, Chet does not mean that the totality of all beers in existence is in the ice-bucket in their backyard. Rather, there is some explicit constraint on the quantified domain. But which one? Buchanan argues that there is a potentially indefinite number of restrictions of the quantifier of this sentence, many of which would be a suitable candidate for what is communicated. For example:

(12) a. Every beer *we bought at the bodega* is in the bucket *in the backyard.*

b. Every beer *we will serve at the party* is in the bucket *decorated in pirate motif.*

c. Every beer *for our guests* is in the bucket *filled with ice.*

d. Every beer *at the apartment* is in the bucket *next to the hot tub.*

Each of these ways of restricting the quantifier expresses different truth-conditions. The problem for the Classical Model, then, is clear: which, if any, proposition is communicated?

Buchanan argues that what's communicated with Chet's utterance (11) is a *restricted proposition type*. This type consists of both a propositional template, a "partial structure which is determined by the lexical meanings of the uttered sentence's constituents in tandem with their syntactic arrangement" (Buchanan, 2010, pp. 357). It also consists of contextual constraints on how to complete that template into a proposition. Many different propositions are potentially tokens of such a restricted proposition type—MacFarlane calls these a "cloud" of propositions. For Buchanan, the speaker displays a certain indifference towards which of these compatible propositions the audience ends up entertaining. Communication succeeds if the audience picks up any single or multiple of these propositions.

Hodgson (2018) is objecting to Buchanan's revisionary account of communication. On his account, a speaker means multiple propositions with their utterance, and the audience has to take up one more of these propositions for them to understand the speaker. Hodgson does this by allowing intentions to be directed at a set of propositions: the speaker intends for their audience to entertain one or more of the propositions in the set. This allows the theorist to maintain most of Gricean pragmatics, Hodgson argues, while accounting for Buchanan's alleged counterexample.

Bowker (2019) argues that the speaker's intention is not to communicate a particular proposition, but instead make a determinate contribution to the common ground of the conversation. *Common ground* is used as a technical term in the sense of Stalnaker (1978, 1988, 1998) and references the shared background knowledge or assumptions of the interlocutors—in much the same way as it is used in this thesis. See section 6.2. The common ground is usually represented as a set of possibilities—all those compatible with what is assumed. The speaker meaning is then analyzed as the effect the candidate propositions have on the common ground. It might be, for example, contextually clear that "beers we bought a the bodega" and "beers for our guests" refer to the same set of beers, in other words, that the corresponding propositions are co-extensional (have the same truth value) at possibilities inside the context set. Uttered, they have the same effect on the common ground. We can then identify what the speaker meant the audience to take up from their utterance as some content p determined by its effect on the common ground.

MacFarlane charitably interprets Buchanan such that the speaker intends

the audience to recognize the proposition type intended, but is indifferent towards which token of that type is picked up and added to the common ground. But this is a problem: which proposition is, ultimately, added to the common ground? It seems that it matters for the purposes of updating the common ground that it is mutually known which proposition the audience ends up entertaining. This is the case, at least, if the propositional type restricts the propositions the audience can accept only so far as to still allow different effects on the common ground. It could, as an edge case, happen that all propositions compatible with the proposition type have an identical effect on the common ground, if accepted. For example, in the account by Bowker (2019), the cloud of propositions has a unique effect on the common ground precisely because of this. But nothing in the setup requires an identical effect, and it might well be that all propositional candidates (12a) to (12d) would have different effects on the common ground, if accepted as the content of the utterance. And then there is no unique effect on the common ground determined by the utterance. Without this effect, Classical Pragmatics (3) is not satisfied. MacFarlane identifies this as the central objection to Buchanan's account. Does the objection have merit? I think it does. To see why, consider the question whether the propositional candidates are co-extensional, given the common ground. Either they are or they are not. MacFarlane is aware of this distinction (2020c, p. 630, n. 32). Suppose they are, i.e. suppose that all candidates in (12) have the same truth values at possibilities in the context set. In this case, any proposition that the audience ends up entertaining has the same effect if added to the context set since it removes exactly those possibilities from the set the proposition is incompatible with. But these are the same for each candidate proposition, by assumption. In this case, which is also the case Bowker (2019) discusses, MacFarlane's objection doesn't apply. Instead, suppose that the candidates are not co-extensional. The beers from the bodega and the beers served to the guests might not actually be the same set. The corresponding propositional candidates then are compatible with different possibilities inside the context set. The result is that there is no unique update on the context set for the propositional candidates. MacFarlane concludes that the assertion then just cannot be viewed as an update to the common ground, and ends on a call:

> For an account like this to succeed, we would need something

> like an algorithm for determining which proposition the audience was responding to (MacFarlane, 2020c, p. 632).

It is just such an algorithm the Coarse-Grained Model aims to provide.

Summarizing, every account of the Multiple Proposition View type rejects one or more of the principles of the Classical Model.

4.3 Rejecting Classical Content

4.3.1 Plan Expressivism

MacFarlane's proposal is to keep Classical Pragmatics intact. Instead, Classical Content is to be rejected: in place of truth conditions, the content is supposed to vary additionally with "one or more non-factual parameters" (MacFarlane, 2020c, pp. 639). These additional parameters are hyperplans. Hyperplans are "[. . .] fully determinate contingency plans covering every possible circumstance and resolving all indecision" (MacFarlane, 2020b, pp. 647). The position is dubbed *plan expressivism* and combines work by Gibbard (2003) and Barker (2002). The content of an assertion is the set of all world-hyperplan pairs such that the asserted sentence is true at the world under the decisions made about the use of language according to the hyperplan. The account is expressivist since the content does not determine truth conditions, but has pragmatic parameters as well.

The content expressed by (5), "Richard is tall", is determined by plans to set a threshold for the gradable adjective *tall*:

$$\{(w, h) \mid \text{the height of Richard in } w \geq \text{the threshold for } \textit{tall} \text{ determined by } h\}.$$

In addition, the common ground consists of world-hyperplan pairs instead of a simple context set.

What is the effect of asserting (5)? MacFarlane (2020b, pp. 649) distinguishes between three cases:

> (C1) If it is already common ground that Richard is 189 cm tall, then an assertion of (5) will be tantamount to a proposal to plan to set the threshold for tall no higher than 189 cm.

(C2) If it is already common ground that we plan to set the threshold for tall at exactly 189 cm, then an assertion of (5) will be tantamount to a proposal to add the factual proposition that Richard is 189 cm tall to the common ground.

(C3) But if the common ground is agnostic about Richard's height (taking it to be between 185 and 195 cm) and undecided about the threshold for tall (not excluding options between 180 and 200 cm), then the update proposed by an assertion of (5) will be neither a plan nor a factual proposition. Rather, it will be a conditional commitment that ties together plans and factual beliefs.

Uttering (5) in case C1 is not uttering a factual proposition, but a proposal for a shared plan: it is a proposal to discard from the context set of world-hyperplan pairs all those with a hyperplan determining the threshold for "tall" to be higher than 189 cm. The effect of an assertion (and acceptance) of (5) is to remove from the common ground all world-hyperplan pairs which are not part of its hypercontent. The context set before the assertion only includes world-hyperplan pairs consisting of worlds at which Richard is 189 cm tall, but there is no consensus about where to set the threshold for "tall". Why is that not a factual statement, on this account? Consider the set Δ_w of "just the worlds" of the context set Δ: $\Delta_w = \{w \mid (w, h) \in \Delta\}$. Asserting (5) in this case does not have an effect on Δ_w. It does have an effect on the set of "just the hyperplans" $\Delta_h = \{h \mid (w, h) \in \Delta\}$, though. We discard from the context set those world-hyperplan pairs consisting of hyperplans assigning a threshold for "tall" higher than 189 cm.

Compare Fig. 4.1. Checkmarks (✓) indicate world-hyperplan pairs which are part of the hypercontent of (5) and thus remain in the context set post assertion. Crosses (✗) indicate those world-hyperplan pairs which are to be discarded from the context set post assertion. The subscript of a hyperplan denotes the threshold set by it, the subscript at a world denotes Richard's height at that world.

If C2 is the case, uttering (5) is making a factual statement, and all world-hyperplan pairs consisting of worlds at which Richard is not at least 189 cm tall are discarded from the context set. This is a factual proposition, since only the set of "just the worlds" Δ_w of the context set is affected. We

	$w_{1.89}$
⋮	
$h_{1.87}$	✓
$h_{1.88}$	✓
$h_{1.89}$	✓
$h_{1.90}$	✗
$h_{1.91}$	✗
⋮	

Figure 4.1: Effect of asserting (5) on the context set in case C1.

discard from the context set of world-hyperplan pairs those consisting of worlds at which Richard is shorter than 189 cm. As a small aside, it seems that if it is common ground that anyone taller than 189 cm is to count as tall, then asserting (5) proposes to add that Richard is *at least* 189 cm tall to the common ground. Compare Fig. 4.2.

	...	$w_{1.87}$	$w_{1.88}$	$w_{1.89}$	$w_{1.90}$	$w_{1.91}$	...
$h_{1.89}$		✗	✗	✓	✓	✓	

Figure 4.2: Effect of asserting (5) on the context set in case C2.

In case C3, the effect on the common ground is neither factual nor a plan to change the threshold. Rather, it is something called a *conditional commitment*. It is the commitment for one of two things. Either that if later in this context we agree on a threshold for tall, this has a factual effect on the context set, since all world-hyperplan pairs containing a hyperplan setting a different threshold will be discarded. Thus the set of "just the worlds" Δ_w will be reduced. Or if later in this context we agree on Richard's height, this will influence our plans about a threshold. If we agree on Richard being 189 cm tall, for example, Δ_h is reduced by those classes of hyperplans which set the threshold for "tall" higher than 189 cm. Compare Fig. 4.3.

How does this account explain successful communication? In one sense,

	...	$w_{1.87}$	$w_{1.88}$	$w_{1.89}$	$w_{1.90}$	$w_{1.91}$	...
⋮							
$h_{1.87}$		✓	✓	✓	✓	✓	
$h_{1.88}$		✗	✓	✓	✓	✓	
$h_{1.89}$		✗	✗	✓	✓	✓	
$h_{1.90}$		✗	✗	✗	✓	✓	
$h_{1.91}$		✗	✗	✗	✗	✓	
⋮							

Figure 4.3: Effect of asserting (5) on the context set in case C3.

the problem becomes trivially easy. The hypercontent associated with an assertion like (5) is context-independent. At least once the "reference class or domain of the degree function has been fixed" (MacFarlane, 2020b, pp. 649). Determining the hypercontent a speaker asserted is thus just a matter of knowing the content of the sentence asserted. The task for the audience to recognize the hypercontent becomes similarly easy. They are able to recognize the content merely in virtue of their linguistic competence.

However, I'd like to levy three points of criticism for the account which justify subsequently developing an alternative model.

First and most obviously, MacFarlane's account gives up Classical Content. This entails to give up truth-conditionality of utterances, and instead to adopt a form of expressivism. Forsaking truth conditions seems to address explaining vague communication, although I will argue below that it doesn't do so convincingly. Whether the result is worth the sacrifice every philosopher has to decide w.r.t. their own goals and purposes. However, giving up truth-conditionality even for perfectly precise utterances which do no suffer from a lack of explanation in the Classical Model seems to be quite costly. To explain many communicative exchanges one does not need to respect idiosyncratic plans to use language and fine tune to our conversation partner. It is enough to point to the fact that both participants are competent speakers of English to explain their ability to share content. Involving

hyperplans in the explanation of every occurrence of communication is to use a sledgehammer to crack a nut. Merely factual statements do not benefit in their explanation from sophisticated planning coordination. It is just very plausible to associate mere truth conditions with merely factual statements.

MacFarlane's account might have the means to address this issue simply by stating that the content of merely factual statements, since they do not rule out any hyperplans, could be seen "as if" we only take into account "just the worlds", i.e. Δ_w. In other words, since they do not propose any change in the set of "just the hyperplans", i.e. Δ_h, the expressivist part of their hypercontent might be ignored in the explanation of its content. Such an account so far, however, seems to be lacking. The Coarse-Grained Model developed in this paper provides a way to use standard truth conditions as contents while still taking into account multiple permissible interpretations of indeterminate expressions.

The second point of criticism concerns case C3 above. The account seems to predict that no factual proposition was communicated. Here is a simple example in which this prediction goes against intuition.

Two officers tasked with assigning new recruits to their bataillons are discussing where to assign Richard. To be eligible for the honor guard, the recruit needs to be taller than 180 cm, as both know. They discuss:

(27) a. Question. Does recruit Richard meet the height requirements to be assigned to the honor guard?

b. Response. Richard is tall.

It seems that, even if slightly unnatural, the response gives a factual answer to the question posed: yes, Richard is indeed tall enough to serve in the honor guard. In this scenario, the common ground pre conversation does neither include any constraints on Richard's height nor constraints on a threshold for using "tall". The officers do not have a consensus with regards to a threshold for "tall". It seems that the question introduces some sort of contextual constraint on the use of language, but it's not obvious why it would introduce a constraint on the use of "tall". By C3, MacFarlane's account however judges that no factual proposition has been communicated with (27b), at best some conditional commitments. This seems to conflict with the intuition that indeed, (27b) is a valid and informative answer to the question. The Coarse-Grained Model to be developed in this paper below

assigns truth conditions to (27b), and thus accords with this intuition.

The third point of criticism concerns coordinating on the common ground: introducing an additional degree of freedom into the representation of the common ground by introducing hyperplans in addition to possible worlds is to introduce additional ways for the presumed background knowledge of speaker and audience to misalign.

In the Stalnakerian tradition, the common ground represents the presuppositions or what the participants in the conversation treat as the presumed common knowledge. See section 6.2.

For each participant, it's those presuppositions which they hold and which they deem their conversation partner to hold, too. Naturally, these presumptions might differ. If we represent the common ground each participant presumes individually, each participant has their "own" context set—the set of those worlds compatible with what they presume to be common ground. If the context sets of the participants do not match up (i.e. are not identical), the context becomes a "defective context", as Stalnaker calls it. The mismatch can lead to a failure in communication. For an explanation of successful communication the mismatch is a challenge also because we cannot easily identify a unique effect on *the* context set, if the participant's context sets don't align. But Stalnaker immediately provides an idea as to how communication can still succeed: "a context is *close enough* to being nondefective if the divergences do not affect the issues that actually arise in the course of the conversation" (Stalnaker, 1978). That is, defective contexts do not have to make explaining successful communication impossible, if the goals or purposes of the conversation are not in jeopardy.

Now, if our theory introduces a lot more complexity into what context sets consists in, there is added potential for these context sets to not align, in other words, for the context to be defective. This is the case for MacFarlane's account. Consider an assertion like (5), "Richard is tall". Let's assume that speaker and audience agree on all factual matters. But the speaker takes it to be common ground that no one in their right mind would use the term "tall" to apply to people shorter than 170 cm. This entails in MacFarlane's account that they take the context set to not include world-hyperplan pairs consisting of hyperplans which set the threshold for tall to be lower than 170 cm. The audience is less lenient in their use of "tall" and thinks any threshold lower than 180 cm to be mistaken. Now the speaker and audience

both take (5) to have a different effect on the context set, simply because their context sets differ (this holds whether we look at pre or post assertion context sets). The defective context thus prevents determining a unique effect on the context set. According to Classical Pragmatics, this is not a case of successful communication, or at least it's not explained by it.

What we need is a way to explain how, even if the context sets misalign, the communication can still succeed by satisfying the purpose of the conversation. The account developed in this thesis will model the purpose of the conversation explicitly and thus yield an explanation for successful communication in spite of, in a sense, misaligned context sets by satisfying the purpose of the conversation.

4.3.2 Complex Propositions

Another account which rejects Classical Content to deal with utterance-indeterminacy is provided by Abreu Zavaleta (2021). I'd like to point out that Abreu Zavaleta's paper is exceptionally well written and is inordinately precise and rigorous. The account might well have a lot of merit and seems particularly interesting. Although it cannot fulfill the aim of this thesis—to defend the Classical Model—as it rejects one of its central principles. I'll give here a brief overview of the account.

Abreu Zavaleta is interested in a problem for a *propositional view of literal communication*—a view very similar to the Classical Model of Communication. This problem is the one described by Buchanan (2010) and discussed in section 4.2.2 above. Chet and Tim are planning a party, and Chet utters (11):

(11) Every beer is in the bucket.

Which proposition is Chet expressing, such that Tim can entertain it? Abreu Zavaleta makes use of MacFarlane's distinction between indexical or *context-dependent* and *context-sensitive* expressions. According to MacFarlane (2009):

- An expression is *indexical* iff its content at a context depends on features of the context.
- An expression is *context-sensitive* iff its extension at a context depends on features of the context.

Thus, for merely context-sensitive expressions, the context only plays a *circumstance-determining* role. Standard indexicals like "I" are context-dependent. Abreu Zavaleta provides the example "weighs 60 kg" for merely context-sensitive expressions. This expression has the same content, independent of context of utterance. But the context may provide means to determine the intended circumstance of evaluation. Such a circumstance would provide the way in which "weighs 60 kg" should be evaluated.

Abreu Zavaleta then introduces the notion of a *proposition**: a function from circumstances of evaluation to truth-values. The corresponding *propositional* view of literal communication* Abreu Zavaleta describes thus:

Expression*

> Every literal assertoric utterance of a declarative sentence expresses a proposition*.

Understanding*

> Communication through a literal assertoric utterance of a declarative sentence is successful only if:
>
> (i) The audience recognizes what proposition* the utterance expresses.
>
> (ii) At least one of the circumstances of evaluation which the speaker assumed the audience would take her utterance to be a description of is such that the audience takes the proposition* the utterance expresses to be a description of it.

Applied to the problem, the account states that (11) expresses a proposition*, namely that every beer is in the bucket. Chet then assumes that Tim takes some certain circumstances of evaluation to be described by his utterance. For communication to succeed, Tim has to actually take Chet's utterance to be a description of at least one of these circumstances. What are these circumstances? Recall from the discussion in section 4.2.2 that there are many different candidate completions for (11) which Tim could take Chet to potentially say, for example:

(12) a. Every beer *we bought at the bodega* is in the bucket.

b. Every beer *we will serve at the party* is in the bucket.

These completions correspond to these circumstances:

(28) a. A circumstance of evaluation which determines that the beers relevant to the truth of Chet's utterance are the ones he and Tim bought at the bodega.

b. A circumstance of evaluation which determines that the beers relevant to the truth of Chet's utterance are the ones which they will serve at the party.

The difference to standard truth-conditional propositions is that circumstances of evaluations are not required to be possible worlds. Instead, circumstances of evaluations can also be *situations* (Barwise & Perry, 1981, 1983; Fine, 2012, 2016). Situations provide, intuitively, partial descriptions of possible worlds. For example, a situation may determine Anna to be 1.7m tall, but nothing about John's weight. In other words, situations may have *incomplete information*. A situation might not determine for a particular object whether it falls into a particular extension. Mathematically, situations are viewed here as points, partially ordered by parthood relation. The reason I titled this section *complex propositions* is that situations are more general (in a certain sense) than possible worlds. Accordingly, situation-based propositions are more general than standard truth-conditional propositions.

How are the circumstances described in (28) cashed out in terms of situations? Let's briefly describe a semantics, following Abreu Zavaleta. I simplify where appropriate. A model consists of a set of situations S, a domain D, a subset D_s of D for every $s \in S$, and an interpretation function I, such that

- if P is a (one-place) predicate, then $I(P)$ is a function from states $s \in S$ to ordered pairs $\langle P_s^+, P_s^- \rangle$, where P^+ and P^- are subsets of D_s.

Intuitively, P_s^+ is the extension of P at a situation s, and P_s^- is its antiextension. Then:

- $P(t)$ is true in s relative to an assignment function g iff $g(t) \in P_s^+$, and
- $P(t)$ is false in s relative to an assignment function g iff $g(t) \in P_s^-$, and
- indeterminate otherwise.

The domain of a quantifier in a situationist semantics is restricted by the situation. Then "every beer" might refer to all the beers in the domain of a situation. More formally,

- $\forall x : \varphi(x).\psi(x)$ is true at a situation s relative to an assignment g iff for all u in D_s such that $\varphi(x)$ is true in s relative to $g_{u/x}$, $\psi(x)$ is true in s relative to $g_{u/x}$.

- $\forall x : \varphi(x).\psi(x)$ is false at a situation s relative to an assignment g iff there is some u in D_s such that $\varphi(x)$ is true in s relative to $g_{u/x}$, and $\psi(x)$ is false in s relative to $g_{u/x}$.

Then the assignment function g is abstracted away in the usual manner. Utterance (11)—"Every beer is in the bucket"—is paraphrased as

$$\forall x : \text{Beer}(x).\text{InBucket}(x).$$

The semantic content of (11) shall then be defined as

$$\langle \{s \mid (11) \text{ is true at } s\}, \{s \mid (11) \text{ is false at } s\} \rangle.$$

This constitutes the proposition* expressed by (11). From my point of view, this implementation of the proposition* view is fairly standard and straightforward and thus convincing.

But how does the proposition* view hold up? I'd like to make two points. First, as mentioned, the proposition* view denies the Classical Model, as it rejects that propositions are truth-conditional in the sense outlined. Thus, the view cannot solve the task this thesis set out to perform, i.e. defend the Classical Model against the challenge from indeterminacy. So far, so good. But second, is the view convincing as a model of communication in its own right? I don't want to make too strong a statement on this point. Let me just voice this worry. One crucial aspect of the view are the circumstances which the speaker takes the audience to be described by their utterance. Another are the circumstance which the audience actually takes the speaker's utterance to be a description of. Communication succeeds only if these match up such that the latter is one of the former. But it seems that nothing in the elaborate proposition* view actually explains what makes the speaker take a certain circumstance to be described by the utterance. Also, nothing in the view explains what makes the audience take a circumstance to be described by the utterance. It's not linguistic competence; competence enables speaker and audience to know the semantic content. It somehow is a feature of the context. It's not the shared background knowledge; background knowledge

enables speaker and audience to judge an utterance true or false. It seems that the rationality of this taking-to-be-a-description-of is left entirely unexplained, which is—for me—not too convincing as a model of communication. The Coarse-Grained Model developed in this thesis has something to say on these points. It also defends the Classical Model of Communication.

4.4 Unsuccessful or Not Applicable

4.4.1 Minimal Content Constraint

This section discusses an account developed by Schoubye and Stokke (2016), which I call the *Minimal Content Constraint*. It does indeed try to recover a unique proposition expressed. The account also is compatible with Classical Pragmatics. The account though, I argue below, does not succeed in determining a unique proposition for mostly technical reasons, and thus fails to defend the Classical Model.

Schoubye and Stokke (2016) develop their framework by making use of the notions *minimal content* and *question under discussion*. The minimal content of a sentence is something like its compositional meaning. See also section 4.2.1 on Semantic Minimalism. A question under discussion represents an antecedently accepted topic of discussion, the answer to which is not unlike the Gricean goal or purpose of the discussion. Propositions are truth conditional and represented as sets of possibilities.

They define *what is said* by a sentence S relative to a context c and a question q_c as a proposition φ that satisfies the following conditions:

- (i) φ entails the minimal content of S in c.[2]
- (ii) φ answers q_c.
- (iii) φ is the maximal proposition satisfying (i) and (ii).

Propositions are sets of possible worlds, and a proposition ϕ entails another ψ iff $\phi \subseteq \psi$. The minimal content of a sentence S in a context c is given by first contextually saturating indexicals and potentially other plausibly context-sensitive expressions like gradable adjectives. The minimal

[2]In Schoubye and Stokke's proposal, the candidate proposition might also be entailed by the minimal content. In the following example and the other cases discussed this part of the condition for what is said does not make a difference, for the most part. I'll indicate explicitly when it does.

content is then just its compositional meaning, if needed supplied by existential generalization. See the case below for an example. Schoubye and Stokke talk about "maximal" and "weakest" propositions, which I interpret as follows: A proposition ϕ is maximal iff it is the maximal element (of a reference class) iff there is no other proposition (of that reference class) ψ s.t. $\phi \subsetneq \psi$. This condition does not imply uniqueness.

Schoubye and Stokke are reconstructing the notion of an answer following Groenendijk and Stokhof (1984a), Hamblin (1973), and Roberts (2012). Given a conversational scenario and a question modeled in terms of question alternatives, *what*, *who*, ..., in short, *wh*-questions are analyzed as a set of propositions which are themselves answers to polar sub-questions of the *wh*-question. For example, "who is awake?" is modeled by a set of propositions such that each proposition answers, for a particular individual, if they are awake:

$$q\text{–alt} = \{\{w \mid \text{Mary is awake at } w\}, \{w \mid \text{Kelly is awake at } w\}, \dots\}.$$

This set of propositions induce a partition. Let's denote this partition with Π_q. For Schoubye and Stokke, a proposition answers such a question if it is a non-empty non-total union of cells of Π_q. I will explore the notion of an answer further in chapter 5.

Illustrate with an example:

(29) a. Question. The space shuttle must be able to carry 35 tons of cargo, endure extreme temperatures, and be capable of withstanding severe cyclonic dust storms. So, what material for the shuttle is sufficiently strong?

b. Response. Steel is strong enough.

Intuitively, (29b) says that steel is strong enough to carry 35 tons of cargo, endure extreme temperatures, and be capable of withstanding severe cyclonic dust storms. Let's assume that the only contextually salient materials we consider are steel, aluminium, and iron. Denote with S the proposition that steel is strong enough to carry 35 tons of cargo and so on. It's negation is $\bar{S}$. Likewise with aluminium: A denotes the proposition that *aluminium* is strong enough to carry 35 tons of cargo and so on. And likewise with *iron*: I denotes the proposition that *iron* is strong enough to carry 35 tons

of cargo and so on. The contextual question under discussion is described by "which material is sufficiently strong to carry 35 tons of cargo [and so on]?" This question partitions the totality of possibilities Ω into cells as depicted in Fig. 4.4. The grid illustrates the partition induced by the question under discussion. The shaded grid cells mark what is said by (29b), intuitively: the union of all S-cells. Cells of the partition are sets of possibilities at which the conjunction of the propositions or its negation depicted hold. So for example, the $\bar{S}AI$-cell is that set of possibilities at which steel is *not* strong enough to carry 35 tons of cargo and so on (but might or might not be strong enough for other things), while aluminium and iron are indeed strong enough to carry 35 tons of cargo and so on.

Ω

SAI	$SA\bar{I}$	$S\bar{A}I$	$\bar{S}AI$
$S\bar{A}\bar{I}$	$\bar{S}A\bar{I}$	$\bar{S}\bar{A}I$	$\bar{S}\bar{A}\bar{I}$

Figure 4.4: What is said by (29b) in context, according to Schoubye and Stokke.

Schoubye and Stokke argue that their account predicts what (29b) intuitively says. They state the minimal content of (29b) as that steel exists and it's strong enough for at least something. To answer the question under discussion, the proposition we are looking for must be a union of partition cells. To be a candidate for what is said, it must also entail the minimal content. This rules out $\bar{S}$ cells, they argue, since there will be possibilities in $\bar{S}$-cells at which steel doesn't exist and as such these cells can't entail the minimal content, since the minimal content requires that steel exists. Their predicted candidate for what is said is represented by the shaded area in Fig. 4.4, according with the intuition described above. In the next section, I will argue that this account has apparent counterexamples, dependent on how we choose to model the common ground of the discourse participants.

In their reconstruction of Roberts' account (2006), Schoubye and Stokke also discuss the picture Stalnaker develops (1978, 1988, 1998) on which "information that is mutually taken for granted by the discourse participants

is referred to as the *common ground*" (Schoubye & Stokke, 2016, pp. 766). See also section 6.2, in which I talk more about Stalnaker's notion. Given a totality of possibilities Ω, the common ground is modeled as that set of possibilities $\Delta \subseteq \Omega$ compatible with common background beliefs. Schoubye and Stokke make use of this notion at a few spots in their paper, strongly indicating that the notion is part of their account.

What does it mean do include Stalnaker's notion of common ground? Schoubye and Stokke are not explicit about it. My interpretation is that instead of applying the conditions for what is said above to the totality of possibilities, we apply it to the reduced context set only. The conditions have to be adjusted accordingly.

We then redefine *what is said* by a sentence S relative to a context c and a question q_c as a proposition φ that satisfies the following conditions:

- (o) $\varphi \subseteq \Delta$.
- (i*) φ entails the intersection of Δ with the minimal content of S in c.
- (ii*) φ answers q'_c
- (iii*) φ is the maximal proposition satisfying (o), (i*), and (ii*).

The question q'_c is intuitively just the question q_c reduced to those possibilities inside the context set. More precisely, there corresponds a unique partition on the context set to the partition induced by the question under discussion: if Π_q is the partition of Ω induced by the question under discussion, then $\Pi_q^\Delta = \{x \cap \Delta \,|\, x \in \Pi_q\} \setminus \{\emptyset\}$ is the unique induced partition on the context set, compare Fig. 4.5.[3]

If the above reconstruction is right about how to incorporate the Stalnakerian notion of common ground in the account by Schoubye and Stokke, then the resulting account cannot deal with case (29). For surely in any plausible context in which the structural properties of steel, aluminium, and iron are discussed it is part of the interlocutor's shared background knowledge that steel *exists* and is strong enough *for something*. This something can be anything, for example, strong enough to not substantially decay after one second. One might imagine a philosophy context in which the existence

[3]Whether we take q–alt and first induce its partition and then intersect with Δ or first intersect with Δ and then induce the partition does not matter, we arrive at the same result.

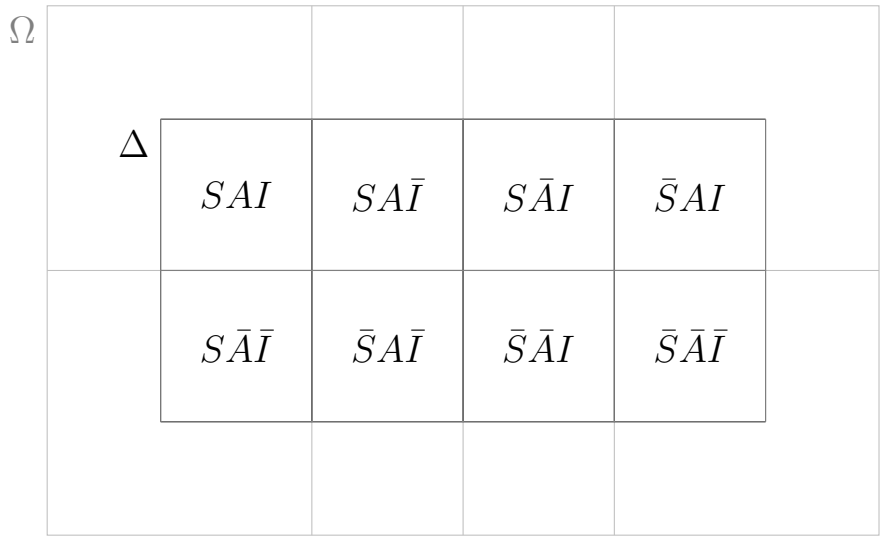

Figure 4.5: The partition Π_q for (29a) on Δ.

of steel is up for debate, but that is an extreme case. If instead it is common ground that steel exists and is strong enough for something, then the context set Δ entails the minimal content. Thus, any union of cells of the partition in Fig. (4.5) entails the minimal content, too. Then we cannot use this entailment criterion to decide between candidate propositions. There is no unique proposition determined, since there are several candidate propositions which provide an answer to the question under discussion but all of them are maximal in the sense given above.

For example, consider the proposition that steel or aluminium or iron is strong enough to carry 35 tons of cargo an so on, intersected with the context set. Let's call this proposition χ_Δ, compare Fig. 4.6.

It satisfies the condition for what is said. To see this, note that it is the union of all cells in Fig. (4.6) except for $\bar{S}\bar{A}\bar{I}$. Clearly, χ_Δ satisfies (o). Since any union of these cells entails the minimal content, so does χ_Δ, and hence χ_Δ satisfies (i*). The partition depicted in Fig. (4.6) is induced by the question variant described in (ii) from the question under discussion in the example. Since χ_Δ is a non-empty non-total union of cells of this partition, χ_Δ answers the question variant and thus satisfies (ii*). Is χ_Δ also the maximal proposition satisfying (o)-(ii*)? We won't find a proposition that satisfies (o)-(ii*) of which χ_Δ is a proper subset. The only candidate which is also a union of partition cells is Δ, but this union is total w.r.t. Δ and

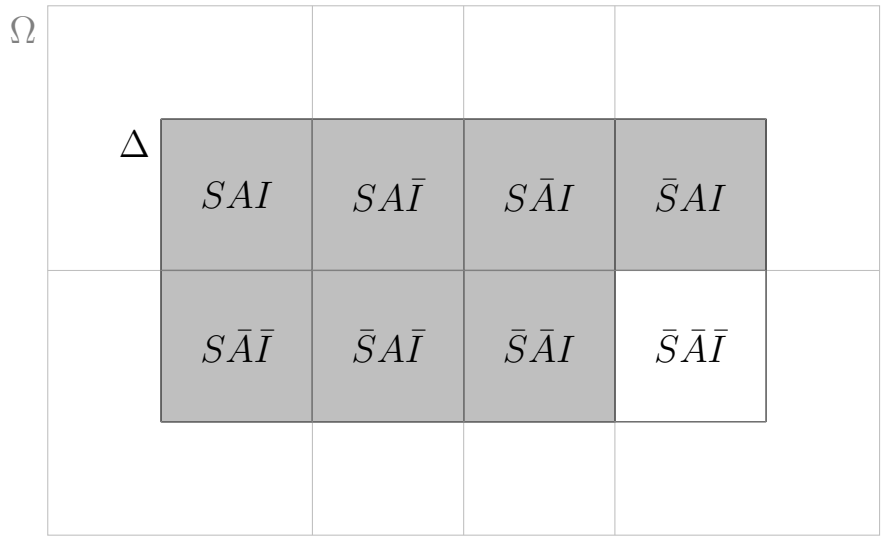

Figure 4.6: The proposition χ_Δ that steel or aluminium or iron is strong enough to carry 35 tons of cargo and so on, intersected with the context set Δ.

thus doesn't qualify as an answer, violating (ii*). Hence χ_Δ is the maximal such proposition, satisfying (iii*). But clearly, this proposition is not what we want our account to predict as what is said by (29b). What's more, any other union of all Π_q^Δ-cells safe one cell satisfies the conditions in the same manner. Thus the account does not predict a single proposition said, if the account is spelled out like described here.

A natural reaction to this argument is to deny that the reconstruction (o) - (iii*) just given is correct, either completely or in parts. There might be some combination or adaption of these conditions which is closer to the intended account. For example, one could deny that (o) is actually required, i.e. that we allow candidate propositions which aren't a subset of the context set (and adopt (iii*) accordingly). In a way, this is quite sensical, since the proposition expressed by a sentence might distinguish between possibilities outside of the context set, too, and it is only the conjunction of common knowledge and what is said which is what we are modeling with the requirement (o). But this change would not solve the problem above. Consider a proposition like χ_Δ, which entails the minimal content and decides the modified question under discussion reduced to the context set. Now take the proposition χ, the proposition that steel or aluminium or iron is strong enough to carry 35 tons of cargo and so on, but *not* intersected with the

context set. Compare Fig. 4.7.

Figure 4.7: The proposition χ that steel or aluminium or iron is strong enough to carry 35 tons of cargo and so on, *not* intersected with the context set Δ.

The Proposition χ satisfies the conditions for what is said just as well. And similarly do other unions of Π_q cells safe one. Thus, it's not possible to find a unique weakest proposition among them for the same reason as above.

What about the other conditions? Condition (i*) is set-theoretically equivalent to (i),[4] so there is no wiggle room here. Condition (ii*) seems to be crucial to the idea of incorporating common ground as a context set. If we get rid of (ii*), but keep (o), many plausible candidate propositions fail to answer the question under discussion, as the propositions are limited to the context set, but the question partition is not, as on the present account an answer needs to be a union of partition cells. The most straightforward option, then, is to disregard the notion of common ground as presented here entirely. However, this route runs into a different problem, which I will outline now.

The problem can be constructed using the case from a well-cited paper by Buchanan (2010). This example is discussed multiple times for different reasons in this thesis, see sections 3.4 and 4.2.2.

An hour before the party is to begin, Tim asks Chet 'Are we

[4]Let's denote the minimal content of S in c with M. It's easily seen that $\varphi \cap \Delta \subseteq M$ iff $\varphi \cap \Delta \subseteq M \cap \Delta$.

> ready to rage?' 'So bro,' Chet responds, 'We are totally ready. The living room totally looks like a pirate ship. The strobe lights are up. Every beer is in the bucket. I just need to find an eye patch to wear with this pirate hat.'

Buchanan asks us to consider just:

(30) Every beer is in the bucket.

What is said by (30)? Let's charitably assume that a question under discussion along the lines of "Are enough beers chilled?" (i.e. in the ice bucket) is perfectly definite and partitions the possibilities clearly into those where there are enough beers chilled and where there aren't, and further assume that "the bucket" is completely non-problematic, too. The minimal content seems to be that every beer *for some non-empty domain* is in the bucket, i.e. that any beer is in the bucket at all. The 'yes'-alternative of the question under discussion entails the minimal content and the account thus predicts as what is said that there are enough beers in the bucket. Although Buchanan argues that there is no unique single proposition said, we sideline this discussion and acknowledge that *this* prediction doesn't constitute a counterexample to its validity. Consider, however, the situation if Chet is slightly more specific:

(31) Every beer *we bought at the bodega* is in the bucket.

The minimal content of (31) seems to be simply that every beer they bought at the bodega is in the bucket. One can quickly see that neither answer to the question under discussion entails the minimal content and not *vice versa*, either: consider possibilities at which (i) there might be enough beers in the bucket but some of the ones bought in the bodega are missing, (ii) all the beers bought at the bodega might not be enough, (iii) there are not enough beers in the bucket yet neither are all those bought at the bodega, or (iv) all the beers bought at the bodega are in the bucket and those are enough beers. Given this configuration, then, the account cannot determine a proposition said, because neither answer to the question under discussion entails the minimal content. A natural thought to mitigate this is to take into account contextual limitations to the possibilities under considerations, for example via a context set. Suppose it is common knowledge how many beers

they bought at the bodega and that those beers are enough as pertaining to the question under discussion, and we only consider possibilities in the context set to find a propositional candidate that answers the question under discussion, similarly to (o) and (ii*) above. The 'yes'-answer to the question under discussion, that there are enough beers in the bucket, still does not entail the minimal content: there are possibilities at which there are enough beers in the bucket but not all the beers they bought at the bodega are in the bucket. However, the 'yes'-answer *is entailed* by the minimal content intersected with the context set: if every beer they bought at the bodega is in the bucket and these are enough beers, as is common ground, then there are enough beers in the bucket (obviously). Changing the account into this direction then could potentially successfully address the problem. However, any way of making the account explicit in such a manner runs into the problem for modeling the common ground discussed in detail just above.

Thus it does not make a difference to the bottom line whether we model the common ground explicitly and restrict the propositions we consider as candidates for what is said to include only those possibilities in the context set: there are plausible cases where no unique proposition can be identified regardless of the choice.

But even if solutions to the issues around the common ground are granted, there is an independent line of problems for the account centering around how to make precise the notion of answering a question. Schoubye and Stokke explicitly take up the definition of a partial answer from Roberts (2012). Yet their working definition is slightly different; they use what I call a coarse weak answer. To make the differences between the notions clear I first briefly give some plausible definitions for answers. What an answer to a question could be will be discussed in much more detail in chapter 5.

I argue that using a coarse weak answer is overly restrictive and the way of mitigating the resulting issues not quite plausible. I finally show that using partial answers does not work at all for Schoubye and Stokke's account.

- A proposition φ is a *complete answer* (CA) to a question q if φ entails a q-partition cell, i.e. $\varphi \subseteq \pi$ for some $\pi \in \Pi_q$.

- A proposition φ is a *partial answer* (PA) to a question q if φ is a positive partial answer or a negative partial answer. That is, if $\varphi \subseteq a$ or $\varphi \subseteq a^c$ for some $a \in q$–alt.

- A proposition φ is a *weak answer* (WA) to a question q if φ entails the complement of at least one cell of the q-partition, i.e. if $\varphi \subseteq \pi^c$ for some $\pi \in \Pi_q$.

For each of these types of answers there is a coarse version that does not carry any additional information: a (complete, partial, weak) answer φ is *coarse* if it is a union of cells of the partition induced by the q-alternatives, or more precisely if there are $a_1, \ldots, a_n \in \Pi_q$ s.t. $\varphi = a_1 \cup \ldots \cup a_n$.

I'll look at two notions in particular: coarse weak answer and partial answer. Consider first the notion of a coarse weak answer, which requires to rule out at least one cell of the question-partition. Schoubye and Stokke employ this notion as a working definition, as is evident from the text: "[...] relative to any partition p on a question q, the union of any non-empty proper subset of p is a partial answer to q"(Schoubye & Stokke, 2016, pp. 768). Part of what's required to determine what is said is to "[...] take the union of any non-empty proper subset of the partition on q_c [...]" (Schoubye & Stokke, 2016, pp. 774). They (I think, mistakenly) label the union of any non-empty proper subset with *partial answer*. In the present discussion, this constitutes a coarse weak answer. I claim that this notion is implausibly restrictive.

To continue the example from the exposition above, let's assume the question under discussion is "Who is awake?" and the utterance made is:

(4b) Kelly has a big day tomorrow.

The response (4b) doesn't give us logical reasons to prefer one answer over the other. Instead, it only gives us non-decisive evidence in favor of one way of answering, depending on the context. Kelly might already be asleep to be in top shape tomorrow, but it might also turn out that she is still preparing and hence awake. Yet on the present account, there is no proposition which entails the minimal content of (4b) and also answers the question under discussion. For the question alternatives are plausibly like described above, and the minimal content is just that Kelly has a big day tomorrow in some way. But the minimal content and its complement are both compatible with any cell of the partition. We'd then be forced to assume that nothing is said with (4b). Schoubye and Stokke (2016, pp. 786) consider a related case:

(32) a. Is Ellen ready for the interview?

b. She's been preparing for weeks.

For (32b), too, the account fails to determine a single proposition. This is because there are possibilities in which Ellen is ready for the interview but she hasn't been preparing for weeks in any fashion, such that the 'yes'-alternative does not entail the minimal content. One way to approach this issue is to see it as part of a strategy to answer (32a) by answering a *different* question:

(33) How long has Ellen been preparing for the interview?

How does answering (33) contribute to answering (32a)? Schoubye and Stokke (2016, pp. 786) propose that this is the task for the common ground: "[. . .] in contexts of this kind, it will typically be common ground that having prepared for a long time is sufficient to count as ready for an interview, while having failed to prepare is sufficient for counting as not ready [. . .]." How can the present account of question under discussions and partitions achieve this result? The only way I can see is this: given the common ground (i.e. a contextually restricted set of possibilities), answering (33) answers (32a). For example, given the common ground, it might be the case that at all possibilities where Ellen has prepared for, say, five days or more are possibilities where they are ready for the interview. If this is so, however, then already (32b) answers the question (32a), and is also a candidate for what is said. For given the common ground, the 'yes'-alternative implies the minimal content of (32b). That is, if we make the strong assumption that the common ground provides for the difference by eliminating all possibilities that would hinder the conversational contribution, we do not need the intermediate step of the secondary question (33). But even then, of course, then we'd run into the objection laid out above.

Synthesizing an additional secondary question that is not obviously a subquestion or in a direct relation to the primary question under discussion seems to be quite *ad hoc*. For it does not seem explanatorily helpful in analyzing *what is said* by an utterance to introduce an implicit question under discussion which is then answered by the utterance, but which is introduced just so the utterance can answer it. That is, in understanding what is said by the question we introduce, we already have to understand what is said by the utterance we introduce the question for. Grindrod and Borg (2019) even diagnoses a vicious circularity: the question under discussion is deter-

mined by the intuitive truth conditions of the utterance, but the intuitive truth conditions of the utterance are explained with the help of the question under discussion.

Does the approach help with analyzing (4)? It does not seem to capture at all the ampliative (i.e. non-deductive) character of the answer that we want to describe. It seems completely reasonable that (4) is used to merely supply evidence towards the question under discussion without actually answering it. This evidential relationship is not a natural explanandum for the present account. Supplying a secondary question under discussion only provides additional explanation if it explains why answering the secondary question is evidence for an answer to the primary. But it doesn't. We don't, arguably, want the common ground to already determine that if Kelly has a big day tomorrow she must be awake, but rather let it be an open question in the conversation. One might simply concede that the ampliative nature of this answer is outside the scope of the account, and I am quite sympathetic to this stance.

What about the second option to define an answer, the partial answer? After all, Schoubye and Stokke cite the exact definition in the paper, would that work?

Reconsider the space shuttle example (29). The minimal content of (29b) is the set of possibilities in which steel exists and is strong enough for something. The minimal content is thus not decisive enough to answer the question under discussion on any notion of answer defined above. Propositions that are entailed by the minimal content consequently are not answers, either. To provide an answer to the question under discussion, the proposition has to, if it is to be a partial answer, entail the S-cells (and thereby ruling out the $\bar{S}$-cells, or vice versa), or similarly for A and I-cells. Let's take, for example, the maximal subset of $\bar{S}$-cells which still entails the minimal content, i.e. all possibilities in which steel is strong enough for something, but not to carry 35 tons of cargo and so on, and call it $\varphi_{\bar{S}}$. This constitutes a partial answer to the question under discussion. Additionally, it must be the maximal such proposition. Additionally, we can generate propositions $\varphi_S, \varphi_A, \varphi_{\bar{A}}, \varphi_I$, and $\varphi_{\bar{I}}$. There is no unique maximal proposition.

Siu (2020) presents a very similar case, but the account here overgenerates. Consider this case, in which the 'yes'-alternative implies the minimal content, if we make the innocuous assumption that being a robot implies

having robot parts.

(34) a. Is Tipper a robot?

b. Tipper has some robot parts in him.

The response (34b) might just be used to answer the question in the affirmative, implicating that since Tipper has robot parts, Tipper is a robot. This would for example be the case if the common ground excluded other possibilities, and the response is part of a strategy to answer the question, like discussed above. On another plausible reading though, (34b) merely provides ampliative evidence that having these robot parts make Tipper a robot. Yet the account predicts that what is said is that Tipper is a robot since that is the maximal proposition answering the question and entailing the minimal content.

Siu poses another objection, but which I don't think is valid:

(35) a. Is everyone who is ready to kill a dragon strong enough to do so?

b. Whoever is ready is strong enough.

They diagnose what (35b) says as that everyone who is ready to kill a dragon is strong enough to do so, or

(36) $(\forall x)$ (If x is ready to kill a dragon then x is strong enough to kill a dragon).

They state its minimal content as that whoever is strong enough for something is ready for something, or

(37) $(\forall x)$ (If $(\exists y)$ (x is ready for y) then $(\exists z)$ (x is strong enough for z)).

Clearly, neither entails the other, and their verdict is that Schoubye and Stokke's account undergenerates. However, a correct existential generalization of (36) seems instead to be

(38) $(\exists y)(\exists z)$ $((\forall x)$ (If x is ready for y then x is strong enough for z)).

Since (38) is implied by (36), the objection doesn't succeed.

It has become obvious, then, that either choice on whether to include the common ground in the account or on which notion of answer to employ leads to substantial objections.

4.4.2 Conversational Exculpature

An anonymous reviewer made me aware of the account by Hoek (2018), commenting that it seems to be quite similar to the one developed by me. The thesis might surely benefit, then, to detail Hoek's well-developed and interesting account and comment on similarities and dissimilarities to the account developed here. I find that Hoek's account is a different spin on a not completely unrelated philosophical method, yet addressing a different phenomenon. Insofar as the accounts are applicable to the same cases, I find that the present account highlights different properties of a conversation, making it intuitively more explanatory powerful.

More concretely, Conversational Exculpature explains cases of pragmatic weakening. These are cases in which a contextual presupposition appears to be subtracted from the literal content of an utterance to result in the apparent utterance-meaning. Hoek's focus is to determine the contextual presupposition and question under discussion, given the literal content and utterance meaning. Thus Conversational Exculpature as developed by Hoek does not address indeterminacy of literal content or utterance meaning. If made to address such cases, the account has to stipulate shared knowledge of contextual presuppositions between speaker and audience in order for communication to succeed. But the account is developed well enough that there is plausible wiggle room, such that speaker and audience do not have to presuppose the exact same proposition. This begs the question, however, why cases of utterance-indeterminacy are plausibly explicated by appeal to a contextual presupposition. The Coarse-Grained Model developed in this thesis provides an alternative, and makes the—I take it—more plausible claim that utterance-indeterminacy is better explicated by modeling the different ways in which the utterance can be made precise, deemed by both speaker and audience to be permissible. I'll detail how Conversational Exculpature works below and give my assessment.

Hoek is particularly interested in explaining cases of pragmatic weakening. Pragmatic weakening is a contrapositive to Grice's pragmatic strengthening: when implicating something other than what is literally said, the speaker says *something more*. In cases of pragmatic weakening, the speaker seems to say *less* than what they literally say or what they literally say would imply. Consider (39):

(39) Ellen wore the same type of hat as Sherlock Holmes.

Hoek argues: uttering (39) semantically presupposes that Sherlock Holmes exists. But the speaker of (39) doesn't intuitively make the commitment that Holmes exists. Rather, they want to say that Ellen wears a deerstalker, the peculiar type of hat Sherlock Holmes wears in the stories. Hoek hypothesizes a hidden proposition which is "subtracted" from the literal content of (39). In this example, this proposition could be that Sherlock Holmes wears a deerstalker. This proposition is a *contextual presupposition* (Simons, 2005, 2013). These are "presuppositions that connect an utterance's literal content to the question under discussion." Speaker and audience must assume this proposition such that the utterance becomes relevant. And this is the central idea: Hoek wants to subtract the contextual presupposition from the literal content such that the result is *wholly about* the topic of discussion (in the sense of Yablo, 2014). And this indeed works well, in my opinion. To see this, we have to take a closer at the formalism Hoek develops. First, some definitions. These follow closely the exposition by Hoek (2018); I added clarifications when necessary.

- A *proposition* is a subset of the set of all possible worlds, i.e. $p \subseteq \Omega$.

Much like you'd expect, I'd take worlds to be primitives here. The negation of a proposition is its complement w.r.t. Ω, i.e $\neg p := \Omega \setminus p$

- A *partial proposition* q is an ordered pair of disjoint sets of worlds $\langle t, f \rangle$ such that q is true at worlds in t and false at worlds in f.

- The *restriction of a proposition* p *to a proposition* q is the partial proposition $\langle p \cap q, \neg p \cap q \rangle$ and is written as $p|q$.

This restriction merely limits the proposition p to only q-worlds, resulting in a partial proposition, as it does not provide binary truth values for worlds at which q does not hold.

- A *subject matter* is a partition on the set of all worlds Ω, i.e. every $w \in \Omega$ is in exactly one partition cell.

- A proposition p is *about* a subject matter S iff p *is* a union of cells of S.

- A proposition p has *no bearing on* S just in case p intersects every S-cell.

This ostensibly means that a proposition p has no bearing on S iff for each cell $c \in S$, $p \cap c \neq \emptyset$. Hoek provides a helpful illustration which I somewhat faithfully tried to recreate, cf. Fig. (4.8). Solid color indicates the true part of a proposition, and diagonal lines indicate the false part of a proposition.

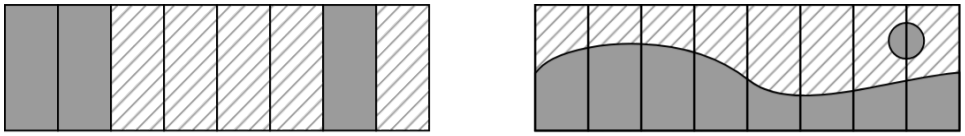

Figure 4.8: A proposition about S (left) and a proposition with no bearing on S.

- A partial proposition r is *about* a subject matter S iff r is a restriction of some proposition about S.

This seems to mean that a partial proposition r is about a subject matter S iff there is a proposition p which is about S (i.e., p is a union of S-cells) and there is a proposition q such that $r = p|q$ (i.e. $r = \langle p \cap q, \neg p \cap q \rangle$).

Hoek adds the following, equivalent definition:

- A partial proposition $\langle t, f \rangle$ is *not about* S iff there is an S-cell containing both t- and f-worlds.

This also means that a partial proposition $r = \langle t, f \rangle$ is about S iff there is *no* S-cell containing both t- and f-worlds iff there is a proposition p which is a union of S-cells and there is a proposition q such that $r = p|q$. This observation might not be entirely obvious but does indeed follow from the definitions.

Proof. See appendix.

- Finally, the *completion* of a proposition $\langle t, f \rangle$ by the subject matter S, written $S(\langle t, f \rangle)$, is well-defined just in case $\langle t, f \rangle$ is about S. Then $S(\langle t, f \rangle)$ is the (partial) proposition such that:

$$S(\langle t, f \rangle) := \langle \{\bigcup c \in S | c \cap t \neq \emptyset\}, \{\bigcup c \in S | c \cap f \neq \emptyset\} \rangle.$$

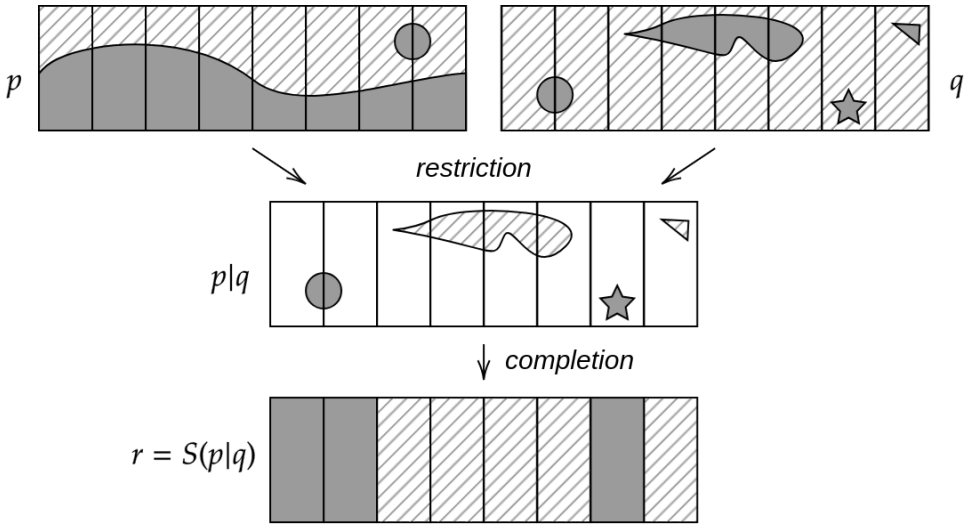

Figure 4.9: Restriction and completion of proposition p.

Hoek provides another very helpful illustration which I again recreated, see Fig. (4.9).

The definition of completion does share some features, as the reviewer pointed out, to my definition of *coarsening*, which will be presented in detail in chapter 5.

Hoek's idea runs along these lines: to determine what is said in the case of pragmatic weakening by, e.g. (39), now referred to as proposition p, one has to find a proposition q to "subtract". The proposition q has to be chosen in such a clever way that the restriction $r = p|q$ is about S, represented by the question under discussion as a partition. This means, as defined above, that the restriction has to be of a proposition about S. But p is not about S, since otherwise we wouldn't need to apply this account at all. So which proposition which is about S is restricted by q such that one can determine what is said? Here Hoek's formalism simply takes the resulting completion, i.e. $r = S(p|q)$. Since then $r|q = p|q$ (as Hoek proves), q is indeed a restriction of a proposition about S.

It works reasonably well in the case of Ellen's hat, although at one point Hoek takes the proposition subtracted to be *the Sherlock Holmes Myth*, which is of course not a proposition. But this, to be fair, is only shorthand for some relevant propositions about Sherlock Holmes. But which ones? It seems that speaker and audience somehow need to agree about this to communicate

successfully. Or if they don't need to, then Hoek needs to explain how communication succeeds anyway. These issues will be addressed in the account developed in this thesis.

The task as I interpret it is to explain what is expressed by an utterance in these peculiar cases of pragmatic weakening. One could argue that these are cases of indeterminacy in opaque contexts as defined in chapter 3. However, I will not make such an argument in this thesis. To determine the proposition expressed, Hoek requires some contextual presupposition which can be subtracted from what has literally been said by that utterance, in order to make the resulting proposition maximally relevant to the purpose of the conversation. The purpose of the conversation is given by the question under discussion, which is assumed to be part of the context. That much is fine, however, the problem remains to determine the proposition subtracted from the literal content. The formal account described above does not help with this issue, as in order to to determine q, it requires to already have determined what is said (r), such that $r|q = p|q$.

Indeed, Hoek gives a slightly different formulation when applying the account to examples: "We start out with a problematic statement that, in context, has a reading r distinct from its apparent semantic content p. The aim is to explain the discrepancy using the present account of conversational exculpature" (Hoek, 2018, sec. IV). In other words, given the literal content of an utterance and what is intuitively said by it, the account describes a method to determine the proposition subtracted from the literal content to arrive at what is intuitively said. So the task might be interpreted differently: given a scenario and an utterance, assume the literal content as well as what is said. Then, determine a reasonable question under discussion and determine the proposition subtracted. Let's assume that determining the question under discussion is unproblematic. Hoek provides additional constraints on q, the proposition subtracted:

- $p|q = r|q$,

where p is the literal content and r what is expressed, and

- q has no bearing on S,

where S is the question under discussion.

It can easily be seen that for any given p, r and S, there are many different propositions q satisfying these conditions and yielding the same $r = S(p|q)$.

Interestingly, Hoek acknowledges this (towards the end of section 3), but takes it as an advantage of the account, as it adds to the account's robustness in the following sense: it's permissible for the audience to take the speaker to use any of several different propositions q, as long as the resulting completion of $S(p|q)$ does not vary. This, then, ensures *successful communication* across different propositions q. This idea is also central to the account developed in this thesis. But it requires the account to revert back to the aim I sketched above, to explain what is expressed by an utterance in cases of apparent pragmatic weakening. Choosing q is ad-hoc, and the audience is at liberty to choose q. The proposition just has to satisfy the set of conditions described above, and result in an answer about S.

In this context, Hoek makes an interesting claim: this robustness does not only hold w.r.t. variations of q, but also w.r.t. variations of S. That is, it's possible to vary the subject matter S without affecting the resulting $S(p|q)$, while keeping p and q fixed. The argument is that since the result r is about S, r will be about any *bigger* subject matter, too. A subject matter being bigger than another here simply means that it is more *fine-grained*. A subject matter T is more *fine-grained* than another S iff every cell of S is the union of cells of T. The audience might just as well take T to be the subject matter instead of S without changing the result. In other words, if T is more fine-grained than S, and q has no bearing on T, then $T(p|q) = S(p|q)$. This is the case since q is chosen such that $p|q = \langle t, f \rangle$ is about S. That is, no S-cells contains both t- and f-worlds. But then also no T-cells contain both t- and f-worlds. Consequently, $t \subseteq T(p|q)$ for $t \in T$ iff $s \subseteq S(p|q)$ for $t \subseteq s$, $s \in S$. In other words, the worlds in the cell of the finer subject matter are included the resulting proposition exactly when the worlds in the "parent" cell of the more coarse-grained subject matter are.

How does the account presented by Hoek compare to the account developed in this thesis? Hoek's account is quite successful, in my opinion, at explaining certain cases of pragmatic weakening, for example the case of Ellen and Sherlock Holmes described above. It's plausibility hinges on the existence of contextual presuppositions which are subtracted from the literal content of the utterance. If one accepts contextual presuppositions as a reasonable tool in the explanatory tool-set, Hoek's account works well, and even allows for certain variations in the subject matter and the contextual proposition without varying the resulting proposition expressed. I see this

as one of the account's major strength. The account developed in this thesis constructs a unique proposition from a variety of contextual factors, similarly allowing for variations in the exact assumptions between speaker and audience while still allowing for communication via a shared proposition. But its range of application is different (and broader). For example, Hoek's account is not applicable to the problem of vague utterances, if vagueness is understood as allowing for many different propositions *semantically* expressed: the account requires definite semantic content. Hoek's paper addresses a case of seemingly imprecise communication. I'd argue that in the case described, neither the account developed in this thesis nor Hoek's account is needed to explain the linguistic phenomenon. Consider (40).

(40) Rob is six foot one.

Hoek argues that the literal content of (40) is that Rob is precisely six foot one. But intuitively, the speaker only commits themselves to something weaker, namely that Rob is somewhere close to, but not necessarily precisely, six foot one. Hoek offers this analysis:

- p: Rob is six foot one.
- q: Rob is some integer number of inches tall.
- S: What is Rob's height to the nearest inch?
- r: Rob is six foot one to the nearest inch (that is, he's between 6'0.5" and 6'1.5").

The question under discussion now makes reference to the standard of precision required. This is plausible, since the standards for numerical ascriptions like these vary from context to context, and representing this within the question under discussion is reasonable. If the literal content of (40) is that Rob is precisely six foot one, than this proposition already entails a unique answer to the question under discussion S. There is no need to employ contextual presuppositions to arrive at the part about the subject matter of what's been said. Simple entailment suffices (and, by the way, for any more fine-grained subject matter, too).

This concludes the discussion of contemporary accounts, and also concludes part I of this thesis. The next parts presents the Coarse-Grained Model of communication. To be able to do so, the next chapter starts by making careful distinctions between different types of questions and answers.

Part II

The Coarse-Grained Model of Communication

Chapter 5

Questions and Answers

5.1 Introduction

What is the meaning of a question? And given its meaning, what are answers to that question? In this section, I'll give an analysis of what an answer might be.

First, be clear what "question" is supposed to refer to. Groenendijk and Stokhof (1997) distinguish three uses: there's the question sentence, a linguistic object, usually with a question mark at the end (Q1). This is called an *interrogative sentence*. Then there is the uttered question, a type of speech act (Q2). Finally there is the meaning of a question, that which can be answered (Q3). So, "what's the meaning of a question?" can be rephrased as "what do we mean (Q3) when we ask (Q2) a question (Q1)?"

Clearly, all three uses are closely related. I'll distinguish between them when necessary, and in other cases the context should make obvious which of the three uses is at issue when using the term "question".

When looking at examples of questions, a few distinctions can readily be made. Cross and Roelofsen (2024) distinguish between *whether*-questions. Intuitively, we answer these with yes or no:

- "Are there more than eight planets in our solar system?"
- "Did you like the play?"

Which-questions allow for more fine-grained responses. For example,

- "Which team won the Euro 2024?"

- "Who is coming to the party?"
- "What's the airspeed velocity of an unladen swallow"?

For the first question, what's asked for is a particular team. The second question can be answered in different ways: naming some of the people who come, naming a group, naming someone who doesn't come, etc. The third question requires some particular unit and magnitude, seemingly, or some range of speed.[1] A further type of question asks *why*:

- "Why did you do that?"
- "Why does the earth not spiral into the sun?"

The first question asks for a *reason* for your action, while the second question asks for a *causal explanation.*

I'll look at some accounts of questions—and their answers—given so far, and subsequently develop my own Coarse-Grained account of answers.

5.2 In the Literature

One of the first authors to give a systematic account of questions and answers was Hamblin (1973). Groenendijk and Stokhof (1997) diagnose principles of Hamblinian question meaning. An answer to a question is a sentence, or statement, and that to know the meaning of a question is to know what counts as an answer to that question. They also add a third principle, the realization of which will be described when discussing partition semantics: the possible answers to a question form an exhaustive set of mutually exclusive possibilities. Hamblin writes:

> So let us turn to semantics. Here we must make some departure, since although we are inclined to class 'who' and 'what' with proper names we can not by any stretch regard them as denoting individuals. But there is a simple alternative: they can be regarded as denoting sets of individuals, namely the set of humans and the set of non-humans respectively. This does not mean, of course, that the formula 'who walks' asserts that the set of human

[1] Of course, the *proper* answer to this question is asking in return whether it's an African or European swallow.

> individuals walks: we must modify other stipulations in sympathy. We shall need to regard 'who walks' as itself denoting a set, namely, the set whose members are the propositions denoted by 'Mary walks', 'John walks',... and so on for all individuals. Pragmatically speaking a question sets up a choice-situation between a set of propositions, namely, those propositions that count as answers to it. (Hamblin, 1973, p. 48)

> Semantically, an answer to a question on a given reading is any statement whose denotation-set on a suitable reading is contained in that of the question. (Hamblin, 1973, p. 52)

For Hamblin, and to some extent Karttunen (1977), the meaning of a *whether-* or *which*-question is the set of all answers to the question. Answers are propositions, so the denotation of a question at a world w is a set of propositions. Propositions are sets of possible worlds. The denotation of a question at a world is a function from worlds to sets of propositions. Questions might have different denotations at different worlds.[2] But quite pointedly, Cross and Roelofsen (2024) lay the finger on the weak point of the account:

> A fundamental problem with these accounts is that they do not specify in more detail what "possible answers" are. (Cross & Roelofsen, 2024)

What method of determining possible answers does Hamblin provide? Hamblin develops his account along the following lines:

(41) Who is coming to the party?

An answer slots a name into the place of *who*. Any single individual of the right semantic type can thus be made to contribute to an answer. For example:

(42) a. Robert is coming to the party.
b. Johnny is coming to the party.

[2]Since different denotations at different worlds won't be relevant to the points made in this section, I'll gloss over the potential differences between meanings and denotations and interchangeably speak of question meaning and the denotation of a question.

c. Mary is coming to the party.

For simplicity, assume that Robert, Johnny and Mary exhaust the contextually salient individuals. Let's denote the propositions expressed by answers (42a), (42b), and (42c) with R, J, and M, respectively. To be clear, R stands for the proposition that Robert comes to the party, J for the proposition that Johnny comes to the party, and M for the proposition that Mary comes to the party. Compare Fig 5.1, which shows these propositions as sets of worlds.

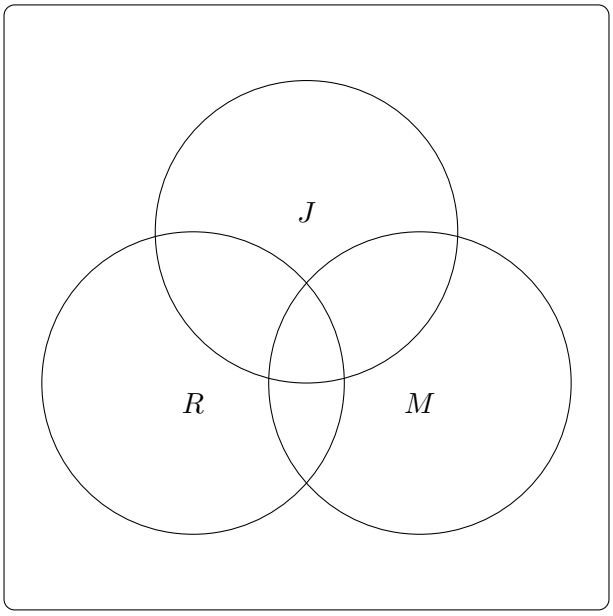

Figure 5.1: Answers to a question in Hamblin's account.

Arguably, R, J, and M do not cover the range of answers to (41), even if it is assumed that Robert, Johnny and Mary exhaust the contextually salient individuals. Intuitively, also (43a), (43b), and (43c) express answers to (41).

(43) a. Only Johnny and Mary are coming to the party.

b. Either Johnny or Mary is coming to the party.

c. Mary is coming to the party, which, by the way, I think will be really cool.

What's going on here?

Utterance (43a) answers the question completely: for each individual, this answer decides whether they are coming to the party. According to

(43a), Robert is not coming to the party, but Johnny and Mary are. But on Hamblin's account, only R, J, and M are answers.

Utterance (43b) does not decide for either individual whether they come to the party. It merely submits the information that at least one of either Johnny or Mary will show up. Does this constitute an answer to the question (41)? Not obviously so, it seems. On the one hand, there is a definite gain of information: the audience learned that it won't happen that neither Johnny or Mary shows up. Depending on the circumstances, this might be important information, for whatever practical reason. In such a circumstance, (43b) does indeed seem to answer (41). On the other hand, since (43b) doesn't decide for either individual whether they come to the party, the speaker here failed to answer the question properly. One could argue that the intuition for judging (43b) an answer to (41) rests on mistaking question (41) for another question, something like "Does either Johnny and Mary come to the party?". But such an argument, since it is surprising, would have to be substantially supported by evidential observations, i.e evidence of people's tendencies to switch out questions with different questions when evaluating thought experiments. In lieu of such deciding evidence, I think it is fair to assume that utterances like (43b) *could* be counted as answers to (41), at least under *some* circumstance. It would therefore benefit a philosophical account to accommodate and explain this type of answer. The account developed in this thesis does just that.

The third utterance (43c) decides for Mary whether they are coming to the party, but yields additional information (namely, that it might be a cool party). I think this *clearly* is an answer to the question "Who is coming to the party?", but interestingly, this case if often overlooked in accounts of answers in the philosophical discussion. Again, the account developed in this thesis will be able to explain this type of answer as well.

A second influential account is given by Groenendijk and Stokhof (1984a): for them, the denotation of a question at a world is its true exhaustive answer; a proposition. The meaning correspondingly is a function from worlds to propositions. All true exhaustive answers form a partition on logical space. Their account accordingly carries the label *partition semantics*. Consider again the example (41) from above, repeated here:

(41) Who is coming to the party?

Suppose that at a world w, Robert is not coming to the party, but Johnny and Mary are. Then the true exhaustive answer to (41) is (43a).

(43a) Only Johnny and Mary are coming to the party.

It is exhaustive, since it specifies for all contextually salient individuals whether they are coming to the party. There is exactly one true exhaustive answer per world. The answers to a question thus partition logical space into equivalence classes. For worlds w, v: $w \sim_q v$ iff q has the same true exhaustive answer at w and v. What are true exhaustive answers? For whether-questions, a proposition denoted by either 'yes' or 'no'. For which-questions, an exhaustive answer is a complete specification: for every contextually salient individual (of the right semantic type), are they coming to the party?

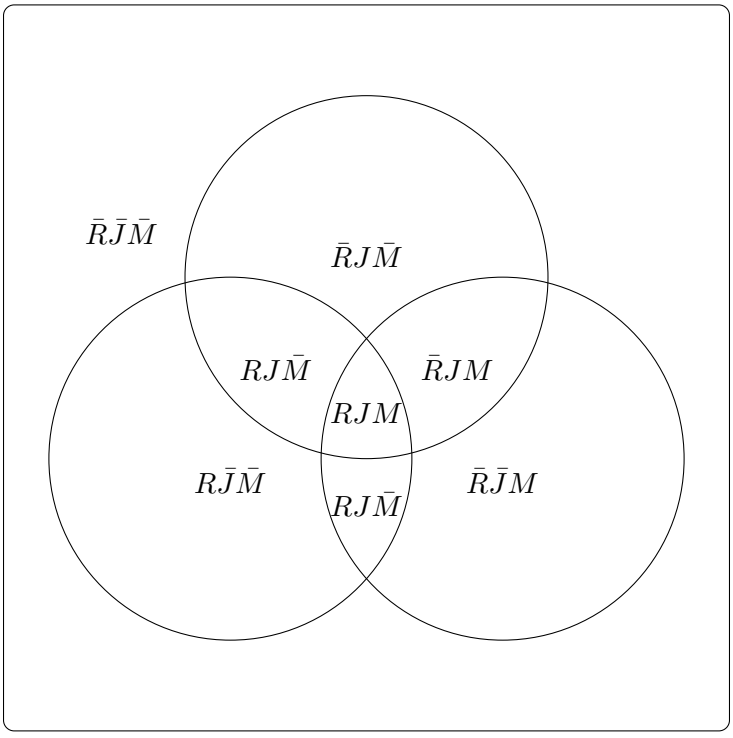

Figure 5.2: Answers to a question partition logical space.

Fig. 5.2 depicts all true exhaustive answers to (41). The notation RJM is shorthand for $R \wedge J \wedge M$ and stands for the proposition that Robert, Johnny, and Mary are coming to the party. A bar above the letter indicates negation, so that, for example, $\bar{R}J\bar{M}$ is shorthand for $\neg R \wedge J \wedge \neg M$ and stands for the

proposition that out of Robert, Johnny, and Mary, only Johnny is coming to the party. Each of the true exhaustive answers constitute a cell of the partition induced by the question (41). To emphasize the fact that it is a partition we're talking about, and to make the following exposition slightly clearer, the partition can equivalently illustrated like in Fig. 5.3.

RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$
$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$

Figure 5.3: Alternative but equivalent illustration of how answers to a question partition logical space.

Limiting the notion of answer to only true exhaustive answers has advantages, most notably in further formal work (e.g. Groenendijk, 1999). But it's also limiting the range of intuitive cases the theory can explain. While example (43a)—"Only Johnny and Mary are coming to the party"—can now be classified as an answer, other options are not (e.g. (43b)—"Johnny or Mary is coming to the party").

This is different in the next account: inquisitive semantics provides the basis for a whole system of logics (Ciardelli, 2022; Ciardelli et al., 2018). Inquisitive semantics builds upon a similar idea again as Hamblin's account. Ciardelli et al. (2013) and Groenendijk and Roelofsen (2009), too, take the denotation of a question to be a set of propositions. These propositions are the answers to the question. But contra Hamblin, they also require two things:

- that the answers cover the whole logical space, and
- that the propositions are downward closed.

The first condition just means that the union of all answers covers Ω, and in fact is identical to *Omega*. This can be achieved by, for example, also counting the negation $\neg p$ of an answer p as an answer. The second condition means that whenever for a question q and a proposition p it holds that $p \in q$,

then also for every proposition $p' \subseteq p$ it holds that $p' \in q$. In other words, all subsets of answers to a question are also answers to that question.

Ciardelli et al. (2018) make this idea more precise in the following way:

- An *information state* s is a set of possible worlds.
- An *issue* is a non-empty, downward closed set of information states.
- An information state s *resolves* an issue I just in case $s \in I$.

Issues contain information states. The states can be overlapping or distinct. Examples of issues are given in Fig. 5.4. This illustration is provided by Ciardelli et al. (2018), which I recreated for the purposes of this thesis. Here, circles represent worlds $w_1 - w_4$. Rounded rectangles represent information states. The downward-closed-property is not included in the illustration.

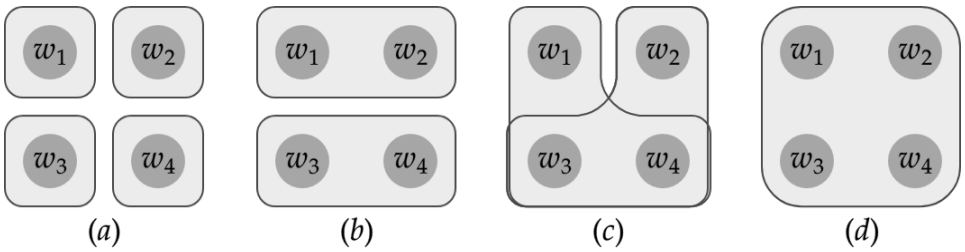

Figure 5.4: Inquisitive issues.

To relate this terminology to the terminology used so far in this chapter: issues in the inquisitive sense correspond, very roughly, to questions. Information states correspond to propositions. An information state resolving an issue corresponds to a proposition answering a question. The correspondence is merely rough because inquisitive semantics defines further conceptual refinements over and above the ones presented here. For instance, propositions are redefined as a more general type of structure than merely a set of possible worlds. A proposition has an informative content (a set of possible worlds) and raises a determinate issue as well. For the purposes of this discussion, these refinements aren't needed and are left out for simplicity.

A lot of care went into analyzing types of interrogative sentences and matching the right issue to the interrogative sentence. Consider the examples given by Ciardelli et al. (2013) and their corresponding issues:

(44) Johnny is coming to the party.

(45) Is Johnny coming to the party?

(46) Is Johnny↑ coming to the party, or Mary↑?

(47) If Johnny is coming to the party, will Mary come too?

(48) Who is coming to the party?

To model these (interrogative) sentences, assume four worlds, one where both Johnny and Mary are coming to the party (11), one where only Johnny is coming to the party (10), one where only Maria is coming to the party (01), and one where neither of them is coming (00). Compare Fig. 5.5. This illustration is again provided by Ciardelli et al. (2018), which I recreated for the purposes of this thesis. Worlds are represented by circles labeled 11,10,01, and 00. Rounded rectangles represent information states. The downward-closed-property is, again, not included in the illustration, but should be kept in mind.

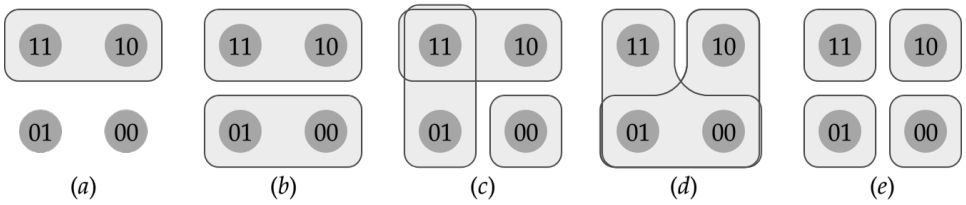

Figure 5.5: Inquisitive issues, again.

Sentence (44) merely expresses (a), the proposition that Johnny is coming to the party. This is just a sanity check to make the illustrations clearer.

Sentence (45) expresses (b). This is an issue dividing the worlds into those at which Johnny is coming to the party and those at which Johnny is not.

Sentence (46) expresses (c). The upwards arrow indicates emphasis of the preceding word when uttering the sentence. Here, the issue is resolved by an information state which entails that Johnny is coming, by an information state which entails that Mary is coming, or by an information state that neither is coming to the party. Note that the set of worlds $\{11, 10, 01\}$, i.e. the set of worlds at which either Johnny or Mary is coming to the party, would *not* resolve the issue (c). This is intentional: the structure of the uttered disjunctive interrogative sentence is such that this set of worlds

wouldn't answer the question. Or, at least, that's the assumption made in the framework.

Sentence (47) expresses (d). The conditional makes the issue slightly more complex and counter-intuitive, and does not contribute to the discussion relevant to this thesis. On the face of it, the question is similar to (45), and just requires a "yes" or "no" answer. However, the overlapping information states pose a challenge. Suppose that 01, i.e. that Johnny is not coming to the party, but Mary is. It seems that the answer to the question is then both yes *and* no. But this merely indicates that conditionals aren't truth-functional, says the inquisitive semanticist. The inquisitive analysis of conditionals requires more elaboration (see e.g. Ciardelli, 2016; Ciardelli et al., 2017).

Assigning issues to these sentences is not always determinate, and leaves room for deliberate choices. Consider (48). Sentence (48) may express (e), in what's called a *mention-all* interpretation, or it may express (c), in what's called a *mention-some* interpretation (cf. Groenendijk and Stokhof, 1984b). Mention-all interpretations require an exhaustive answer, while mention-some interpretations do not. Notice, however, that in neither interpretation, the information state that Johnny is *not* coming to the party resolves the issue. It is not entirely clear to me how the structure of the question necessitates a choice between these two particular interpretations, as opposed to the interpretations developed in the following section. In any case, which different interpretations are valid interpretations seem to be extraneous to the inquisitive framework.

5.3 In the Coarse-Grained Model

The Coarse-Grained Model developed in this thesis provides a notion of answer which can also address the above concerns, and additionally sheds some light on the relations between different types of answers.

Concretely, the account takes as a starting point the Hamblinian question meaning: for a given interrogative sentence (i.e., the question as a linguistic object), take a set of propositions as its denotation. These propositions are called *question-alternatives*. One way to illustrate this is to follow Hamblin in taking as question-alternatives all propositions resulting from substituting the salient range of names for the interrogative pronouns. For example,

the question-alternatives of "Who is coming to the party?" is the set of propositions such for each contextually salient individual, some proposition in the question-alternatives decides whether they are coming to the party. It seems like this is a good first approach to systematically tackle this issue. However, the Coarse-Grained Model is not bound to any particular way of setting up the question-alternatives. All that is required, technically, is there to be some privileged set of propositions which pick up crucial features of the question. For this reason, the account is compatible also with *why*-questions, if the meaning of the question is given as a set of propositions. So, for example, the meaning of "why did you do this?" could be given by a set of propositions specifying separate reasons why you did this. The account then would proceed as detailed below. But this would have to be worked out in a different paper.

Back to the question "Who is coming to the party?". If the salient individuals are, say, Robert, Johnny and Mary, then the question-alternatives are the set consisting of the propositions that Robert is coming to the party, that Johnny is coming to the party and that Mary is coming to the party:

$$\begin{aligned} q\text{–alt} = \{&\{w \mid \text{Robert is coming to the party at } w\},\\ &\{w \mid \text{Johnny is coming to the party at } w\},\\ &\{w \mid \text{Mary is coming to the party at } w\}\}. \end{aligned}$$

We can write this shorthand as q–alt $= \{R, J, M\}$. Depending on the purpose, one can require that question alternatives are closed under negation, but nothing important for this account depends on this decision. Relative to a context, the propositions might stand in any logical relation to one another. For example, it might be contextually known that whenever Robert is coming to the party, Johnny stays home. Without any context, these propositions are expected to be logically independent, although one could think of some far-fetched exceptions, like some individual might necessarily be a proper part of another. The question-alternatives mirror the simple Hamblian answers to a question, cf. Fig. 5.6.

The question alternatives induce a partition. Each partition cell decides for every question alternative whether the alternative is the case. This can be conveniently illustrated like in the case of partition semantics, compare

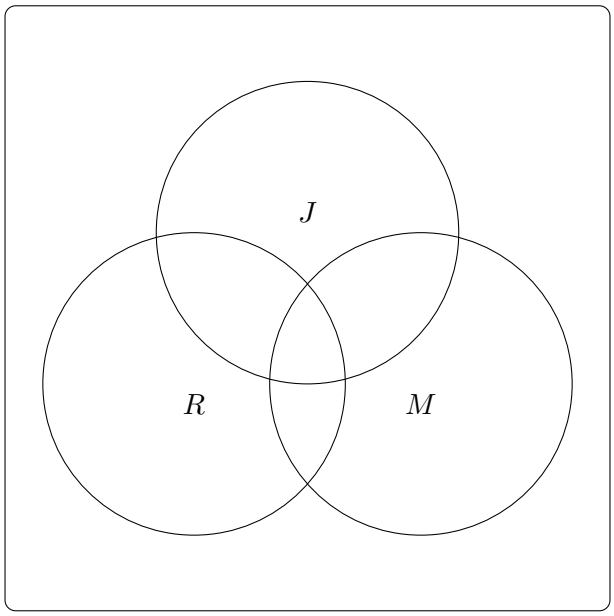

Figure 5.6: Question-alternatives in the Coarse-Grained Model.

RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$
$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$

Figure 5.7: Question-alternatives partition logical space.

Fig. 5.7.

So for example, one particular partition cell only includes worlds at which Robert is coming to the party, but Johnny and Mary are not. Compare Fig. 5.8, in which this partition cell is shaded.

Partition cells represent any maximally specific Boolean combination of question-alternatives. Unions of partition cells can represent partial information. For example, there is a union of partition cells such that at each world in this union, Robert is coming to the party, but whether Mary or Johnny are coming is not decided. Compare Fig. 5.9, in which this union is shaded.

RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$
$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$

Figure 5.8: Robert is coming to the party, but Johnny and Mary are not.

RJM	$RJ\bar{M}$	$R\bar{J}M$	$\bar{R}JM$
$R\bar{J}\bar{M}$	$\bar{R}J\bar{M}$	$\bar{R}\bar{J}M$	$\bar{R}\bar{J}\bar{M}$

Figure 5.9: Robert is coming to the party, but whether Johnny and Mary are coming is not decided.

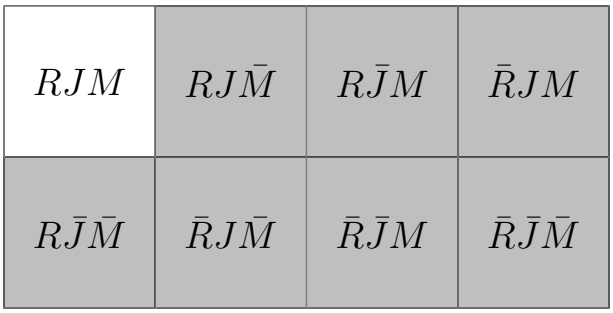

Figure 5.10: Robert, Johnny, and Mary are not all coming to the party.

Similarly for any other Boolean combination. Another example: the union of all cells without the cell at which all three are coming to the party represents the information that it is not the case that all three come to the party, but not more information than that. This union does not answer, for instance, whether Robert comes to the party. Compare Fig. 5.10, in which this union is shaded.

It seems natural, then, to distinguish broadly between three different types of answers to a question. These correspond, intuitively, to degrees to which a question is answered. If only a single partition cell is excluded, there is some information gain w.r.t. the question, but only very little. This type of answer is the weakest. If the answer decides for some individual whether they come to the party, we gain more information than from the weakest answer. This type of answer partially answers the question. Then there is a single partition cell. This is the maximal informational gain w.r.t. the question. The answer is complete. Complete and partial answers have already been described by Roberts (2012). In addition, some cases require for *positive* information. For example, if the assistant asks their supervisor what they should be doing next, it might not be sufficient for the supervisor to state "Don't bother sending the email to professor Heisenberg." Even though this would constitute a partial answer, since it decides whether the task of sending an email to prof. Heisenberg should be done next. But the assistant still does not know what to do next. A *positive* answer like "Finish your research paper!" is much more helpful.

Orthogonal to the distinction between complete, positive partial, negative partial and weak answers is a second one: is the answer precisely a union of partition cells, or does the answer merely entail a union of partition cells?

One can easily verify that any proposition entails a unique minimal union of partition cells. Which relations hold among these different notions of an answer?

To discover these relations, the notions must be made reasonably precise. So let's start with some definitions. As above, I assume that for a question literal suitable question-alternatives, i.e. a set of propositions, is given. That is, there is a set q–alt which contains some propositions $p_1, \ldots, p_n \in \Omega$, where Ω is some set and can be thought of as a totality of possibilities. We assume here finitely many partition cells. The question-alternatives q–alt induce a partition Π_q, also called q-partition, on Ω such that each partition cell decides each question alternative. This means that for propositions $p_1, \ldots, p_n \in q$–alt, a partition cell $\pi \in \Pi_q$ is defined as $\pi = x_1 \cap \ldots \cap x_n$, where $\pi \neq \emptyset$ and either $x_i = p_i$ or $x_i = p_i^c$. The proposition p_i^c is the complement of p_i w.r.t. Ω. With this simple toolkit, we can define the different types of answer:

Definition. A non-empty proposition φ is a *complete answer* (CA) to a question q if φ entails a q-partition cell, i.e. $\varphi \subseteq \pi$ for some $\pi \in \Pi_q$.

Definition. A non-empty proposition φ is a *positive partial answer* (PPA) to a question q if φ entails at least some q-alternative, i.e. if $\varphi \subseteq a$ for some $a \in q$–alt.

Definition. A non-empty proposition φ is a *negative partial answer* (NPA) to a question q if φ entails at least some complement of a q-alternative, i.e. if $\varphi \subseteq a^c$ for some $a \in q$–alt. Equivalently, if $\varphi \cap a = \emptyset$.

Definition. A non-empty proposition φ is a *partial answer* (PA) to a question q if φ is a positive partial answer or a negative partial answer. That is, if $\varphi \subseteq a$ or $\varphi \subseteq a^c$ for some $a \in q$–alt.

Definition. A non-empty proposition φ is a *weak answer* (WA) to a question q if φ entails the complement of at least one cell of the q-partition, i.e. if $\varphi \subseteq \pi^c$ for some $\pi \in \Pi_q$.

The types of answer stand in certain fairly straightforward relationships to one another. See Fig. 5.11 for a relationship diagram of complete (CA), positive partial (PPA), negative partial (NPA), partial (PA) and weak (WA) answers. A line with an arrow head from one type of answer to another indicates entailment in direction of the arrow. For example, a positive partial

answer (PPA) is also a partial answer (PA) and a weak answer (WA). A complete answer is always either a positive partial answer or a negative partial answer, or both.

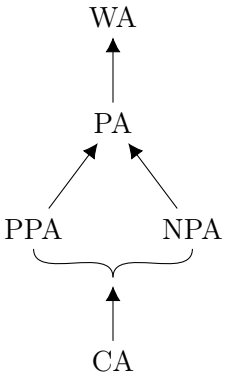

Figure 5.11: Relationship diagram of different types of answers. See the appendix for proofs.

For each of these types of answers there is a coarse version that is exactly a union of q-partition cells:

Definition. A complete (positive partial, negative partial, partial, weak) answer φ is a *coarse* complete (positive partial, negative partial, partial, weak) answer if it *is* a union of cells of the partition induced by the q-alternatives, i.e. if there are $\pi_1, \ldots, \pi_m \in \Pi_q$ s.t. $\varphi = \cup \pi_i$.

With these, we can state the following basic remarks:

Remark. For each answer φ and a question q there corresponds a unique *coarsening* $C(\varphi)$: the minimal union of $q-$partition cells s.t. $\varphi \subseteq C(\varphi)$.

Compare Fig. 5.12 for a visual representation of coarsening. Proposition φ, represented by the blob-shape, is a subset of some totality of worlds. The partition Π, represented by the grid, is induced by a question q. Finally, $C(\varphi)$, shaded grey, is the coarsening of φ for q.

Proof. Elementary, see Appendix.

Remark. If a proposition φ is a complete (positive partial, negative partial, partial, weak) answer to a question q, then the coarsening $C(\varphi)$ is a *coarse* complete (positive partial, negative partial, partial, weak) answer to the question q.

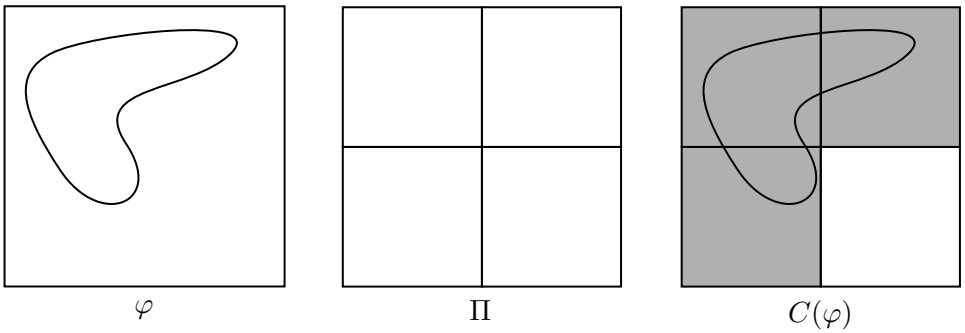

Figure 5.12: Visual representation of a coarsening.

Proof. Elementary, see Appendix.

The relations between notions of answers are mirrored in their coarse counterparts. Additionally, any coarse complete (positive partial, negative partial, partial, weak) is, by definition, also a complete (positive partial, negative partial, partial, weak) answer. See Fig. 5.13 for a relationship diagram of complete (CA), positive partial (PPA), negative partial (NPA), partial (PA) and weak (WA) answers, as well as coarse complete (CCA), coarse positive partial (CPPA), coarse negative partial (CNPA), coarse partial (CPA), and coarse weak (CWA) answers.

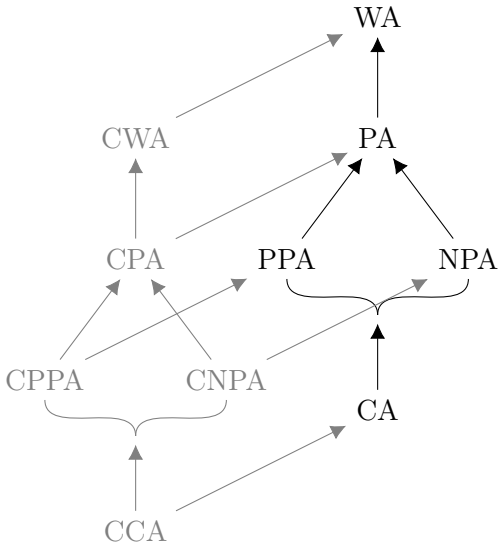

Figure 5.13: Relationship diagram of different types of answers, including coarse answers.

How do the accounts of Hamblinian question meaning, question meaning

in partition semantics, and inquisitive question meaning relate to the coarse-grained question meaning laid out here? Hamblinian question meaning is the starting point of coarse-grained question meaning. And it survives as part of the taxonomy: coarse partial answers are equivalent to an answer in Hamblinian question meaning. However, with a caveat. While coarse-grained question meaning also allows coarse answers with more information than a partial answer to count as partial answers, they would not count as answers on the simple Hamblinian account.

Partition semantics requires the single answer to a question at a world to be its true exhaustive answer. An answer in partition semantics is thus always a complete (exhaustive) answer. Groenendijk and Stokhof (1997) actually use the terminology "complete and precise" answer already, which they use, as far as I can see, in the sense of *complete* and *coarse* laid out here.

For the standard inquisitive question meaning, the case is just as straightforward. Inquisitive semantics requires question meaning to consist of downward closed sets. In other words, inquisitive question meaning requires to also count each subset of an answer as an answer. This is exactly what the notion of a (non-coarse) partial answer provides. If $p \in q$–alt, then also $p' \in q$–alt for every non-empty $p' \subset p$. In other words, any proposition which entails a proposition or its complement from the question-alternatives is a partial answer to the question.

To illustrate this exposition, let's look at a concrete example, and limit our view to coarse answers to the question "Who is coming to the party?". To make matters manageable, further assume the only contextually salient individuals are Johnny and Mary. Take J to stand for the proposition that Johnny is coming to the party, and M the proposition that Mary is coming to the party. Negation is indicated with a bar: $\bar{J}$ says that Johnny is not coming to the party, and $\bar{M}$ says that Mary is not coming to the party.

What are the possible answers to the question, according to Coarse-Grained Question Meaning? The question induces a partition with four partition cells. At worlds in the first partition cell, Johnny and Mary are both coming to the party. I'll denote this cell with JM. In the second partition cell, only Johnny is coming ($J\bar{M}$). In the third, only Mary is coming ($\bar{J}M$). And in the fourth, neither Johnny or Mary is coming to the party ($\bar{J}\bar{M}$). There are then four complete answers, corresponding to the

partition cells: $JM, J\bar{M}, \bar{J}M, \bar{J}\bar{M}$. Other answer types can be constructed by taking different unions of these partition cells. There are only four such unions which correspond to partial answers (besides the complete answers): The proposition that Johnny is coming to the party ($JM \cup J\bar{M}$) and the proposition that Mary is coming to the party ($JM \cup \bar{J}M$), the proposition that Johnny isn't coming to the party and the proposition that Mary isn't coming to the party. Other non-total non-empty unions of partition cells constitute a weak answer.

By virtue of their set-structure these different answers for a given question come with a natural order. This order is called an *inclusion order* or simply partial order on the set of partition cells. The answers form a lattice. Compare Fig. 5.14 for an illustration as a Hasse diagram. Fig. 5.14 and the following figures have been directly inspired by Rutherford (1965).

Partial answers induce substructures on this order. The maximal element of this substructure is the union of partition cells identical to the question alternative this partial answer evaluates. Compare Fig. 5.15 and 5.16. In Fig. 5.15 partial answers deciding for Johnny whether he comes to the party are highlighted. For example, JM is a partial answer, but also $JM \cup J\bar{M}$. In Fig. 5.16, partial answers instead deciding for Mary whether she comes to the party are highlighted.

The structures described here might look familiar, and indeed they correspond exactly to a free boolean algebra. An example of such an algebra is given in Fig. 5.17.

The inclusion order creates a straightforward order on the degree of informativeness w.r.t. the question. Complete answers are the most informative, while partial answers are more informative than some weak answers. The different answers are separated into levels of informativeness, here represented by horizontal rows of answers. The number of different answers on one such level is a function of the total partition cells n and the number k of cells the answers of that row consists of. It can be calculated simply by taking the binomial coefficient $\binom{n}{k}$.

Why are these different notions of answer important or relevant to a philosophical discussion? Several reasons. First, the account is more general than Hamblian question meaning, question meaning in partition semantics, and inquisitive question meaning. Additionally the account *unifies*, i.e. subsumes in a single picture, all of Hamblian question meaning, question meaning in

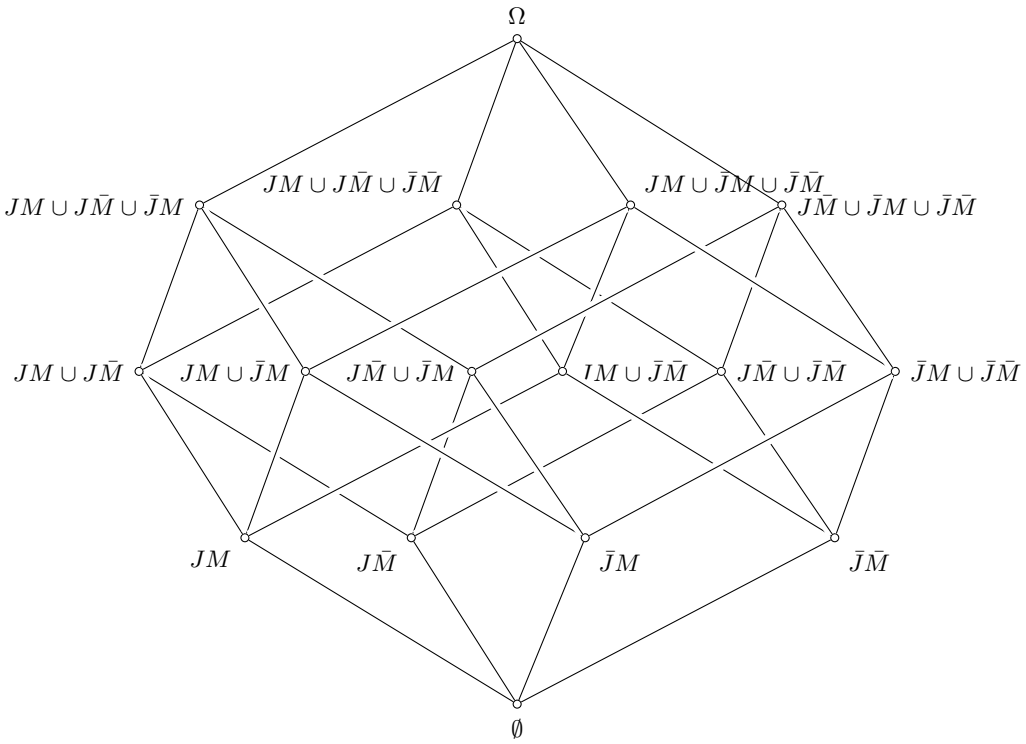

Figure 5.14: Answers to the question "Who is coming to the party?" ordered by inclusion.

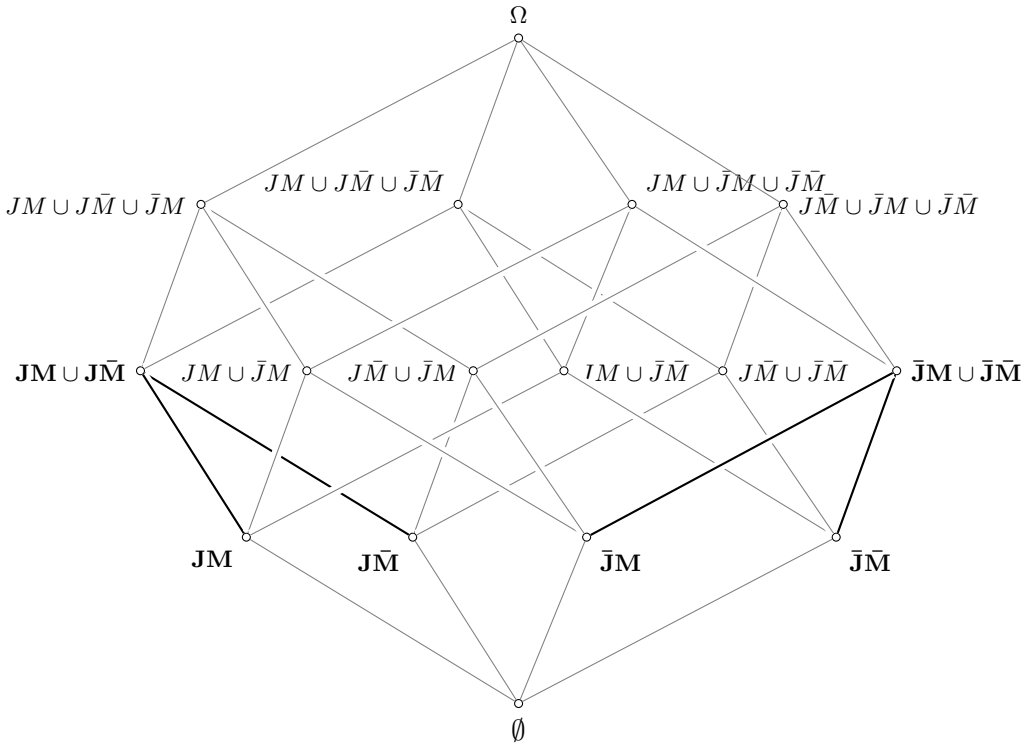

Figure 5.15: Answers to the question "Who is coming to the party?" ordered by inclusion. Partial Answers that decide whether Johnny is coming to the party are highlighted.

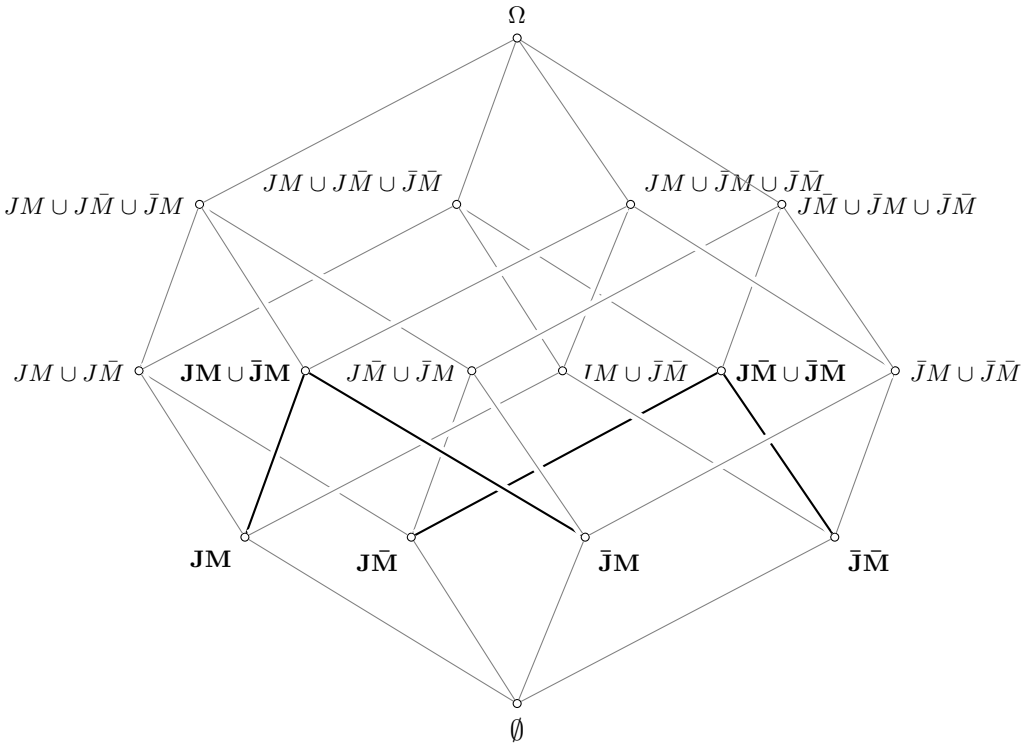

Figure 5.16: Answers to the question "Who is coming to the party?" ordered by inclusion. Partial Answers that decide whether Mary is coming to the party are highlighted.

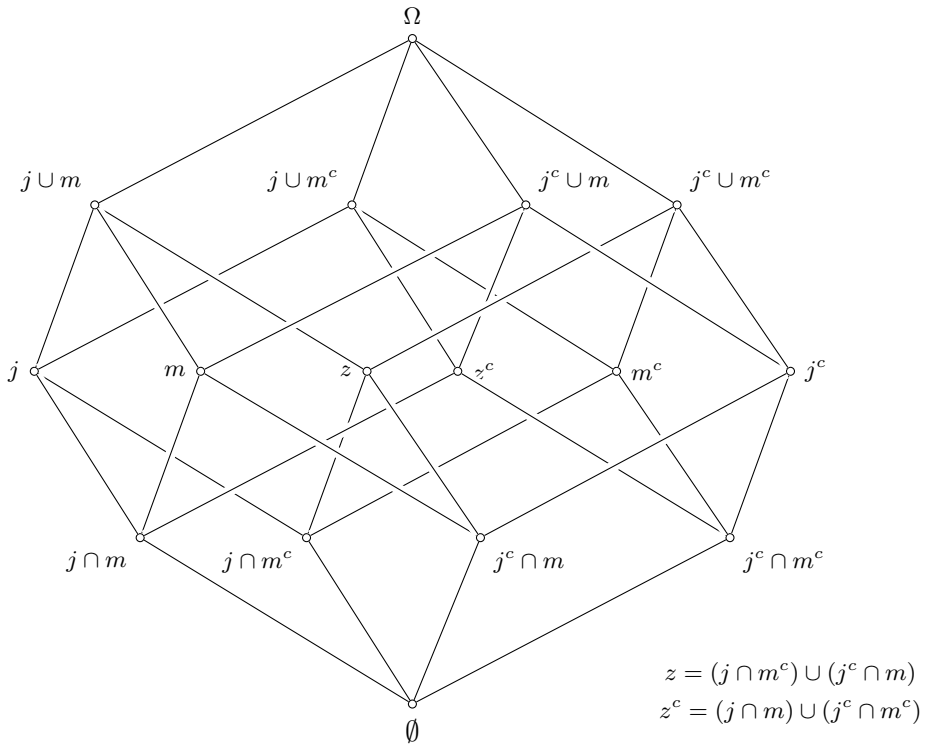

Figure 5.17: Free Boolean algebra generated from j and m.

partition semantics, and inquisitive question meaning. A second reason: as seen in the discussion of Schoubye and Stokke's account in section 4.4.1, careful analysis of the structure of answers reveals that there isn't always a unique answer determined by the "weakest proposition". If "weakest proposition" is taken to be the maximal element of set a partially ordered by the subset relation, it becomes clear that there is a weakest proposition which still answers the question only in certain benevolent cases.

5.4 Application of Coarse-Grained Answers

Finally, coarse-grained question meaning enables the Coarse-Grained Model of communication to apply a great deal of nuance when explaining cases of utterance-indeterminacy. Here is one such case: consider the a scenario in a lab. There is an arrangement of test tubes like in Fig. 5.18. This illustration is provided by MacFarlane (2020a), which I recreated for the purposes of this thesis.

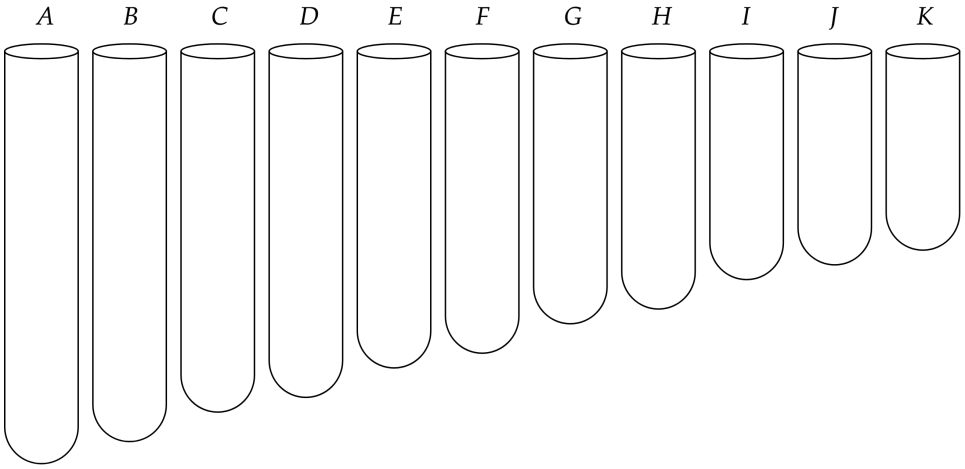

Figure 5.18: MacFarlane's test tubes.

The assistant (As) asks their supervisor (Su): "which test tube to you want me to clean?" The question-alternatives can be described by

$$\begin{aligned} q\text{–alt} = \{&\{w \mid \text{Su wants As to clean test tube A at } w\}, \\ &\dots \\ &\{w \mid \text{Su wants As to clean test tube K at } w\}\}. \end{aligned}$$

The supervisor responds to their assistant:

(49) I want you to clean this test tube.

Without indicating which of the test tubes A through K they refer to. The assistant has no way of determining the intended test tube, if there is any. Communication fails.

Suppose, however, the question was a more general "What do you want me to do?" The question alternatives could, depending on the context, then include all sorts of additional potential tasks the assistant might perform.

$$
\begin{aligned}
q\text{–alt} = \{&\{w \mid \text{Su wants As to clean test tube A at } w\},\\
&\ldots\\
&\{w \mid \text{Su wants As to clean test tube K at } w\},\\
&\{w \mid \text{Su wants As to write an e-mail to prof. Heisenberg at } w\},\\
&\{w \mid \text{Su wants As to } \Phi \text{ at } w\},\\
&\ldots\}.
\end{aligned}
$$

Then (49) does not constitute a partial answer, for there is no question-alternative that the answer decides. It might be any of the test tubes. But (49) *does* constitute a weak answer, for it excludes all partition cells which decide the test-tube-cleaning question-alternatives for test tubes A to K in the negative. After all, there is *some* test tube the supervisor wants the assistant to clean. Does (49) answer the question? That seems to depend on our standard for an answer. If we require a partial answer, it does not, the communication may fail. If we allow weak answers, then the communication does indeed succeed. How is this standard set? I will not provide an answer in this dissertation, but there might an interesting pragmatic issue to contextually determine such a standard. The Coarse-Grained Model can model this subtle difference, but it is agnostic w.r.t. which standard to choose.

For a third example, consider again the assistant asking the supervisor "What do you want me to do?" The question-alternatives, let's assume, include some list of duties:

$$\begin{aligned} q\text{–alt} = \{\{w \mid & \text{ Su wants As to } \Phi_1 \text{ at } w\}, \\ \{w \mid & \text{ Su wants As to } \Phi_2 \text{ at } w\}, \\ & \ldots \\ \{w \mid & \text{ Su wants As to } \Phi_n \text{ at } w\}\}. \end{aligned}$$

The supervisor answers with:

(50) Don't do Φ_{17}.

The response (50) constitutes a negative partial answer, as it negatively decides whether Su wants As to Φ_{17}. But it seems that the assistant still does not know what to do. The information is not enough to act upon. Even though the response is a partial answer. It seems that here, there are pragmatic reasons why a *positive* partial answer is required. Any positive partial answer gives the assistant information to act upon.

So much for questions and answers. The next chapter discusses further concepts necessary to then state the Coarse-Grained Model in the subsequent chapter.

Chapter 6

Central Concepts

6.1 Introduction

This chapter discusses the notions of *common ground*, which is an established notion in the literature. Further, the chapter takes up the notions of *subject matter* and *question under discussion*, which are shown to be closely linked. Finally, the chapter introduces the notion of *permissible interpretation*, which is an idiosyncrasy of this thesis. Yet, the ideas involved in the notion are well supported by established research. With these central notions as well as those of *questions*, *answers*, and *coarsening* described in chapter 5, the Coarse-Grained Model can be stated in the following chapter.

6.2 Common Ground

Already briefly discussed in chapter 2, Robert Stalnaker developed in a series of papers a method to model conversational contexts and conversational contributions using broadly the tools of truth-conditional analysis (Stalnaker, 1974, 1975, 1978, 1986, 1988, 2002, 2004). These notions were in part developed to analyze the contents and effects of speech acts, and in particular the effect an assertion may have on the contextual situation it was uttered in (most explicitly in Stalnaker, 1978). The content of an assertion is affected by its context, and in turn the assertion acts upon the context, too. Stalnaker develops his notions to analyze a single step in the shifting dynamics of a conversation. Stalnaker notes among his major influences the work of David Kaplan (Kaplan, 1968, 1989) as well as Paul Grice (Grice, 1989).

Propositions are assumed to be represented by sets of possible worlds. In a conversation, participants have certain presuppositions. Stalnaker cashes out the notion of presupposing that p as the disposition to act as if to assume or believe that p (Stalnaker, 1978). For the present purposes, the exact definition of a presupposition does not matter much. These presuppositions are taken to be the common ground of speaker and audience. The common ground is represented by a *context set*—the set of possible worlds compatible with what is common ground. For example, if some propositions P and Q constitute the common ground, then the context set is their intersection. Sometimes, the labels *common knowledge* or *mutual knowledge* are used. These terms have technical definitions; Vanderschraaf and Sillari (2023) distinguish: a proposition P is mutual knowledge between speaker and audience if both speaker and audience know that P. P becomes *common* knowledge if the speaker also knows that the audience knows that P, and the audience also knows that the speaker knows that P, and so on, iterating potentially indefinitely. Stalnaker makes it clear that knowledge of P is not required in order for P to be counted as common ground (Stalnaker, 1978). The proposition just needs to be presupposed. Presupposing does not require believing, though, and much less requires knowing. Presuppositions do not have to be mutual, either, and here the notion becomes subtle. For Stalnaker, each discourse participant—for present purposes, speaker and audience—have their own presuppositions. For these presuppositions to count as common ground, any discourse participant, let's take the speaker, also has to presuppose that the audience shares those presuppositions. That is, the speaker assumes that their own presuppositions are mutual. So, what the speaker takes to be common ground is what the speaker presupposes to be the case, if the speaker also assumes that the audience shares that presupposition. This requires the presuppositions to be assumed *common* in the above sense described for common knowledge. The speaker's assumed common ground can be represented by a context set. For the audience, the same reasoning holds. They have their own context set. In other words, the speaker takes some information to be common ground, while the audience takes some—potentially different—information to be common ground. Stalnaker immediately notices that this may become problematic, and defines: In a *nondefective context*, the speaker's context set and the audience's context set align perfectly. In a *defective context*, both sets come apart, with

potentially destructive consequences to their ability to communicate with one another: since context influences content, the audience might take the speaker to say something different to what the speaker intended to express. For example, Smith tells Jones: "Adams took care of their trouble," talking about the recent weight gain/loss issues Adams has been experiencing. But Jones isn't aware of those, and instead remembers the trouble Adams had trying to better their golf handicap. Jones takes up something else entirely about Adams than Smith intended to express. This consequence does not *have* to be catastrophic, though: Stalnaker defines a context *close enough* to being nondefective "if the divergencies do not affect the issues that actually arise in the course of the conversation" (Stalnaker, 1978). Stalnaker leaves this last notion intuitive without making it more precise.

Now, assuming a nondefective context, an assertion made in that context potentially alters the context set: if what has been asserted is accepted, that is, henceforth presupposed by the participants, it gets added to the common ground. The context set is updated to reflect this new information. All possible words incompatible with what has been said are discarded from the context set. This highlights the dual interaction between context and content: the context of an utterance influences its content, and the content in turn influences the subsequent context. Context is "both the object on which speech acts act and the source of the information relative to which speech acts are interpreted" (Stalnaker, 1998). This analysis gives an idea of how the dynamics of a conversation might unfold. In this thesis though, as mentioned, I employ Stalnaker's common ground not as an analysis of the dynamical development of a conversation in time, but rather as an analysis of a synchronic "snapshot" of the content and effect of a single utterance.

To sum up: speaker and audience presuppose propositions, and this information makes up the common ground. What they presuppose can come apart, in more or less problematic ways. What's presupposed helps determine what is expressed with an utterance, and the utterance's content in turn influences the subsequent presuppositions.

6.3 Question under Discussion and Subject Matter

When my friend excitedly exclaims:

(51) Romania beat Ukraine 2:1!

Referring to the match in the UEFA Euro 2024, she claims something false. The match ended 3:0 in favor of Romania. But intuitively, there is something right about her claim. And there is: she is right about the fact that Romania won the match, but not about the result. What explains this intuition? Examined merely truth-conditionally, (51) expresses a proposition. If propositions are sets of worlds, the actual world is not part of the proposition expressed by (51). That Romania beat Ukraine is logically implied by (51). So maybe even though her claim is false, the intuition that she said *something* right is explained by appeal to what's logically entailed by her claim. But this cannot be the whole story: Her utterance (51) also implies that either Snow is yellow or it is not. But it seems rather odd to say that she is right about *that* when uttering (51). Yablo (2014) introduces some influential ideas: My friend is indeed right about something when she utters (51). She is right about *part* of what she expressed. There is a proposition she is right about which is part of the content she expressed. This proposition is that Romania beat Ukraine. Yablo picks up an example by Goodman (1961): *Maine is prosperous* is about New England, Goodman claims. Why? Because it logically implies that *of the New England states, at least one is prosperous.* But mere implication, as Goodman proposes, overgenerates: *Maine is prosperous* is not about Texas just because it logically implies that *Maine is prosperous, or Texas is.* Yablo (2014, pp. 15) proposes the following condition to determine part-hood.

A proposition B is part of a proposition A the inference from A to B is both:

(i) truth-preserving: A logically implies B, and

(ii) aboutness-preserving: A's subject matter includes that of B.

So in addition to logical implication (i), another condition is required. This condition (ii) requires a particular relation between the subject matters of both propositions. But what is the subject matter of a proposition? And what does it mean for one subject matter to include another? These questions are central to Yablo's influential account of *aboutness*, or *topic*, or subject matter.

For Yablo, a subject matter induces an equivalence relation—or partition on logical space. Logical space is treated here as the set of all possible worlds. The subject matter partitions worlds into equivalence classes if they are "in-

discernible where that subject matter is concerned" (Yablo, 2014, pp. 26). Consider as an example the subject matter of how many stars there are. Then two worlds are in the same equivalence class precisely if they have equally many stars. What, then, are subjects? Any kind of thing by which one can distinguish possible worlds, or "a system of differences, a pattern of cross-world variation" (Yablo, 2014, pp. 27). Yablo gives examples: the western hemisphere, Queen Victoria, the nineteenth century. But not *literally* the nineteenth century: "living in the nineteenth century is not the same as living in an equivalence relation" (Yablo, 2014, pp. 26). Worlds w and v are grouped into the same equivalence class by the subject matter *the nineteenth century* if the nineteenth century in w is an intrinsic duplicate to the nineteenth century in v (Lewis, 1988). The metaphysically rather iffy "intrinsically indiscernible" will have to be left unexplained at this stage, as it leads too far afield. Subject matters described like this are slightly difficult to conceptualize, I must admit, but Yablo provides a simpler way of thinking about them: as sets of propositions, which are the answers to *questions*. As explained in detail in chapter 5, questions partition logical space in a particular way. This particular way might well be conceptualized as a subject matter; and it's quite straightforward to understand subjects matter this way. How does the subject matter *the number of stars* partition logical space? Take the question "how many stars are there?" and the set of propositions making up answers to this question. A proposition for any n that there are n many stars. How does the subject matter *the nineteenth century* partition logical space? One generic way to ask is "what happened in the nineteenth century?", let's call this question q. Answers to this question state for any given event whether it happened in the nineteenth century. Worlds w and v are in the same partition cell if and only if everything that happened in w in the nineteenth century also happened in v, and *vice versa*. This question necessitates a rather fine-grained partition, as one can imagine; there are many, many events which happened (and didn't happen) in the nineteenth century. A more specific question is "Which major historical events happened in central Europe in the 1870s?", let's call this question q'. This question is more specific—and its corresponding partition is *less* fine-grained than the more general question. In terms of subject matter, Yablo says that the subject matter given by question q (the more general one) *includes* the subject matter of question q'. The partition given by q is

a *refinement* of the partition given by q'.

Let's use these conceptual tools to explain the introductory example (51). Recall that my friend excitedly exclaimed:

(51) Romania beat Ukraine 2:1!

And intuitively, she got something right about the match, namely that Romania indeed beat Ukraine. She got the exact result wrong, though. What's the subject matter? One plausible candidate can be represented by the question "what was the result of Romania against Ukraine?"—partitioning logical space into equivalence classes of worlds.[1] Worlds are in the same equivalence class precisely if the result is the same, regardless of further details about the match. For example, it does not matter who scored, or whether there were any red cards, etc. Now consider just "Romania beat Ukraine. The question we might plausibly take to represent its subject matter is "did Romania beat Ukraine?", which partitions logical space into those worlds where Romania won and those where Romania didn't (including draws). This latter partition about whether Romania won is less fine-grained than the former partition about the exact result. But each result completely determines whether Romania won or not. In other words, the former partition is a refinement of the latter. The subject matter of "Romania beat Ukraine 2:1" *includes* the subject matter of "Romania beat Ukraine". And with this, the second condition (ii) above is satisfied. The proposition that Romania beat Ukraine is part of the proposition that Romania beat Ukraine 2:1, because it is truth-preserving (i) as well as aboutness-preserving (ii).

Why is the inference from (51) (that Romania beat Ukraine 2:1) to the proposition that Romania beat Ukraine 2:1 or snow is yellow not aboutness-preserving? To answer this, consider the subject matter of the proposition that Romania beat Ukraine 2:1 or snow is yellow. Assume that the proposition that snow is yellow is a contingent (if false) proposition and also logically independent of (51). The resulting partition of the disjunction is *more fine-grained* than the partition generated by the subject matter of (51). This is

[1] Potentially, one could also say something right when saying that Poland beat Ukraine 3:0, instead of Ukraine. The question representing the subject matter should therefore be more general, something like "who beat whom with which result?", referring to that particular match. I leave out this complication to make things more straightforward. But the argument stands in any case: such a question would generate a partition which includes the partition given by "What was the result of Romania against Ukraine?". By transitivity of the partition-inclusion relation, the same holds for the partition given by "Did Romania beat Ukraine?".

because any partition cell answering "what is the result of the match Romania against Ukraine?" is further subdivided into cells in which snow is yellow and cells in which snow is some other color. If this is right, however, the subject matter of (51) does not include (i.e. not is a refinement of) the subject matter of the proposition that Romania beat Ukraine 2:1 or snow is yellow. Thus, the inference is not aboutness-preserving.

Subject matters are a way of explaining the intuition about this case. Questions can be used to clarify subject matters. Questions can determine a partition, and subject matters can be represented as a partition. Thus questions in this sense are closely related to subject matters. Questions can take the role of determining the subject matters of propositions. Conversations, too, have subject matters; conversations are about something; they have a co-operative purpose (compare the discussion of Gricean pragmatics in section 2.4). This purpose can be pragmatic. One way to represent such a purpose is informational: there is a *question* the conversation is supposed to answer. It's common to represent such an informational purpose as a *question under discussion*. In the literature, this is often abbreviated as *QUD*, which I will not adopt. It's a pragmatic feature of the conversation, something embedded in its context. The idea is that a conversation is supposed to answer a certain question, and this question is used by the participants of the discussion to infer contextual clues or fill in missing information. The Coarse-Grained Model developed in this thesis makes heavy use of the notion of conversational purpose as a question under discussion. With such a question, indeterminate language can be made determinate in ways detailed in chapter 7.

The Coarse-Grained Model is, of course, not the first account to use a question under discussion as a tool in pragmatics. As far as I see it, it is the first to successfully defend the Classical Model of Communication with the help of a question under discussion. I already discussed some of those in chapter 4. Hoek (2018) employs subject matter in form of a partition on logical space. Subject matter helps to explicate cases of pragmatic weakening, by determining the contextual presupposition subtracted from what has been literally said. Schoubye and Stokke (2016) intend to determine *what is said* by an utterance in context with the help of a question under discussion. Their account applies at least in cases of *indeterminacy due to opaque contexts*, and the example (29) discussed in this thesis is their main case. The

idea is that what is said is a way of contributing to the purpose of the discussion. Their account falls short, however, as I point out in detail in chapter 4. The Coarse-Grained Model developed in this thesis manages to overcome the shortcoming of this account. Bowker (2022) uses a question under discussion to give a pragmatist truth-conditional response to context-switching arguments (see section 4.2.1). Leitgeb (2014, 2017) develops a thesis about the relationship of full belief and degrees of belief: a rational agent believes a proposition if and only if she assigns a stably high subjective probability to it. A *stably* high subjective probability is cashed out as stably high under probabilistic conditionalization on the members of a set $\mathcal{Y}$. That is, for any given subjective probability in a proposition P, it's not enough for full belief for the subjective probability to exceed some threshold r, it must additionally exceed that threshold conditional on a set of (perceptual, hypothetical, testimonial, etc.) propositions $\mathcal{Y}$. To see some analogy to a question under discussion, recall that chapter 5 defines the denotation of a question as a set of propositions which in some sense give an answer to the question. By defining stably high degrees of belief w.r.t. a set $\mathcal{Y}$, the theses spelled out in Leitgeb's account are relative to that set. The set is determined by the "context of reasoning of the *rational agent who has the beliefs*. Such a context will then include the agent's 'question' or partitioning of possibilities (...)" (Leitgeb, 2017, p. 99). This observation leads naturally to research questions: what kind of interrogative sentences (the denotation of which is a set of propositions) characterizes some particular set $\mathcal{Y}$ determining the relation between full belief and degrees of belief? Questions Leitgeb has in mind in the above quote read something like "how brave or cautious do I want to be?" (Leitgeb, 2017, p. 98), which does not in an obvious way fit with the account of questions and how they generate partitions in the Coarse-Grained Model (see chapter 5). Further, does the analysis of *answers* to a question given in the Coarse-Grained Model exhibit any applicability to the issue of stably high credence of belief? This latter research question aims at further distinguishing between sets of propositions to conditionalize on. For example, a given set $\mathcal{Y}$ determines a set of complete answers, a set of partial answers, a set of weak answers, and so forth. Unfortunately, this thesis is not the place to delve into these questions; this will have to wait for future work. Schaffer and Knobe (2012) as well as van Elswyk (2020), inspired by Lewis (1996), employ a question under discussion to address epistemic contextu-

alism. In particular, they use a question under discussion to explicate the idea that claims about knowledge-ascription change truth conditions with contexts. Skordos et al. (2022) employ a question under discussion in the cognitive developmental sciences: they hypothesize that children's understanding of quantified statements depends on a contextually given question under discussion.

The use of questions under discussion is, one might expect, not uncontroversial. Critical voices come, among others, from Picazo (2022) and Buchanan and Ian Schiller (2022). Since their criticisms can also be levied against the Coarse-Grained Model, I address the issues raised by these authors in chapter 9.

6.4 Permissible Interpretations

A further important part of the prerequisite concepts for the Coarse-Grained Model are *permissible interpretations*. This concept is used to model an agent's attitude towards the meaning of an utterance. When an utterance exhibits indeterminacy (see chapter 3), it is unclear which truth-conditions it expresses. For a particular agent, several different truth-conditions could potentially be expressed by the utterance. Any such set of truth-conditions deemed permissible by the agent is labeled as a permissible interpretation.

To make matters slightly more precise, let's think of truth-conditions as propositions. Propositions are assignments of binary truth-values to possible worlds. In this thesis, given a totality of possible worlds, propositions are total functions from possible worlds to the truth-values *true* and *false*. Sets of possible worlds correspond one-to-one to such an assignment, which is why I often interchangeably use either conceptualization of a proposition. Interpretations assign propositions to sentences. Since often the utterance in question is fixed, I also speak interchangeably of an interpretation as a proposition. An agent might consider several different interpretations as potentially expressed by an utterance of that sentence. The interpretations deemed potentially expressed in this way make up the set of permissible interpretations for that agent. Does that mean that the set of permissible interpretations is entirely subjective and agent-relative? In this thesis I claim that no, a rational agent is not at liberty to deem any interpretation permissible. Instead, there are certain objective constraints that interpretations

have to satisfy. This thesis is not the place to develop a full theory of what amounts to a permissible interpretation, but I will mention some guiding ideas. Generally, a permissible interpretation should coincide with what is conventionally established. Using English competently limits the range of permissible interpretations. "(...) one cannot use sentences such as 'Sue is ready' or 'Smith weighs 80 kg' to say that Louise is German or that the king of Sweden is a poor driver. There are, in other words, strict constraints on what sentences can mean" (Schoubye & Stokke, 2016, p. 787). A permissible interpretation does not wildly diverge from the literal meaning of a sentence. Instead, often expressions allow for contextual variations, and even the complete context does not always seem to suffice to uniquely determine a proposition. One way to account for the *origin* of permissible interpretations is to develop them from the grounds up in a theory about conventional agreement. Famously, David Lewis' doctoral thesis (Lewis, 1969) investigates conventional meaning as regularities in the outcomes of co-operation problems. Another way of thinking about restricting permissible interpretations was discussed in section . Buchanan's (Buchanan, 2010) *restricted proposition type* is a "partial structure which is determined by the lexical meanings of the uttered sentence's constituents in tandem with their syntactic arrangement".

The above does not make precise what exactly a permissible interpretation amounts to, but should give some indication. A way to make the notion more precise is provided by Leitgeb (2022). Here, interpretations are classified as *admissible* (the distinction between the choice of words *admissible* and *permissible* is, I think, negligible) if they satisfy certain contraints. Leitgeb states that the set of admissible interpretations forms something akin to a theory in sense of the non-statement view of scientific theories, described by Suppes (2000). What are these constraints? Leitgeb (2022, p. 905) describes these as given by *metasemantic facts* and *metasemantic laws* (paraphrasing):

(i) All linguistic facts concerning the competent usage of predicates and singular terms (individual constants, individual variables, functions terms) in a first-order language L.

(ii) All non-linguistic facts that are relevant as to whether the atomic formulas in L are satisfied.

(iii) All metasemantic laws taken together that concern the atomic for-

mulas, and hence the predicates and singular terms in L.

Simply put, facts in (i) constrain the truth-conditions of sentences by way of language competence, and facts in (ii) determine whether the truth-conditions hold at a possible world. Laws in (iii) give expression to particular metasemantic theories by which one assigns meaning. For example, a linguistic fact about the competent usage of a predicate like "tall" would require that whenever an interpretation i includes an object a in the extension of "tall", then i also includes all objects b in the extension of "tall" for which i includes the tuple $\langle b, a\rangle$ in the extension of "taller than". A corresponding non-linguistic fact requires that an object b is in fact of greater physical height than an object a whenever i assigns "taller than" to $\langle b, a\rangle$. Leitgeb gives the example of Kripkean baptism for proper names: "For all proper names a in L, for all objects d, if present usage of a in L is suitably causally connected to an act of baptism in which d was named an a, then $d \in Uni(I)$ and $I(a) = d$" (Leitgeb, 2022, p. 906). For context, Leitgeb talks about the intended interpretation I, which is an extensional interpretation of a first-order language and thus has a domain of objects Uni and assigns objects to proper names. This thesis on the contrary talks about *intensional* interpretations of a minimally formalized natural language. But it's straightforward to model many expressions of natural language discussed in this thesis as part of a simple propositional language with intensional interpretations. In this way, it's also straightforward to apply the Kripkean example to the type of interpretations used in this thesis. Leitgeb uses these contraints to develop a theory based on Ramsey-sentences about interpretations for semantically indeterminate language. But I'd wager they could be lifted from Leitgeb's theory to be used in the present thesis. It's one way to make explicit which interpretations should count as permissible for a rational agent to entertain. Leitgeb's theory even allows for the constraints to not uniquely determine a set of permissible interpretations, such that there is room for rational agents to entertain different sets of permissible interpretations. This becomes apparent when Leitgeb quotes Williamson (1994, p. 209): "meaning may supervene on use in an unsurveyable and chaotic way".

Another way to constrain which interpretations are permissible is by developing a theory of the diachronic dynamics of an agent's set of permissible interpretations. For example, consider a context in which a certain threshold is established for a vague predicate like "tall" such that anything below

that threshold is definitely not tall. Then one might formulate the requirement that a rational agent discards from their set of permissible interpretations those interpretations which still assign "tall" to something less tall than the contextually established threshold. This line of thinking strongly indicates that the set of permissible interpretations is to be considered context-dependent. A further way of constraining permissible interpretations is by prescribing how to retain contextual constraints between contexts. Does the contextually established threshold for "tall" carry over at all into other contexts? But this, alas, leads to far afield from the topic at hand. These considerations about permissible interpretations, then, shall suffice for the present moment.

I'll end the chapter by briefly mentioning two objections one might have to the philosophical validity of the notion.

The multitude of permissible interpretations describing multiple propositions potentially expressed by an utterance might appear awfully like the Multiple Proposition View. As discussed in section 4.2.2, the Multiple Proposition View takes the speaker to express a range of propositions with a single utterance, instead of just one. Similarly, the Coarse-Grained Model takes the audience to entertain multiple propositional interpretations of what the speaker might have expressed. But the similarities end when it comes to what is communicated: as is described in chapter 7, on the Coarse-Grained Model there is still a single proposition communicated, even though the audience deems many different interpretations permissible. The many permissible interpretations help determine what the speaker said: they necessitate a single answer to the question under discussion. This way, communication can successfully be explained with appeal to only a single communicated proposition. The Coarse-Grained Model thus is not a Multiple Proposition View. It does, however, take the initially appealing *intuition* of the multiple propositions and makes it compatible it to the Classical Model of communication.

A different objection notes that if permissible interpretations are objectively constrained in the way described in this section, the account seems to be unable to explain cases of mere implication which go against the literal meaning of the sentence uttered. For example, using irony, one can communicate propositions which even contradict the conventional meaning of a sentence. But how could permissible interpretations allow for such a content? This, indeed, is a trickier issue, and reasonably restricts the range of

application of the Coarse-Grained Model. I have more to say on this issue in chapter 9.

Chapter 7

The Coarse-Grained Model of Communication

7.1 Introduction

This chapter finally presents the Coarse-Grained Model. I'll take concepts and notions worked out in the previous chapters and use them to develop the model this thesis is meant to defend.

The Coarse-Grained Model extends the Classical Model presented in chapter 2. Again, the Classical Model is made up of the following principles:

Classical Pragmatics

(1) The content of an assertion is a (single) proposition.

(2) Uptake consists in recognizing the proposition asserted.

(3) If the assertion is accepted, its content is added to the conversational common ground.

Classical Content

Contents are ways the world might be (truth conditions).

The Coarse-Grained Model extends the Classical Model by providing a method to determine a proposition communicated. The model will yield a unique proposition expressed for a given utterance and a given context. This works in particular for cases which exhibit the types of apparent utterance-indeterminacy defined in chapter 3. Let's take stock. The central concepts the Coarse-Grained Model makes use of are the following.

- A context set of possible worlds representing the (assumed) shared background presuppositions of speaker and audience. See section 6.2.
- A question under discussion, representing the purpose or goal of the conversation. A question under discussion determines a set of propositions, called the question-alternatives, and a partition on logical space. See section 6.3.
- Coarse-grained question meaning, and in particular the definition of what amounts to an answer to the question under discussion. See chapter 5.
- The coarsening of a proposition is just that part of the proposition directly relevant to the question under discussion. See chapter 5.
- Permissible interpretations, which are those interpretations of the utterance deemed compatible with the utterance's literal meaning and contextual factors. Interpretations assign truth values to sentences at possible worlds. See section 6.4.

The resulting model connects these concepts in a simple manner. There are several points in this process where a decision between alternative ways of making the model more precise are necessary. I argue below that these decisions are well-motivated by directly implementing the Gricean Maxims. These decisions, then, are not *ad hoc*, and instead underline the pragmatic character of the model.

7.2 Implementing the Gricean Maxims

It's not exactly breaking new ground to emphasize the connection between Grice's Maxim of Relation—be relevant!—and discourse structured by questions. Already the *akademisch proefschrift* of Groenendijk and Stokhof (1984a) stresses the applicability of questions as a tool to purposes of conversations. Hoek (2018) motivates his theory of Conversational Exculpature by this observation. See section 4.4.2. Nevertheless, it's instructional and it gives additional import to the Coarse-Grained Model to show how parts of the model are direct implementations of Gricean Maxims, rather than ad-hoc choices for the formalism's sake.

The starting point is Grice's Cooperative Principle:

> Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. (Grice, 1975)

The guiding idea is that conversational contributions are assumed to further the purpose of the conversation. The purpose of a conversation can plausibly be explained as the joint pursuit of informational gain: find an answer to a question that is relevant to the interests of either discourse participant. This question is often referred to as the question under discussion. See section 6.3.

The Coarse-Grained Model deals with the simple act of literal assertoric communication in a person-to-person conversation as a speech act which is intended to influence the situation in which it is performed. The model accounts for indeterminacy by allowing the audience to entertain multiple potential propositions compatible with what they take the speaker to express. For example, a vague expression might semantically restrict its content to a range of propositions. Correspondingly, the audience has a range of permissible interpretations they might take the speaker to express. In addition, contextual factors can limit the range of interpretations further. The interpretations each assign a proposition to the sentence uttered. This leads to multiple propositions potentially expressed. Which unique proposition does the audience have reason to take the speaker to express?

First, the permissible interpretations have to be in some sense compatible with the literal meaning of the sentence type uttered. For example, "Snow is white" is not permissibly interpreted as that roses are red. Compare section 6.4 for more discussion.

Second, the Gricean maxims give normative clues to prefer some interpretations. Consider Grice's Maxim of Relation—simply "be relevant". Being relevant for Grice is to address the purpose of the conversation. Being relevant in the Coarse-Grained Model is answering the question under discussion. The model thus prefers such interpretations which are informative w.r.t. the question under discussion.

For example, consider the exchange:

(52) a. Question: Tomorrow is the final exam! Do you think Tipper is prepared?

b. Response: Tipper is ready.

Suppose that (52a) introduces a question under discussion into the conversation. The literal meaning of (52b) leaves open what it is that Tipper is ready for. On the Coarse-Grained Model, assuming the speaker to be cooperative entails assuming that they answer the question under discussion, Interpreting (52b) as saying, e.g. that Tipper has finished eating does not answer the question under discussion. Interpreting (52b) as saying that Tipper is ready for the exam, however, answers the question under discussion.

Furthermore, the maxim is evaluated against the body of information the speaker and the audience take for granted in the conversation. This body of information is taken to be common ground. Contents are propositions, and propositions are simple Kripkean sets of possibilities or possible worlds. The common ground is modeled as the set of those worlds compatible with what is shared background presuppositions. Making an assertion is *inter alia* a proposal to add to the common ground (cf. section 6.2). Thus, the model is interested in those interpretations which not just answer the question under discussion, but answer the question presupposing the information in the common ground. The context set includes those possible worlds of which the discourse participants think could be the actual world. In other words, worlds in the context set are *epistemically possible* from the joint perspective of the partitipants, provided that the context is not defective. This entails also that a preferred interpretation shall make the utterance true at some world in the context set, thus satisfying Grice's first Maxim of Quality—"Do not say what you believe to be false".

Grice's second Maxim of Quantity—"Do not make your contribution more informative than is required (for the current purposes of the exchange)"—introduces an additional requirement. The speaker is assumed to not provide superfluous information. To this end, on the Coarse-Grained Model, only the information directly relevant to the question under discussion is taken to be the content of the utterance. Information is, in a sense, subtracted from the permissible interpretations. See sections 6.3 and 4.4.2 for more discussion on this idea. The model then finds an answer to the question under discussion which is just as much or more coarse-grained than any particular preferred interpretation. This ensures satisfying the Maxims of Quantity. I will return to how this is achieved and give a more precise description below.

Although Grice notes that the Maxims of Manner are "of less urgency", they provide additional reasons for the audience. Consider "avoid obscu-

rity of expression" and "avoid ambiguity". By taking the speaker to follow these maxims, the audience has reason to assume that the context provides enough clues to infer a determinate content for the sentence uttered containing indeterminate expressions. That is, the audience has reason to believe that the speaker intends to communicate some determinate content, even if the expression might be imprecise, vague, ambiguous, and so on. Taken together, these criteria license the audience to infer what the speaker intends to express.

7.3 Putting it All Together

Let's start by modeling the process of "decoding" the content communicated. That is, I give a simple model to explain *what the audience infers the speaker to express*. The strategy is to be quite conservative with what has to be shared between speaker and audience. I require few things: that the factual common ground is non-defective. Additionally, there needs to be a consensus about the purpose of the conversation. But as to how certain expressions are understood in detail by speaker and audience there can be a significant discrepancy as long as consensus about satisfaction of the purpose of the conversation is not jeopardized. Speaker and audience thus "share some content," but that sharing is not made more difficult by using indeterminate language. Let's develop this idea and make it more precise.

The audience, in decoding the utterance of the speaker, does not always know its propositional content. The model accounts for this by allowing the audience to entertain multiple potential propositions compatible with what they take the speaker to express. For example, a vague expression might semantically determine, depending on the theory of choice, a range of propositions as its content. Correspondingly, the audience has a range of permissible interpretations they might take the speaker to express. In addition, contextual factors might limit the range of permissible interpretations further. The permissible interpretations each assign a proposition to the sentence uttered. This leads to multiple propositions potentially expressed. If these propositions agree on a way to answer to the question under discussion, the audience has reason to infer that this answer is the content the speaker intends to express. The answer multiple propositions agree on is for successful communication, on this account, a single proposition. Both audience and

speaker are able to coordinate on an effect on the common ground. This way, communication can succeed and be explained with a single proposition communicated, but does not require perfect alignment between the speaker's use of indeterminate expressions and the audience's understanding of these expressions.

Questions under discussion in this model are elementary questions asking *whether*, *which*, *who*, and so on. See section 6.3 for more on questions under discussion and chapter 5 for more on how to model questions and their answers. The content of a question is a set of propositions, each of which answers the question in some sense. For example, "who is coming to the party?" intuitively might be answered by naming a relevant person who is awake. Here, it is modeled by a set of propositions, called "question-alternatives", such that each proposition answers, for a particular individual, if they are awake:

$$\begin{aligned} q\text{–alt} = \{&\{w \mid \text{Johnny is coming to the party at } w\}, \\ &\{w \mid \text{Mary is coming to the party at } w\}, \ldots\}. \end{aligned}$$

Each partition cell definitely decides for each question-alternative whether it is the case, generating a partition. Let's denote this partition for a question q with Π_q.

A non-empty proposition φ is an *answer* to a question q if φ rules out at least one cell of the partition induced by the q-alternatives, or more precisely if $\varphi \cap \pi = \emptyset$ for some $\pi \in \Pi_q$. This corresponds to a *weak answer* defined in chapter 5. Different notions of answer can be employed, such that model determines different propositions for what is communicated. The reasons for the choice of notion on an answer can be varied, but for the purposes of presenting the Coarse-Grained Model, a *weak answer* shall suffice. In the above example, "Johnny or Mary are coming to the party" would express a proposition which is an answer to "Who is awake?", since it rules out the case in which both Johnny and Mary are not coming to the party, even though it does not decide for either Johnny or Mary whether they are coming to the party.

For each non-empty proposition φ and question q, there exists a unique *coarsening* $C_q(\varphi)$ (cf. section 5.3): the minimal (by the subset relation) union of partition cells s.t. $\varphi \subseteq C_q(\varphi)$. The context set Δ contains all

possible worlds compatible with the common ground. For present purposes, propositions, answers and question are defined w.r.t. Δ, s.t. $\varphi \subseteq \Delta$ etc. It's usually clear which question is relevant, so I'll omit the subscripts for partitions and coarsenings where possible.

The ways in which a audience interprets what a speaker intends to express are represented as interpretations: functions from sentences to propositions. Not all logically possible interpretations are relevant, instead, the permissible interpretations are constrained by competent use of English, rules of composition, and so forth. But variations in the use of indeterminate terms is allowed: for sentences containing indeterminate expressions, the audience has many permissible ways of interpreting what the speaker expresses. Let's call the set of these interpretations I_h (cf. section 6.4). Some of these interpretations may provide an answer to the question under discussion. Call these interpretations *question-answering*. Usually it is clear which uttered sentence is at issue, so one may also speak loosely of an *interpretation* to be an answer to a question and having a coarsening, when technically it's the proposition assigned by the interpretation to the uttered sentence. So, for each interpretation $i \in I_h$ there is a unique coarsening $C(i)$. The union of all coarsenings of the question-answering (w.r.t. to the question under discussion) interpretations provides what the audience has reason to take the speaker to intend to express. This is the coarse-grained answer to the question under discussion which is compatible with all interpretations that at least provide some answer to the question under discussion. Compare Fig. 7.1 for an illustration. The coarsenings $C(i), C(j)$ and $C(k)$ are shaded grey. The proposition the audience infers the speaker to express is $\varphi = C(i) \cup C(j) \cup C(k)$. Note that this illustration makes use of two ways to represent propositions.[1] One way is by *shading*: For example, coarsening $C(i)$ is shaded gray for worlds w_1 to w_4. This represents the set $\{w_1, \ldots, w_4\}$, i.e. $C(i)$ is true at worlds w_1 to w_4. The other way is by stating the truth values directly. For example, interpretation i assigns *true* (T) to w_2 and w_4, and *false* (F) to all others. This is equivalent to the proposition $\{w_2, w_4\}$.

Couldn't one just take the weakest coarsening as what's expressed? Why take the union? In many cases, there is no unique weakest coarsening. This is a problem for the minimal content constraint account by Schoubye and

[1] I thank an anonymous reviewer for bringing to my attention that the ways of illustrating propositions should be clarified.

	π_1		π_2		π_3		π_4	
	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
i	F	T	F	T	F	F	F	F
j	F	F	F	T	F	F	F	F
k	F	F	F	F	F	F	T	T
φ	T	T	T	T	F	F	T	T

Figure 7.1: Example for permissible interpretations i, j, and $k \in I_h$, i.e. from the perspective of the audience. They give different answers to the question represented by the partition π_1 to π_4.

Stokke (2016) (compare section 4.4.1). Consider the example in Fig. 7.1. There is no unique weakest coarsening. But the union of all coarsenings fits this requirement. For more discussion, see section 9.8.

Distinguish three basic cases: either what the audience infers the speaker to express, let's call it φ, is the context set, i.e. $\varphi = \Delta$. In this case, the question-answering interpretations cannot agree at all on ruling out a partition cell and the utterance fails to answer the question under discussion. Or $\varphi = \emptyset$. In this case, no interpretation makes the utterance true at any possibility in the context set. This is the case if the utterance can in no permissible way be interpreted to be compatible with the common ground, and for the conversation to proceed we'd need to revise some of our knowledge or reject the assertion. I won't deal with this case here. Or φ is a proper subset of the context set Δ, which should be the ordinary case of successful communication. The interpretations agree to rule out some cell of the partition, and thus the audience takes the speaker to contribute to the conversation by adding to the common ground. To summarize:

Thesis *What the audience infers the speaker to express* with a literal assertoric utterance in a context is the union of all coarsenings of the question-answering interpretations the audience deems permissible.

Simply put, taking together all coarsenings of the propositions assigned by the question-answering interpretations yields what the audience takes the speaker to say. For communication to succeed in the Classical Model,

the proposition the audience takes the speaker to express must also be the same single proposition the speaker intends to express. As discussed extensively above, indeterminate expressions pose problems to this success as the audience might have in mind a different proposition or multiple different propositions than the audience ends up entertaining, even though intuitively the communication is successful. A central claim made here is that while the speaker may indeed have different permissible interpretations in mind, the difference does not affect the success of the communication if the permissible interpretations of both speaker and audience "give the same answer" to the question under discussion.

What the speaker intends to express is modeled with the same tools just proposed, but in this case, the permissible interpretations are ways in which the speaker interprets indeterminate expressions, let's call this set I_s.[2] Some of these interpretations answer the question under discussion. What the speaker intends to express is then just the union of coarsenings of these question-answering interpretations.

Thesis *What the speaker intends to express* with a literal assertoric utterance in a context is the union of all coarsenings of the question-answering interpretations the speaker deems permissible.

Consider the example presented in Fig. 7.1. The speaker deems permissible slightly different interpretations compared to the interpretations the audience deems permissible. See Fig. 7.2, which illustrates the speaker's interpretations. Interpretations l and m are permissible interpretations from the perspective of the speaker, thus $l, m \in I_s$. The coarsenings $C(l)$ and $C(m)$ are shaded grey. The proposition the speaker intends to express is $\psi = C(l) \cup C(m)$.

With these two definitions, the central statement of the model is the following.

The Coarse-Grained Model of Communication

(1) Communication is explained according to Classical Pragmatics and Classical Content.

(2) The content of an assertion is *what the speaker intends to express*, and uptake is *what the audience infers the speaker to express.*

[2]Note that this is a rational reconstruction of the speaker's intention as opposed to a model of the speaker *comes up* with which words to use to express what they intend to.

	π_1		π_2		π_3		π_4	
	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
l	F	T	F	F	F	F	T	F
m	F	F	F	T	F	F	F	F
ψ	T	T	T	T	F	F	T	T

Figure 7.2: Continuing the example in Fig. 7.1.

In the example, what the speaker intends to express, ψ, is the same proposition that the audience infers the speaker to express, φ. Communication succeeds. To see this, recall the definitions of Classical Pragmatics and Classical Content. There is a single proposition asserted as described by what the speaker intends to express. Uptake of this proposition is explained as what the audience infers the speaker to express. The common ground is an explicit part of the explanation, and the effect on the common ground of this assertion is clear.[3] The content of an assertion is still truth conditional since it is a proposition and propositions are assumed to be truth conditional.

7.4 Reasons to Adopt the Coarse-Grained Model

Why is the Coarse-Grained Model convincing? I'd like to offer four reasons.

The first reason is simply that the Coarse-Grained Model is conservative w.r.t. the Classical Model of Communication. It does not reject any principle of the Classical Model and thus benefits from and is consistent with a wealth of philosophical theory about communication.

The second reason takes up again and expounds on what has been described in the beginning of this chapter. The account explains why the audience has reason to infer what the speaker intends to express. Potentially, there are many propositions compatible with the permissible interpretations the audience could take the speaker to intend to express. Why is it exactly

[3]Does this update of the context raise any additional rational requirements on the speaker and audience? They should reevaluate their permissible interpretations. At the very least, if an interpretation does not make true an utterance just accepted at any possible world in the context set, it is to be discarded.

that proposition determined by the account that the audience should take the speaker to intend to express? The reason is given by the audience taking the speaker to communicate according to the central maxims of Gricean pragmatics. The maxims thus justify the technical choices made in the account. Consider Grice's Maxim of Relation – simply "be relevant". Being relevant for Grice is to address the subject, focus, or needs of the exchange. Being relevant in the Coarse-Grained Model is answering the question under discussion. Thus only such interpretations which are informative w.r.t. the question under discussion are considered. Next, consider Grice's second Maxim of Quantity—"Do not make your contribution more informative than is required [for the current purposes of the exchange]". Since I allow arbitrary propositions to answer a question, and not only precisely unions of cells of the partition generated by the question, these propositions might provide additional information which is not required by the purpose of the exchange, i.e. information which is not directly relevant to answering the question under discussion. That's why it's the *coarsenings* of the propositions assigned by the interpretations that make up what the speaker intends to express and what the audience infers the speaker to express. The coarsenings are just that part of the informative exchange directly required to answer the question under discussion. Grice's first Maxim of Quantity—"Make your contribution as informative as is required"—provides grounds to take a coarsening as the *minimal* union of partition cells s.t. the proposition is a subset of this union. Any larger union of partition cells is lacking some information w.r.t. the question under discussion. Further, let's take the union of these coarsenings. One cannot leave out the coarsening of some interpretation which answers the question under discussion. This would exclude from the union of coarsenings some partition cell which is part of the coarsening of that interpretation. One would then take the utterance to say more than it says—one would add arbitrary content to the utterance. Under some permissible interpretation, the utterance is true at possible worlds in that cell just arbitrarily excluded. One would thus violate Grice's first Maxim of Quality – "Do not say what you believe to be false". In a straightforward way, then, the Gricean maxims directly translate into explicit features of the present account.

The third reason for why the present account is convincing is that it can be made to give an explication of a context "close enough to being nondefective" in a somewhat similar sense that Stalnaker uses (cf. section 6.2).

The account explains why communication can succeed in a context even though speaker and audience might have different presumptions about the use of language. Communication succeeds if the purpose of the conversation is recognized by speaker and audience to be satisfied in the same way by the utterance. Stalnaker defines a context to be defective when the presumed common ground of speaker and audience misalign. In Stalnaker's technical sense, the context is defective when what the speaker deems to be the context set is not identical to what the audience deems to be the context set. The present account instead assumes that their context sets are indeed identical, but allows for cases in which the interpretations deemed permissible misalign. In the technical sense of this account, that is the case if $I_h \neq I_s$. The context thus can be said to be defective, if in a slightly different meaning from Stalnaker's. Such a context can be "close enough to being nondefective" if it still allows for successful communication. On this account, this is the case if what the audience infers the speaker to express indeed is the the same proposition that the speaker intends to express. And this is the case if the difference between the interpretations the audience deems permissible and the interpretations the speaker deems permissible does not lead to different answers to the question under discussion. More technically, this is the case if the union of all coarsenings of interpretations in I_h which are an answer to the question under discussion is identical to the union of all coarsenings of interpretations in I_s which are an answer to the question under discussion. Less technically, communication succeeds if the different uses of language between speaker and audience are irrelevant for practical purposes. By making precise the Gricean maxims for a certain type of conversation, then, the present account also provides an explication of a context close enough to being nondefective to "not affect the issues that actually arise in the course of the conversation" (Stalnaker, 1978).

The fourth reason for why the present account is convincing is that it provides a method to explain successful communication in the Classical Model for all problematic cases presented in the thesis so far. This will be demonstrated in chapter 8 below.

Chapter 8

Applications

8.1 Introduction

This chapter takes the most prominent examples brought up throughout this thesis and discusses the issues from the point of view of the Coarse-Grained Model. The claim is that the Coarse-Grained Model solves the issues convincingly, and promises to be a good general candidate model to determine communicated content. The chapter is structured by categorizing the issues by the types of indeterminacy they incur, see chapter 3.

Before going into the problematic cases taken from the literature and discussed throughout this thesis, let's get a feel for the Coarse-Grained Model by applying it to the problem discussed in the introduction, repeated here.

Suppose you and your partner are getting ready to go for a walk, and are unsure about whether to take the umbrella. Your partner exclaims:

(3) It's pouring!

Without deliberating much, you'll take the umbrella with you.

How does the Coarse-Grained Model approach to explain this case of successful communication? A first step is to establish a question under discussion. This leaves some leeway to the theorist, but some candidate questions are clearly more plausible than others. For example, "What's the weather like?" is a plausible question, but "Did you see that ludicrous display last night? What's Wenger doing sending Walcott on that early?" is a less plausible candidate. I take this general observation for granted. Let's assume there is a question under discussion theorists might agree on. For the present case, quite plausibly, "Is it raining?" seems to recommend itself. If you know

the answer to that question, you know whether to take an umbrella, simple as that. So the question under discussion divides the possible worlds in the context set into those where it's raining and those where it's not. That this division is unproblematic is further assumed. I deal with the issue of a potentially indeterminate question under discussion in section 9.5. Now, let's consider permissible interpretations for "It's pouring". The introduction mentioned the following: a short drizzle, a hailstorm, your place at the limit of the field of influence of a rain cloud? Of course, other interpretations are possible, especially those where it's prototypically pouring. The important point is that given the context, all these interpretations imply that it is raining. Accordingly, they all answer "Is it raining?" in the positive. How does the Coarse-Grained Model determine the proposition communicated? Consider several possible worlds: there is a short drizzle at w_1, a hailstorm at w_2, we're on the limit of a rain cloud at w_3, and it's cloudy but dry at w_4. It's raining at w_1, w_2, and w_3, but not at w_4. The question under discussion accordingly partitions this set into $\{w_1, w_2, w_3\}$ for *yes* and $\{w_4\}$ for *no*. Consider further the interpretation of (3) that there is a short drizzle (i_a), that there is a hailstorm (i_b), and that we're on the limit of a rain cloud (i_c). Compare Fig. 8.1 for an illustration. The coarsenings of the interpretations are shaded gray.

	yes			*no*
	w_1	w_2	w_3	w_4
i_a	T	F	F	F
i_b	F	T	F	F
i_c	F	F	T	F
φ	T	T	T	F

Figure 8.1: Interpretations for (3).

All coarsenings agree on answer φ to the question under discussion. Yes, it's raining. The union of the coarsenings just is the *yes*-answer to the question under discussion. Thus, the Coarse-Grained Model determines a unique proposition communicated.

Similar strategies can be applied to the other cases discussed in this thesis. The rest of this chapter demonstrates how it's done.

8.2 Indeterminacy in Opaque Contexts

Let's deal with Schoubye and Stokke's example (29), repeated here.

(29) a. Question: The space shuttle must be able to carry 35 tons of cargo, endure extreme temperatures, and be capable of withstanding severe cyclonic dust storms. So, what material for the shuttle is sufficiently strong?

b. Response: Steel is strong enough.

Schoubye and Stokke argue that the question (29a) introduces a question under discussion into the conversation. The question simply is (53):

(53) What material for the shuttle is strong enough for carrying 35 tons of cargo, enduring extreme temperatures, and withstanding severe cyclonic dust storms?

They state that intuitively, what is said by the answer (29b) is something like (54).

(54) Steel is strong enough for carrying 35 tons of cargo, enduring extreme temperatures, and withstanding severe cyclonic dust storms.

In the present discussion, the corresponding proposition to (29b) is the candidate for what the speaker intends to express, what the audience infers the speaker to express, and the effect on the common ground, too. To model the example, Schoubye and Stokke ask us to assume for simplicity that contextually, only steel, aluminium and iron are salient options for what is strong enough to carry 35 tons of cargo and so on. The question then partitions the context set Δ like in Fig. 8.2. S is the proposition that steel is strong enough for carrying 35 tons of cargo and so on, A is the proposition that aluminium is strong enough for carrying 35 tons of cargo and so on, and I is the proposition that iron is strong enough for carrying 35 tons of cargo and so on. A bar over the letter indicates negation.

As argued extensively in section 4.4.1, the account by Schoubye and Stokke does not actually succeed in determining a unique proposition. The

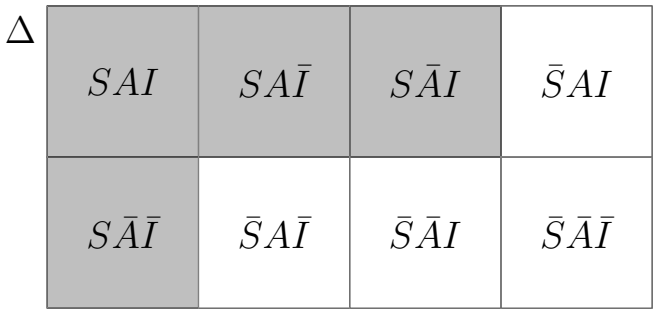

Figure 8.2: The question (29a) generates a partition on the context set Δ. The account determines the shaded area as the communicated content.

Coarse-Grained Model, however, does so: the question under discussion provides the structure the audience is able to infer the speaker's intention by. Even if the audience might be unsure about the precise range of the ways in which the speaker uses the expression "is strong enough", they have reason to infer that the speaker intends to answer the question under discussion in a particular way.

Which interpretations are permissible interpretations, for the audience? Interpreting (29b) as saying that *iron* is strong enough for carrying 35 tons of cargo and so on, for example, is not permissible, since using "steel" to refer to "iron" is not using English competently. Interpreting (29b) as saying that *steel* is strong enough for carrying 35 tons of cargo and so on, however, is permissible. Let's in the following call this interpretation i_{35}. So is interpreting (29b) as saying that steel is strong enough for a lot of other more demanding tasks in addition to carrying 35 tons of cargo an so on (i_{35+}). Interpreting (29b) as saying that steel is strong enough for anything at all (i_{any}), which would amount to a very weak content for the utterance, can be permissible, too. However, i_{any} is not question-answering: that steel is strong enough for anything at all (and doesn't, for example, spontaneously combust immediately after production) seems to be a plausible item of the common ground of a conversation in an engineering context. And if the proposition is already part of the common ground (or entailed by it), it is true at all worlds in the context set. Thus, i_{any} does not answer the question under discussion and is not considered in determining what the speaker intends to express, what the audience infers the speaker to express, or the effect on the common ground.

See Fig. 8.3 for an illustration. Coarsenings are shaded grey. Interpretation i_{any} does not exclude any partition cell and thus does not provide an

answer to the question under discussion. The Coarse-Grained Model licenses the audience to infer as the content of the utterance the union of the coarsenings of all reasonable interpretations w.r.t. to the question under discussion. In this example, this is exactly the union of all S-cells of the partition.

	SI		$S\bar{I}$		$\bar{S}I$		$\bar{S}\bar{I}$	
	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
i_{35}	T	T	T	T	F	F	F	F
i_{35+}	T	F	T	F	F	F	F	F
i_{any}	T	T	T	T	T	T	T	T
φ	T	T	T	T	F	F	F	F

Figure 8.3: Interpretations for (29).

The content of all S-cells is, in this context, just the proposition that steel is strong enough for carrying 35 tons of cargo and so on. The Coarse-Grained Model hence determines a single proposition for what the audience infers the speaker to express. What the speaker intends to express can be determined simply to be the same proposition, and thus also determines the effect on the common ground, rendering the communication successful. This is the case even if the speaker has deems different interpretations permissible, as long as the interpretations agree on an answer to the question under discussion. The interpretations agree, as laid out in chapter 7, if what the speaker intends to express and what the audience infers the speaker to express determine the same proposition. That what the speaker intends to express and what the audience infers the speaker to express determine the same proposition is not a coincidence. Very few assumptions are required: that both speaker and audience use English competently, that they agree (also implicitly) on a question under discussion, and that the context is not defective in a major way. Then taking only the *coarse-grained* part of the permissible interpretation allows for some leeway towards speaker and audience disagreeing on which ways of making the utterance more precise. Relevant is only that part of the interpretations which is just *about* the question under discussion. This reduces the "likelihood" (intuitively speaking, not meant technically)

that what the speaker intends to express and what the audience infers the speaker to express come apart.

One simple way of conceptualizing the idea is this: given a question, utterance, and a context set, the coarse-grained model determines a partial function from interpretations to a proposition. This function is, usually, *not* injective. This means that for any particular proposition the Coarse-Grained Model determines there are, usually, many possible configurations of interpretations mapped onto that proposition. Both speaker and audience deem permissible a particular, possibly different, configuration of interpretations. The communication succeeds if both configurations are mapped onto the same proposition by the function the Coarse-Grained Model provides.

Section 3.3 discussed a host of other examples. This particular type of indeterminacy has been the subject of much debate, as examples abound. Recall the example about Tipper being ready, repeated here:

(52) a. Question: Tomorrow is the final exam! Do you think Tipper is prepared?

b. Response: Tipper is ready.

What does the Coarse-Grained Model have to say about this example? First, it seems that the question is already phrased sufficiently precise. Is Tipper prepared enough to pass the final exam? The question partitions the context set neatly into those worlds where Tipper is prepared and those where he isn't. Note that this is not the same question as whether Tipper actually passes the exam, as there might be all sorts of mishaps preventing Tipper from succeeding, even though he was prepared. Or Tipper might get lucky and pass the exam without being adequately prepared.

Let's assume that the reference to Tipper is not at issue here. The Coarse-Grained Model should determine that "is ready" in the response is used to communicate that Tipper is prepared for the exam. Semantically, there are potentially many interpretations compatible with the response (52b). For example, that Tipper is ready to dance, that he is ready to for the Olympic sprint semi-final, or simply that he is ready for something, etc. One of the permissible interpretations says that Tipper is ready to take the exam tomorrow. Even interpretations might be permissible according to which Tipper is ready for anything the world might throw at him, even though it's unlikely that anyone would entertain such an interpretation. Take the

interpretation of (52b) that Tipper is ready to dance. Absolutely, and also arguably given the common ground (except if we are talking about a dancing exam or some such thing), that interpretation is logically independent from the question-partition. This simply means that the interpretation will assign *true* to some worlds in the *yes*-partition cell as well as the *no*-partition cell, and similarly assigns *false* to some worlds in the *yes*-partition cell as well as the *no*-partition cell. In other words, the interpretation that Tipper is ready to dance will not answer the question under discussion, and thus not count towards what the audience infers the speaker to express. It's similar for most of the other interpretations, except for those which directly answer the question under discussion—a coarse complete answer—or entail such an answer—a complete answer. Arguably, given the context, any such interpretation will answer the question under discussion in the affirmative. The coarsenings of the question-answering interpretations will thus coincide with the *yes*-partition cell, as does their union. The Coarse-Grained Model determines the *yes*-partition cell as what the audience infers the speaker to express.

But what about an interpretation such as that Tipper is ready for absolutely nothing (call it i_n)? If i_n were among the permissible interpretations, it would answer the question under discussion in the negative, since at no world where Tipper is ready for absolutely nothing is he also ready for the exam. The result would be that the coarsening of i_n is to be included in the union of all question-answering permissible interpretations, making the Coarse-Grained Model predict the whole context set as what the audience infers the speaker to express. This would not look good for the Coarse-Grained Model. The Coarse-Grained Model has three options here: the first is denying that i_n is compatible with it's semantic content, on the grounds that it introduces a negation which is not present in the syntactic structure of the sentence uttered. The second adds some contextual factor which limits the range of permissible interpretations. A systematic account of these contextual factors is lacking at this point, however, and would need to be worked out to be convincing. The third option is to employ the common ground. If it is common ground that Tipper is ready for something at all, there is no world in the context set at which Tipper is ready for absolutely nothing. Consequently, i_n assigns *false* to all worlds in the context set, such that it isn't question-answering and does not count towards what the audience

infers the speaker to express.

Further examples discussed include:

(6) a. Steel is strong enough. [to be used in the space shuttle]
b. That lamp is cheap. [relative to other lamps]
c. Mutual knowledge is relevant. [to communication]
d. Biden is too old. [to run again for office]
e. Gregor was merely a bookkeeper. [as opposed to an accountant]
f. The princess is late. [for the party]
g. The king has arrived. [at the palace]
h. Al has finished. [speaking]

These examples can be addressed in a similar fashion to (29) and (52) above. The strategy that emerges is clear: first, we need to make some plausible assumptions about the context. This includes the shared knowledge between speaker and audience, and importantly, the purpose of the conversation. By varying the question under discussion, the utterance can be finagled into the context-appropriate content. But it is not completely arbitrary; so *finagling* is maybe not the best expression. The utterance's content is limited by several factors, including the semantic content of the sentence uttered. After determining these parameters, the Coarse-Grained Model can be applied.

Take example (6b), "That lamp is cheap." The intended but not explicitly stated content compares that particular lamp to some other contextually relevant lamps. For example, a couple goes into a lamp store, with the goal to buy a lamp within their budget. (I recently learned that there actually are lamp stores and that people have specific lamp budgets.) Partner A asks: "Which lamp can we afford?", and the contextually salient objects are then just some selection of lamps in the store. Partner B responds with "That lamp is cheap," pointing to a particular lamp. Partner B does, intuitively, not express that the lamp is cheap looking or cheaply made, but simply gives their opinion on whether that lamp fits their budget, answering the question in the affirmative. Does the Coarse-Grained Model give this result? To simplify, treat "cheap" here not as a vague predicate with borderline cases, but as a binary yet context-dependent predicate. See section 8.4 in this chapter for a discussion of vague predicates. Consider competing

interpretations: that the lamp is looking cheap, that the lamp is cheaply made, that the lamp is cheap, compared to a car, or that the lamp is cheap enough to afford. Arguably, except for the interpretation that the lamp is cheap enough to afford, none of these interpretations answer the question under discussion. For it might or might not be that they cannot afford the lamp, even though it looks cheap, or even though it is cheaply made, or even though it is cheap compared to a car. Given these permissible interpretations, the Coarse-Grained Model determines the *yes*-partition cell as what the audience infers the speaker to express. Similar to the Tipper-related example discussed above, there could be, hypothetically, also a permissible interpretation which answers the question in the negative. This interpretation says something like that the lamp is indeed cheap, but not cheap enough for them to afford. Here just the same reasoning as in the Tipper-related example applies, and the Coarse-Grained Model is not forced to accept this type of interpretation.

8.3 Indeterminacy due to Contextual Indifference

What in the case of a contextually unresolved quantifier domain in (11), repeated here?

(11) Every beer is in the bucket.

The Coarse-Grained Model determines a single proposition communicated. The key step is to identify the purpose of the conversation. What's the question under discussion they are trying to resolve? Tim asks Chet: "are we ready to rage?" which, among other things, is the case if enough beers are chilled when their guests arrive. A plausible assumption is that they'd like there to be at least a beer for every guest they expect to come. The question under discussion could then be

(55) Are there enough beers chilled such that each guest can have at least one?

This question neatly divides all possible worlds in the context set Δ into those where there are enough beers and those where there aren't. Note that this can be the case even if it's not common ground how many guests will show up. It's clear from the example that Chet's answer is in the affirmative. So, his indifference towards which proposition might be expressed, as

Buchanan diagnoses it, indicates that no matter which of the specifications (12a)-(12d), repeated here, one considers, they all answer the question under discussion in the affirmative.

(12) a. Every beer *we bought at the bodega* is in the bucket *in the backyard.*

b. Every beer *we will serve at the party* is in the bucket *decorated in pirate motif.*

c. Every beer *for our guests* is in the bucket *filled with ice.*

d. Every beer *at the apartment* is in the bucket *next to the hot tub.*

For example, it might be common ground that they clearly bought enough beers in the bodega earlier. And if all of those are indeed chilled, enough beers are chilled, obviously.

In the Coarse-Grained Model, we can represent the candidate completions of the original utterance with different interpretations, say i_a, i_b, i_c, and i_d for (12a), (12b), (12c), and (12d) respectively. They all have in common that their coarsenings are exactly the affirmative answer to the question under discussion. Note that this model allows for the interpretations to assign different intensions for these candidates, to the effect that it does not have to be common ground that, say, the beers they bought at the bodega are exactly the same as they plan to serve their guests. Still, a single proposition can be assigned to the utterance in context. For illustrative purposes, I'll add interpretations i_{uni} for "every beer in the universe is in the bucket" and i_{any} for "any beer at all is in the bucket", which both don't provide an answer to the question under discussion and thus aren't question-answering. Compare Fig. 8.4.

The question "are there enough beers chilled such that each guest can have at least one?" is represented by partitioning the set of possible worlds. Coarsenings are shaded grey. Interpretations i_{uni} and i_{any} don't provide answers to the question under discussion: i_{uni} assigns *false* to all worlds in the context set, and the coarsening $C(i_{\text{uni}})$ assigns *true* to all worlds in the context set. Interpretations i_a, i_b, i_c, and i_d assign *true* to some worlds in the *yes*-partition cell, and *false* to all worlds in the *no*-partition cell. As a result, the coarsenings $C(i_a), C(i_b)$, and $C(i_c)$ each coincide with the *yes*-partition cell, and so does their union. The Coarse-Grained Model determines the *yes*-partition cell as what the audience infers the speaker to express.

	yes				*no*	
	w_1	w_2	w_3	w_4	w_5	w_6
i_a	T	F	T	T	F	F
i_b	T	T	F	T	F	F
i_c	F	T	T	T	F	F
i_d	F	T	F	T	F	F
i_{uni}	F	F	F	F	F	F
i_{any}	T	T	T	T	T	F
φ	T	T	T	T	F	F

Figure 8.4: Interpretations for (11).

For these purposes it does not matter, then, which of these propositional candidates the speaker and audience actually entertain, as long as they give the same verdict w.r.t. the question under discussion. The model determines the affirmative answer to the question under discussion as what the speaker intends to express, what the audience infers the speaker to express and the effect on the common ground, thus explaining successful communication within the Classical Model.

8.4 Semantic Indeterminacy

MacFarlane's account was criticized in section 4.3.1 for determining that (27b), repeated here, doesn't communicate factual information.

(27) a. Question. Does recruit Richard meet the height requirements to be assigned to the honor guard?

b. Reponse. Richard is tall.

Does the Coarse-Grained Model fair better? The question (27a) provides a question under discussion, which neatly partitions the context set into those worlds where Richard meets the height requirements and those worlds where Richard doesn't. We can model the permissible interpretations of (27b) with

a range of interpretations assigning truth values as if the predicate "tall" was perfectly precise. This way of modeling is reminiscent of an intensional version of supervaluationism, see section 3.2.

A single interpretation has a similar effect to MacFarlane's hyperplan thresholds: each interpretation gives a precise standard for when to count Richard as tall.

Note that the Coarse-Grained Model does not require interpretations to be consistent across judgments about different individuals or even different utterances about the same individual. It is consistent with the Coarse-Grained Model for interpretations to employ some standard for "tall" for one utterance and quite another standard for a different utterance. For example, if an interpretation judges Richard to be tall if and only if Richard is 1.8m tall or taller does not entail that the interpretation must judge Mike to be tall if and only if Mike is 1.8m tall or taller. If and how the permissible interpretations are retained, modified, or discarded, is not part of the Coarse-Grained Model. Of course, working out a full account of how conversations with vague predicates develop is an interesting problem in its own right, but won't be part of this thesis.

Suppose such precise interpretations $i_1, \ldots, i_n$ assign truth values according to a threshold for "tall" proportional to n. Some of these interpretations answer the question under discussion, namely those which interpret Richard as being as tall or taller than the height requirements. Compare Fig. 8.5 for an illustration. The possible worlds are sorted by Richard's height. Interpretations are sorted by threshold for "tall". Coarsenings are shaded grey. The partition induced by the question under discussion is represented by the vertical line. Interpretations i_m and i_{m+1} don't provide answers to the question under discussion. With interpretations sorted like this, it's easy to see why for gradable adjectives it's quite straightforward for an utterance including vague expression to contribute to the conversation. All the question-answering interpretations agree on a way to answer the question under discussion, i.e. the union of their coarsenings still provides an answer. Note that it's still perfectly possible for the speaker and the audience to have different ideas of how to apply the predicate "tall". Using a vague expression to answer a precise question does not entail that the predicate is now contextually made precise. Rather, it gives reason in this context to not apply "tall" to someone measuring under 1.80 m. No statement about a precise

cut-off point is made. The Coarse-Grained Model thus determines unique truth conditions communicated for (27b) even though it contains a vague expression.

So much for the applications of the Coarse-Grained Model. The next chapter discusses several objections to the model.

		no		*yes*			
	...	$w_{1.78}$	$w_{1.79}$	$w_{1.80}$	$w_{1.81}$	$w_{1.82}$	...
⋮							
i_m		T	T	T	T	T	
i_{m+1}		F	T	T	T	T	
i_{m+2}		F	F	T	T	T	
i_{m+3}		F	F	F	T	T	
i_{m+4}		F	F	F	F	T	
⋮							
φ		F	F	T	T	T	

Figure 8.5: Assignments of truth values by permissible interpretations to (27b).

Chapter 9

Challenges and Objections

9.1 Introduction

This chapter discusses some of the most salient objections to the Coarse-Grained Model. Some of the objections come from anonymous reviewers of papers I submitted; I have indicated wherever that is the case. Of course, not every conceivable challenge is addressed, and some responses are not worked out to the detail they'd deserve. The chapter starts out discussing challenges that the Coarse-Grained Model is either not sufficient or not necessary. Sufficient and necessary for what, exactly? For explaining the intuitive judgments in cases of successful or failing communication. These challenges are closely related to the idea that the Coarse-Grained Model is supposed to be a conceptual analysis of successful communication. This latter claim, however, is not made in this thesis. Nevertheless, the Coarse-Grained Model is intended to explain wide range of cases. One way of making this clearer is the following. The claim that the Coarse-Grained Model is sufficient can be, roughly, represented as the claim that: for all propositions φ and for all cases c, if the Coarse-Grained Model determines φ as the proposition communicated in c, then φ is the proposition intuitively communicated in c. And accordingly, the claim that the Coarse-Grained Model is necessary can be represented as the claim that: for all propositions φ and for all cases c, if φ is the proposition intuitively communicated in c, then the Coarse-Grained Model determines φ as the proposition communicated in c. The challenges then provide descriptions of cases in which the Coarse-Grained Model fails to be sufficient or necessary, respectively. It's also worth point-

ing out that many of the other challenges discussed in this chapter entail that the Coarse-Grained Model is not sufficient or not necessary. But for the sake of exposition, these challenges are listed under different headings, which try to pick out their core characteristics.

9.2 Coarse-Grained Model Not Sufficient

One might object that the model is too generous in diagnosing successful communication. In other words, that the result of the model is not sufficient to determine successful communication. In his discussion of Stalnaker's diagonalization strategy, MacFarlane (2020a) argues that diagonalization renders successful a case of communication which intuitively does not succeed, and one might suspect the Coarse-Grained Model of a similar overgeneration. Consider:

(56) I want you to clean this test tube.

Assume that the context in this case does *not* yield a referent for "this test tube". The speaker doesn't indicate or demonstrate in any way which test tube is referred to. MacFarlane argues that communication fails, and that diagonalization does nothing to explain this failure. Stalnaker's diagonalization strategy won't be discussed in this thesis. The Coarse-Grained Model can provide more explanation as to what the conditions for a communicative failure are. Here, communicative success is relative to the goal or purpose of a conversation. The goal is determined by the context. Consider different questions under discussion: if the question is "what do you want me to do?", then (56) provides some information towards answering that question, because it excludes other contextually salient activities the speaker might want the audience to do besides cleaning test tubes. It is in this context that the Coarse-Grained Model would diagnose successful communication: the assertion provides an answer to the question under discussion. If, however, the purpose of the conversation is much more focused, e.g. the question under discussion is something like "which test tube do you want me to clean?", then asserting (56) does not constitute successful communication on the Coarse-Grained Model. This is because the viable interpretations do not agree on any answer to the question under discussion. To see this, let's assume that the question partitions the set of possible worlds by test tube that the speaker

of (56) wants to be cleaned. Then for each relevant test tube, there is a viable interpretation under which that test tube is referred to with "this test tube". The union of their coarsenings does not provide an answer to the question under discussion, as it just yields the whole context set. Given this purpose, then, the account does not determine a proposition communicated and thus does not render this case of communication successful. One might object, though, that even for the first purpose, represented by "what do you want me to do?", successful communication would determine some concrete thing to do, and intuitively this communication fails as well, contrary to what the Coarse-Grained Model determines. The model does provide another parameter to adjust for this type of demand: the notion of an answer to a question. As presented, the notion of answer is rather weak and "easy" to satisfy for a proposition. There are, however, natural additional requirements that an answer plausibly has to satisfy. Answers could be required to decide between question alternatives, for example, or even could be required to determine a single partition cell. With a more demanding notion of answer chosen as a parameter for the Coarse-Grained Model, it can be easily demonstrated that the account does not determine a proposition communicated, as the viable interpretations do not provide an answer in the stricter sense to the question under discussion. This charge that the model overgenerates can thus be rebutted.

Buchanan and Ian Schiller (2022) argue, too, that a question under discussion-based approach is not sufficient: consider a case in which the audience takes the speaker to answer the question under discussion, but the speaker intends their utterance to not answer the question under discussion, and instead intends for different contextual factors to determine what they said. Buchanan and Schiffer's example is a question under discussion "How are things with your uncle?" and the response "It's hot". Here, the audience takes the speaker to answer the question under discussion, supposing that *it's hot where her uncle lives* is an appropriate answer. The speaker, however, talks about their present situation, intending to say that *it's hot here right now*. This communication seems to fail, but a question under discussion-based account seems to sanction the exchange as successful. On the Coarse-Grained Model, this is again simply a failure of co-operation: the questions under discussion speaker and audience assume do not align. The context is deficient.

9.3 Coarse-Grained Model Not Necessary

Buchanan and Ian Schiller (2022) discuss the account by Schoubye and Stokke (2016), which is a predecessor and inspiration to the Coarse-Grained Model, see section 4.4.1. Yet Buchanan and Schiller's criticism could also be levied against the Coarse-Grained Model. They claim the question under discussion is not necessary to determine the content expressed. In other words, the Coarse-Grained Model might undergenerate. If people state a perfectly normal declarative sentence p and do not answer the question under discussion, whichever it might be, they are still plausibly described as having said *that p*. A question under discussion does not at all help in determining what has been said in this case. It is questionable how *co-operative* this type of answer is, and Grice's Cooperative Principle is at the core of the Gricean analysis of speech act content. See also sections 9.4 and 9.9 below.

Buchanan and Schiffer construct an example in which, they claim, the question under discussion is shared between audience and speaker: "Suppose that Carla has just asked Oscar 'What non-alloyed material is strong enough to build my radio transistor out of?' to which Oscar replies by uttering 'Aluminium is strong enough' in the (mistaken) belief that aluminium is not an alloy." (Buchanan & Ian Schiller, 2022, p. 75, minimally edited). Buchanan and Schiffer argue that intuitively, Oscar says that aluminium is strong enough to build Carla's radio transistor out of. But they claim that this proposition will not count as an answer to the question under discussion. Their reasoning: the question-alternatives will be

$$\begin{aligned} q\text{–alt} = \{\{w \mid & \text{ Copper is strong enough to ... at } w\}, \\ & \{w \mid \text{ Iron is strong enough to ... at } w\}, \dots\} \end{aligned}$$

for all non-alloyed metals. Since aluminium is an alloy, it will not figure in the intensional description of any question-alternative and thus Oscar's utterance doesn't constitute an answer to the question under discussion. Intuitively, it's clear what Oscar said, but the question under discussion-based account fails to diagnose properly. The Coarse-Grained Model, I claim, is able to deal with this case. The range of individuals (in this case copper, iron, etc.) to build question-alternatives from is limited by contextual salience. This requires some cooperation, and Oscar would accordingly need to be aware that aluminium is an alloy for the question-alternatives to be contextually limited

to actually non-alloyed metals. This makes the context, again, deficient: for Oscar, aluminium is a non-alloyed metal, and the question-alternatives he assumes thus include a set of worlds at which aluminium is strong enough to build Carla's radio transmitter out of. Oscar answers the question under discussion as he interprets it, explaining why it appears that he says that aluminium is strong enough. Yet, communication does not succeed, as there is no single unique proposition determined both by Oscar's and by Carla's version of the question under discussion. This is intuitively appealing: there seems to be something wrong with Oscar's answer to Carla's question.

9.4 Question Under Discussion Not Available

A crucial assumption of the model holds that there is a question under discussion implicit in the context. This assumption might be denied. What if there is no such question, and instead completely new information is uttered (and the subject of the conversation potentially changed)? The Coarse-Grained Model cannot determine content without a question under discussion, and it seems that cases changing the subject lack such a question. A possible strategy to address this is to employ *implicit questions* (e.g. Schoubye and Stokke, 2016, although Picazo, 2022 argues against their usefulness). These questions are generated *by* an assertion, such that some utterances can introduce new questions under discussion they immediately answer. In a way, this line of response is congruent with the claim that any sentence has a subject matter (see section 6.3). Subject matters are what the sentence is about. They are modeled as a partition. Consequently, any sentence could be seen as a proposal to introduce the sentence's subject matter as a question under discussion. Of course, in this scenario, the question under discussion is of no help with determining the content of the sentence, as determining the content is a prerequisite to determining the sentence's subject matter. The upshot is that the Coarse-Grained Model is *compatible* with cases in which there does not seem to be an obvious contextually salient question under discussion, such that the objection is rebutted.

9.5 Question Under Discussion Indeterminate

Further problems arise when the existence of the question under discussion is granted, but it itself could contain indeterminate expressions and run the risk of not being part of the common ground, i.e. speaker and hearer deem the question under discussion to partition the context set in a different way. The Coarse-Grained Model would not be applicable. But even in this case, under certain circumstances there might be a unique way in which the indeterminate utterance answers the indeterminate question under discussion. That is, there might still be a single proposition communicated. A detailed answer to these and similar objections will have to wait for further work.

I'll provide some ideas on how to proceed. The first line of defense for the Coarse-Grained Model is calling upon the context: There are plausible cases where, while the question under discussion might be indeterminate, the indeterminacy does not continue into the context set. In other words, there are cases where while it might be indeterminate to which question-partition cell a possible world belongs, there is no such indeterminacy for worlds in the context set. In such a case, there might just be no problem to the Coarse-Grained Model at all. But granted, there are also cases where the question under discussion is indeterminate, even focusing on just the context set. Whether the Coarse-Grained Model can still determine a single proposition for an utterance depends on several factors. For one, it depends on the way the question under discussion is indeterminate. For another, it depends as well on the answers to the question under discussion.

First, what does it mean for a question under discussion to be indeterminate, in the Coarse-Grained Model? The simplest way of modeling an indeterminate question under discussion is to represent the content of a question literal differently. Not as a (set of propositions generating a) partition, but as multiple partitions. Take the simple case of two non-identical partitions which don't include each other (in the sense of subject matter, see section 6.3). Bracket the question how this partition is generated from a set of propositions for now. Suppose it's indeterminate whether Π_a or Π_b best represents the content of the question. The situation might look something like in Fig. 9.1.

In a situation like this, it would make sense to evaluate the result of the Coarse-Grained Model individually, i.e. for each partition separately. It

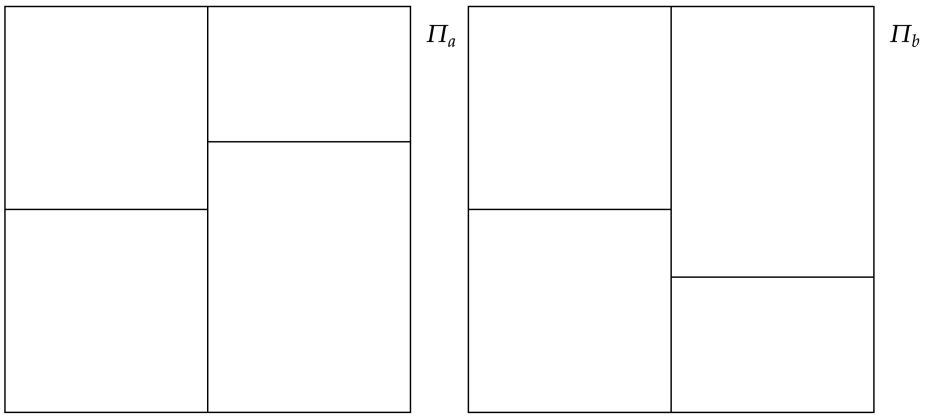

Figure 9.1: Partitions representing an indeterminate question.

might turn out that both processes yield the same proposition. For example, the Coarse-Grained Model might determine φ from both Π_a and Π_b. Compare Fig. 9.2.

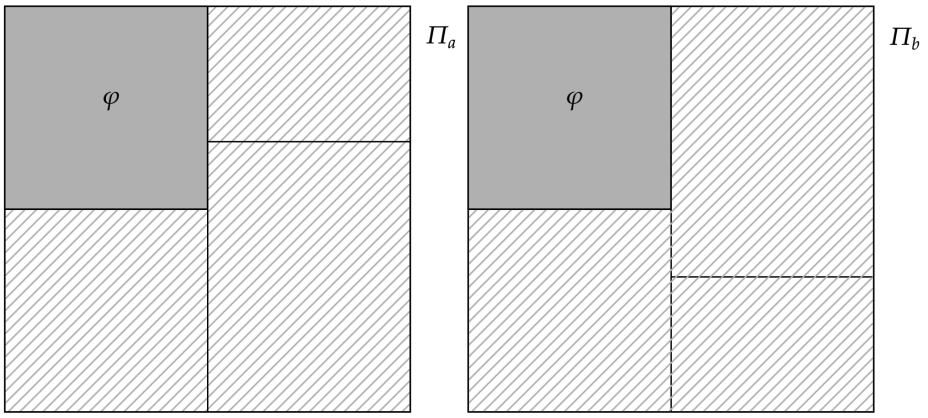

Figure 9.2: The same proposition φ resulting from the Coarse-Grained Model for partitions representing an indeterminate question.

Here, the indeterminacy of the question does not force all answers to be indeterminate. There is no problem of "higher-order indeterminacy" in this case. The Coarse-Grained Model has the resources to resolve the issue. It seems that this move would also work for more involved indeterminacy, but that's a task for future work.

The case is not so simple if we consider a situation in which the question

under discussion exhibits vagueness. For example: "Is Fred bald?" could be represented by a range of different partitions $\Pi_1, \ldots, \Pi_n$. Compare Fig. 9.3.

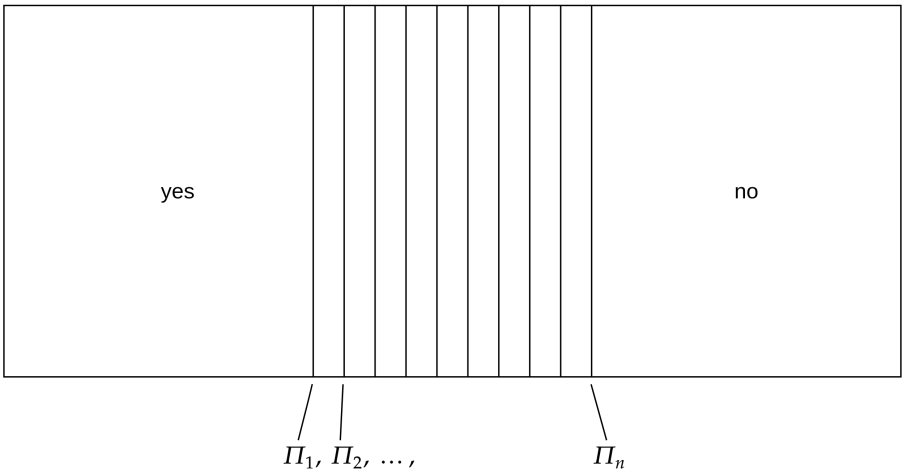

Figure 9.3: Partitions representing a vague question.

Here, the partitions $\Pi_1, \ldots, \Pi_n$ all consist of exactly two cells. Suppose that the Coarse-Grained Model yields a 'yes'-answer for each partition. Then the resulting propositions $\varphi_1, \ldots, \varphi_n$ have the property that $\varphi_1 \subseteq \ldots \subseteq \varphi_n$. Compare Fig. 9.4 and 9.5.

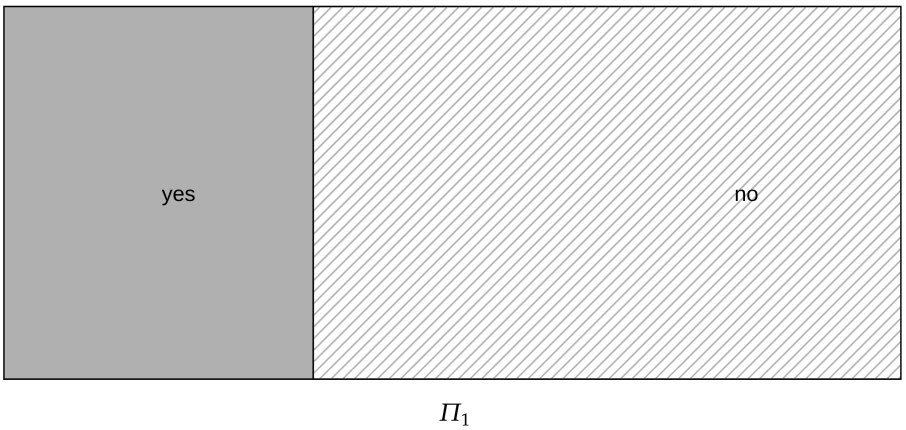

Figure 9.4: Answer φ_1 to the question represented by partition Π_1.

It might be tempting to propose that in this case, one should take the union of the resulting propositions, which is φ_n. After all, this represents

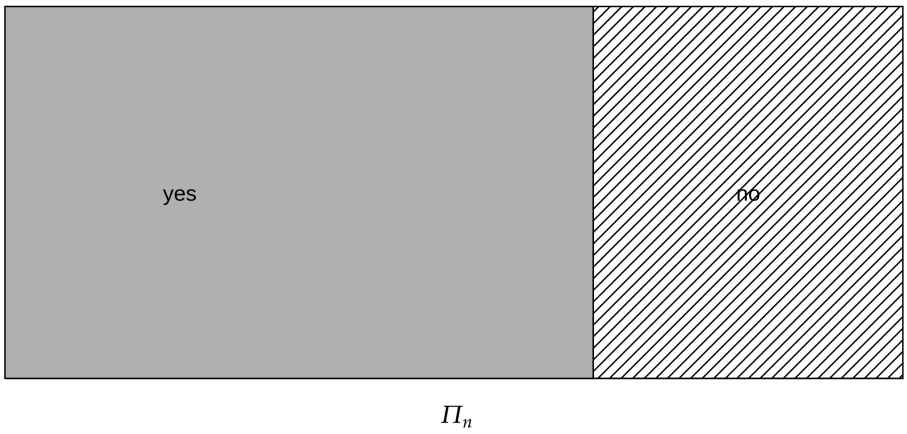

Figure 9.5: Answer φ_n to the question represented by partition Π_n.

the information all answers to the different questions can agree upon. While this seems sensible at first glance, it ignores the problem of higher-order vagueness. In short, this proposal requires that there is a determinate point between being definitely bald and being a borderline case of bald. Otherwise, φ_n could not be determined.

So this approach does not yield a completely satisfying result.

There are more problems with indeterminate questions. In particular, speaker and audience might have different interpretations of the question under discussion. That is, the speaker assumes one way of partitioning the context set, the audience another. I'd venture to guess that this problem could be addressed in much the same way: there are situations in which the Coarse-Grained Model yields a single proposition, despite different interpretations of the question under discussion between subject and audience. These cases would look much like in Fig 9.1 above.

Indeterminate questions certainly post a host of issues to the Coarse-Grained Model. But the Coarse-Grained Model is not without its means of defense; there are plausible ways for future work to address these issues.

9.6 Conversational Exculpature

Conversational Exculpature created by Hoek (2018) seems to display somewhat similar features to the Coarse-Grained Model developed in this paper. In particular, the use of a question under discussion as a partition and the

notion of subtracting information w.r.t. to the question under discussion in Hoek's account seem to run parallel to the concept of coarsening. Central differences between the account by Hoek and the one developed here, like already mentioned in section 4.4.2, lie in their explanatory aims and applicability: Hoek seeks to explain cases of the linguistic phenomenon *pragmatic weakening*. These are cases in which a contextual presupposition appears to be subtracted from the literal content of an utterance to result in the apparent utterance-meaning. Hoek's focus is to determine the contextual presupposition and question under discussion, given a literal content and utterance meaning. Thus, Conversational Exculpature does not address indeterminacy of literal content or utterance meaning. It wouldn't be applicable in its present form, for instance, to cases of vagueness, if vagueness is understood as *semantic* indeterminacy. Conversational Exculpature requires the literal content of a sentence to be determinate, i.e. a known unique proposition.

9.7 Weakest Proposition Good Enough

One might further ask why the Coarse-Grained Model is at all needed in the cases discussed here.[1] Would it not suffice to take as the proposition communicated the weakest proposition that answers the question under discussion? For example, in Fig. 8.3, this might simply be i_{35}, the proposition that steel is strong enough to carry 35 tons of cargo and so on. The intuition about something like this objection is on the right track, as the Coarse-Grained Model aims to find the weakest question-answering proposition compatible with the permissible interpretations. But there are problems with this overly simple approach which the Coarse-Grained Model successfully addresses. The simple approach would not work since in many cases there are multiple propositions which answer the question under discussion, but none of these propositions is the weakest. This is due to the fact that partition cells are only partially ordered by inclusion. One can strengthen this objection by specifying that one takes the weakest proposition *out of the propositions given by the permissible interpretations*. And indeed, this works for Schoubye and Stokke's example as described here. But, depending on which interpretations are taken to be permissible, this does not work generally: in the abstract example given in Fig. 7.1, there is again no weakest

[1] I thank an anonymous reviewer for this objection.

proposition out of the permissible interpretations that answers the question under discussion. This holds in cases where what is communicated with an indeterminate expression is not one of the way of making it precise, but an answer to the question under discussion compatible with each question-answering way of making it precise. For a more concrete example, consider this case of indeterminate reference. For our dinner party, I need to plan whether to prepare a vegan dinner, depending on who's coming. You state:

(57) I invited my friend.

I know your friends, and know that all of them are vegan, which I don't know of, say, your colleagues. But I don't know who in particular you are referring to. Was it Robert, Johnny, or Mary? It appears that communication must fail, because the referent cannot be resolved. But often, statements like this still manage to communicate enough to satisfy the purpose of the conversation. They are "good enough", which is what the Coarse-Grained Model aims to predict. Assume that it's contextually clear that you invite exactly one guest. The permissible interpretations are that you invited Robert, that you invited Johnny, and that you invited Mary. "My friend" refers to one of these three people. Between these permissible interpretations, there is no weakest one. Compare Fig. 9.6. Interpretation i_r says that you invited Robert, i_j that you invited Johnny, and so on. At w_r, you invited Robert, at w_j, you invited Johnny, and so on. The question under discussion here can be summarized as "Should I prepare a vegan dinner?". But the communication might still succeed, as you gave me the information that I indeed have to prepare a vegan dinner. Even though I don't know which referent you actually had in mind, if any.

9.8 Why Take the Union of Coarsenings?

A similar worry asks why it is required to take the union of the coarsenings of all permissible question-answering interpretations.[2] After all, suppose I have several coarsened interpretations, each of which answers the question under discussion in exactly the same way. Why take the union of these coarsenings? Indeed, in the particular case that all coarsened interpretations answer the question under discussion in the same way, the coarsening of any

[2] I thank an anonymous reviewer for this objection.

	yes			*no*	
	w_r	w_j	w_m	w_x	w_y
i_r	T	F	F	F	F
i_j	F	T	F	F	F
i_m	F	F	T	F	F
φ	T	T	T	F	F

Figure 9.6: Assignments of truth values by permissible interpretations to (57).

interpretation is equivalent to the union of all coarsenings. But again, this does not hold generally: in Fig. 7.1, the individual coarsenings are not equivalent to the their union. For a more concrete example, consider the previous case (57), but this time the question under discussion is "Who did you invite for dinner?". Suppose that besides Robert, Johnny, and Mary also some colleagues which you wouldn't call your friends are contextually salient as potentially getting invited. Compare Fig. 9.7. Interpretation i_r says that you invited Robert, i_j that you invited Johnny, and so on. At w_r, you invited Robert, at w_j, you invited Johnny, and so on. The question under discussion here is "Who did you invite for dinner?". Note that this time, the question under discussion partitions the context set in a more fine-grained manner. No longer are the coarsenings of the permissible interpretations identical to their union.

9.9 Information is Lost

Another worry centers around losing information.[3] Suppose the question under discussion is "Do you have any money?" and the answer simply states "I have five bucks". Coarsening the answer to be only about the question under discussion, i.e. effectively answering "yes" instead of the more specific "five bucks", reduces the amount of information communicated. This might not be appropriate. See section 9.4 for a sketch of a response. This objection

[3] I thank an anonymous reviewer for this objection.

	r	j	m	c_1	c_2
	w_r	w_j	w_m	w_{c1}	w_{c2}
i_r	T	F	F	F	F
i_j	F	T	F	F	F
i_m	F	F	T	F	F
φ	T	T	T	F	F

Figure 9.7: Assignments of truth values by permissible interpretations to (57).

can be strengthened, however, to also include some indeterminacy. Consider instead the answer "I have around five bucks", leaving unspecified the exact amount. It's not so clear that all one says here is that one has any money at all, as the answer to the question under discussion would suggest. But it's also not so clear what intuitively *is* said. Would it be felicitous to state if one has only three dollars? Or does one need at least four, but no more than six? Any decision on a precise cut-off seems to be somewhat arbitrary. The Coarse-Grained Model at least provides some principled account about which proposition is expressed.

9.10 Merely Evidential Information

Consider the following exchange, already discussed in section 4.4.1:

(4) a. Question. Who is awake?

b. Response. Kelly has a big day tomorrow.

The partition generated by the question under discussion is assumed to determine for some contextually relevant individuals whether they are awake. But (4b) does not, on the Coarse-Grained Model, answer this question: it is plausibly contextually possible[4] that Kelly has a big day tomorrow but not awake just as much as that Kelly has a big day tomorrow but is awake. It

[4]A proposition is contextually possible if it is compatible with what is common ground.

just seems more likely that she will be asleep already given that she has a big day tomorrow.

Does the response, intuitively, answer the question? Here, intuitions can diverge. One intuition is that no, the response does not provide an answer to the question under discussion. The response might, however, give evidence towards a (partial) answer. Given the response, it becomes more likely that Kelly is not awake. On the account of answer in the Coarse-Grained Model, (4b) would not qualify. Does that mean that nothing was said? No, it could be a case of changing the subject and implicitly introducing a different question under discussion (see section 9.4 above). Alternatively, one could devise a probabilistic (or in any case ampliative, non-monotonic) extension of answers to the Coarse-Grained Model. In such an extension, one could compare the level of confirmation for different answers to the question under discussion the response provides. Then, one considers whether such a response constitutes an answer in a new, weaker and defeasible sense. Unfortunately, developing this idea has to be pushed to future work.

Another intuition is that indeed, the response (4b) does answer the question "Who is awake?" negatively for the case of Kelly. On this understanding, uttering (4b) answers (4a) in a direct sense as a partial answer. But as just established above, on the Coarse-Grained Model the response does not always qualify as an answer. For a response, one could argue that this intuition comes with the assumption that it's contextually clear that Kelly will go to bed early—and consequently be asleep—if she has a big day tomorrow. Maybe that would be a typical thing for her to do, as she is very diligent. In this case, the response answers the question for the partition reduced to the context set. The Coarse-Grained Model gives exactly the right result in this case.

9.11 Implicature

There is a tension between the Coarse-Grained Model and merely implicated or implied content in the sense of Grice (1967). Recall that the interpretations deemed permissible by a rational agent are not just arbitrary interpretations, but they are objectively constrained. One of the constraints, plausibly, is semantic: the literal meaning of the sentence or of expressions of the sentence limits the range of permissible interpretations. For example,

proper names cannot just arbitrarily refer. "Smith weighs 80 kg" cannot be used to say that Louise is German. What is said by an utterance is, on the Coarse-Grained Model, constrained by the permissible interpretations and the question under discussion. What is said by an utterance is then also constrained, to some extent, by the semantic meaning of the uttered sentence. As is well known, and described in some detail in section 2.4, Grice divides speaker meaning into *at least* what is literally said and what is implicated. This distinction can of course be divided up further, but it's enough to state the challenge to the Coarse-Grained Model. What is implicated goes beyond what the sentence literally said. The speaker assumes that all sorts of contextual factors enable the audience to infer what the speaker implicates.

The Coarse-Grained Model foremost determines *what is communicated*. During the course of the thesis, I employ the phrases "what is said", "what is expressed", and so on, as well as "what the audience infers the speaker to express". Thus, naturally, it seems that the Coarse-Grained Model determines *what is literally said*, and how literal communication succeeds. However, the permissible interpretations merely constrain and do not uniquely determine a proposition communicated. Instead, the question under discussion together with other contextual clues put the audience in a position to infer what's communicated. From this point of view, when what is communicated differs substantially from the literal meaning of a sentence, it seems that what the Coarse-Grained Model determines is rather *what is implicated* by an utterance. Both types of communicated content are relevantly different, but it's unclear which of the two the Coarse-Grained Model actually addresses, if any. Or so goes the challenge.

Now, the Coarse-Grained Model is in large part motivated by implementing Grice's Cooperative Principle and the corresponding maxims (see section 7.2). These are principles intended to aid in determining *what is implicated* by an utterance. The question under discussion is supposed to model the informational purpose of a conversation. This suggests that the Coarse-Grained Model determines *what is implicated* first.

I think its instructional to take a look at the examples discussed in this thesis with respect to Grice's distinction. Consider again Schoubye and Stokke's case (29):

(29) (29a) Question. The space shuttle must be able to carry 35 tons of cargo, endure extreme temperatures, and be capable of with-

standing severe cyclonic dust storms. So, what material for the shuttle is sufficiently strong?

(29b) Response. Steel is strong enough.

It's appears that what is literally said with (29b) is *not* that steel is strong enough to carry 35 tons of cargo and so on. But, it also does not seem to be *conversationally implicated* that steel is strong enough to carry 35 tons of cargo an so on. Section 2.4 describes a test, and the canceling seems to be infelicitous. It is not quite the same with "Every beer is in the bucket" and "Richard is tall", as the reader might verify themselves. So it's indeed not quite clear which of these categories apply to the Coarse-Grained Model.

To stress again the tension mentioned above, the class of implicatures include cases of *irony*. Grice's example: A says that 'X is a *fine* friend', even though it's common knowledge that A recently was betrayed by X. Grice pragmatically analyzes this case by describing A as *flouting*, i.e. intentionally and obviously violating, the Maxim of Quality—"say what you believe is true". For the Coarse-Grained Model, this poses a problem, as it seems that saying that X is *not* a fine friend at all is not expressed by any permissible interpretation. Consequently, there is no answer to a conceivable question under discussion—for the present purposes, let's assume that it's "is X a good friend to A", compatible both with what is common ground and the permissible interpretations of A's utterance. If no permissible interpretation makes the utterance true at any world in the context set, the "fault" could lie with either the range of permissible interpretations or the common ground. In the case of irony, the fault seems to lie with the range of permissible interpretations. The Coarse-Grained Model would thus need an extension in which to decide in a principled manner that irony occurs and then systematically update the set of permissible interpretations. I don't see a simpler solution for the Coarse-Grained Model at present.

What about what is literally said? Cannot the Coarse-Grained Model aid at all in determining what is literally said? I'd actually argue that it can, in some cases. Those are the cases discussed in section 9.4 above, which require an impromptu *implicit* question under discussion, in lieu of a contextually salient one. An implicit question represents the sentence's subject matter, as far as the subject matter is determinate, and as such is heavily constrained

by its semantic properties. Arguably, any proposition (i) derived from the permissible interpretations of an utterance and (ii) answering a question under discussion identical to the sentence's subject matter is reasonably close the sentence's literal meaning. In these cases, then, the Coarse-Grained Model plausibly adds to determining of what is literally said.

Chapter 10

Conclusion

This thesis argues that for many cases of successful communication, there is a way to determine a unique proposition communicated. With this proposition, the principles of the Classical Model can be saved; and communication can be explained conservatively. This result is surprising, since seemingly all other approaches so far either reject some principle of the Classical Model or are unsuccessful at systematically determining a unique proposition communicated. Thus this thesis provides a novel perspective on a problem which has not been successfully addressed so far.

The model developed is dubbed the *Coarse-Grained Model*, as it takes a coarse-grained approach to determining the proposition communicated. The model saves the Classical Model because it adds to the Classical Model a method of determining a unique proposition. Additionally, all principles of the Classical Model in explaining communicative success are upheld. The proposition is determined by an interplay of contextual common ground, interpretations deemed permissible, and contextual purpose in form of a question under discussion. While the interactions of context and content of an utterance seem dynamic, the explanandum of the Coarse-Grained Model is not a dynamically unfolding conversation. Rather, the Coarse-Grained Model looks at a sort-of synchronic snapshot of a conversation.

This thesis claims that the Coarse-Grained Model convincingly solves the issues of the cases of indeterminacy in communication. The thesis thus contributes to the recent literature about the truth-conditional content of utterances, which, as discussed, most other accounts stand skeptical towards. The account of *answers* in this thesis is novel, seems to naturally follow from

basic assumptions about the content of an interrogative literal, and might be useful to other theories down the road.

The thesis is an extended argument for the main claim that even in cases of utterance-indeterminacy, it's reasonable to explain successful communication by appeal to a single set of truth-conditions. To recap, the claim is already clearly stated in the introduction. The Classical Model of communication, which the thesis sets out to defend, is then developed and placed in historical context. Different types of indeterminacy are introduced, which all seem to have the same consequence, namely that there is no unique proposition communicated determinable. The thesis then looks in great detail at previous philosophical accounts addressing some sort of indeterminacy in communication. But all these accounts either reject the Classical Model in some form or other, or do not, the thesis argues, succeed in determining a single proposition. But fret not; the Coarse-Grained Model saves the day. Or that, at least, is claimed. To develop the Coarse-Grained Model, it is instructive to first introduce the technical prerequisite concepts the model employs. In particular, a whole chapter is devoted to explicating what answering a question amounts to. This is of great import for the model, since the model revolves around a question under discussion. Finally, the model is stated, and applied to the central cases discussed so far in the thesis. These cases led most other authors to abandon the Classical Model. The Coarse-Grained Model instead embraces it and finds unique propositions communicated in all cases discussed. Finally, the most salient challenges are successfully rebutted, save those which require an extension of the model.

But there are also some less obvious challenges or at least discussion-worthy issues, which I will briefly describe in this conclusion. By divorcing the possible worlds from use of language, the Coarse-Grained Model gives up central assumptions of (read: is more general than) 2-D semantics. In the 2-D semantics developed by, for example, Stalnaker (1978), intensions of expressions can vary with the world they are uttered in. The *context of utterance* can vary. For example, at world w, "water" refers to H_2O, whereas at another world v, "water" refers to a chemical compound XYZ. It is a fact *of that world* w that "water" refers to H_2O. On the Coarse-Grained Model, there is no fact of world w which determines the meaning of "water", at least it is not an obvious part of the model. When making a claim about language, like "water means H_2O", that claim is a claim about facts of the world. But on

the Coarse-Grained Model, this claim seems more to say something about which interpretations are to be deemed permissible. Thus it is not quite clear what the content of such a claim would be—but it seems to be non truth-conditional, in violation of the Classical Model! Indeed, the Coarse-Grained Model would have to do some substantial work to address this issue. One route to go is to claim that worlds *partially* determine the meaning of expressions, but not of the whole language. Whatever can objectively be said to constrain permissible interpretations could then, for example, be "part of the fabric" of the world in which the conversation takes place. But variations in truth-conditions for indeterminate expressions are still allowed, even at the same world.

A different challenge comes from the observation that the content the Coarse-Grained Model assigns to an utterance does not seem to be compositionally constructed from the syntax of the sentence and the meaning of the sentence's parts. This, however, is one of the main advantages of truth-conditional analyses of content. Here, the response is simply that the Coarse-Grained Model most likely does not determine what is literally said, but merely what is implicated (see section 9.11). What is implicated by an utterance, however, is not generally taken to be compositionally constructed from the meaning of the uttered sentence's parts, so there is would be no problem here.

In any case, this shall be all that is discussed in this thesis, and as shown a lot of further work is required.

Appendix A

Proofs

- Observation in section 4.4.2: a partial proposition $r = \langle t, f \rangle$ is about S iff there is *no* S-cell containing both t- and f-worlds iff there is a proposition p which is a union of S-cells and there is a proposition q such that $r = p|q$.

Proof. If a proposition p is a union of $S-$cells c_i, i.e. $p = \bigcup c_i$, then so is $\neg p = \bigcup c_j$. Since S is a partition, $c_i \cap c_j = \emptyset$ for all S-cells. So there is no S-cell c such that $c \cap t \neq \emptyset$ *and* $c \cap f \neq \emptyset$. In other words, each cell is fully contained in either t, f, or their intersection is empty. So no S-cell is containing both t- and f-worlds. On the other hand, if there is no S-cell containing both t- and f-worlds, then it's straightforward to construct a proposition p which is a union of S-cells and a proposition q such that $\langle t, f \rangle = \langle p \cap q, \neg p \cap q \rangle$. Take $q = t \cup f$. By assumption, for each cell of S, $c \cap t = \emptyset$ or $c \cap f = \emptyset$ (or both). Construct p such that $p = \{\bigcup c \in S | c \cap f = \emptyset\}$. Then, $p \cap q = t$ and $\neg p \cap q = f$. $\square$

- Relationships between types of answers, Fig. 5.11 in section 5.3.

Proof. $CA \Rightarrow PPA$ or NPA: $\varphi \subseteq \pi$ for some $\pi \in \Pi$, so $\varphi \subseteq a$ or $\varphi \subseteq a^c$ for each $a \in q$–alt, since the partition is generated from q-alternatives. If $\varphi \subseteq a$, then φ is a PPA. If $\varphi \subseteq a^c$, then φ is an NPA. $\square$

$PPA \Rightarrow PA$, $NPA \Rightarrow PA$ by definition.

$PA \Rightarrow WA$: Let q–alt $= \{a_1, \ldots, a_n, b_1, \ldots, b_m\}$, with $n \geq 1, m \geq 0$. Let $\varphi \subseteq a_1 \cap \ldots \cap a_n$, so φ is a partial answer. Let $\chi = a_1^c \cap \ldots \cap a_n^c$. Clearly, $\varphi \subseteq \chi^c$, which implies $\varphi \cap \chi = \emptyset$. Let $\psi = \chi \cap b_1 \cap \ldots \cap b_m$. Note that ψ is

a partition cell. Clearly, $\psi \subseteq \chi$, which implies $\chi^c \subseteq \psi^c$. Then $\varphi \subseteq \chi^c \subseteq \psi^c$, which implies $\varphi \subseteq \psi^c$. Since ψ is a partition cell, φ is a weak answer. □

- Remark in section 5.3: for each answer φ and a question q there corresponds a unique *coarsening* $C(\varphi)$: the minimal union of $q-$partition cells s.t. $\varphi \subseteq C(\varphi)$.

Proof. Construct the coarsening, then show that it is unique. Each element of Ω is in exactly one partition cell of Π. For each $w \in \varphi$, call π_w that partition cell s.t. $w \in \pi_w$. The coarsening $C(\varphi)$ is the union of these partition cells, i.e. $C(\varphi) = \cup \pi_w$ for all $w \in \varphi$. Obviously, $\varphi \subseteq C(\varphi)$. If $C(\varphi)$ is not unique, then there is a coarsening χ such that either (i) there is a $\pi \subseteq \chi, \pi \in \Pi$ with $\pi \not\subseteq C(\varphi)$ or (ii) there is a $\pi \subseteq C(\varphi), \pi \in \Pi$ with $\pi \not\subseteq \chi$. (i) For all $w \in \pi$, where $\pi \not\subseteq C(\varphi)$, $w \notin \varphi$. But then, χ cannot be minimal, i.e. not a coarsening. (ii) There is a $w \in \pi$, where $\pi \not\subseteq \chi$, but $w \in \varphi$. Then, $\varphi \not\subseteq \chi$, i.e. χ is not a coarsening. □

- Remark in section 5.3: if a proposition φ is a complete (positive partial, negative partial, partial, weak) answer to a question q, then the coarsening $C(\varphi)$ is a *coarse* complete (positive partial, negative partial, partial, weak) answer to the question q.

Proof for complete answers. Others directly analogous. Let $\varphi \subset \pi, \pi \in \Pi$. It's coarsening $C(\varphi)$ is π. Clearly, $\pi \subseteq \pi$, hence π is a coarse complete answer. □

- Relationships between types of answers, Fig. 5.13 in section 5.3.

Proof. $CCA \Rightarrow CA, CPPA \Rightarrow PPA, CNPA \Rightarrow CPA, CPA \Rightarrow PA, CWA \Rightarrow WA$ by definition.

$CCA \Rightarrow CPPA$ or $CNPA$: exactly analogous to $CA \Rightarrow PPA$ or NPA above, except $\varphi = \pi$.

$CPPA \Rightarrow CPA, CNPA \Rightarrow CPA$ by definition.

$CPA \Rightarrow CWA$: exactly analogous to $PA \Rightarrow WA$ above, except $\varphi = a_1 \cap \ldots \cap a_n$, so φ is coarse. □

Bibliography

Abreu Zavaleta, M. (2021). Communication and indifference. *Mind & Language*, *36*(1), 81–107.

Aristotle. (1985). *The Complete Works of Aristotle, Volume One* (J. Barnes, Ed.). Princeton University Press.

Aristotle. (2001). *The Basic Works of Aristotle* (R. McKeon, Ed.). Modern Library.

Armstrong, J. (2023). Truth and Imprecision. *Analytic Philosophy*.

Ashworth, E. J. (1981). "Do Words Signify Ideas or Things?" The Scholastic Sources of Locke's Theory of Language. *Journal of the History of Philosophy*, *19*(3), 299–326.

Austin, J. L. (1961). A Plea for Excuses'. In J. O. Urmson & G. J. Warnock (Eds.), *Philosophical Papers*. Clarendon Press.

Austin, J. L. (1962). *How to Do Things with Words* (M. Sbisá & J. O. Urmson, Eds.). Clarendon Press.

Bach, K. (1994). Conversational Impliciture. *Mind & Language*, *9*(2), 124–162.

Bach, K., & Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.

Baldwin, T. (2006). Philosophy of Language in the Twentieth Century. In B. C. Smith & E. Lepore (Eds.), *The Oxford Handbook of Philosophy of Language* (pp. 60–99). Oxford University Press.

Barker, C. (2002). The Dynamics of Vagueness. *Linguistics and Philosophy*, *25*(1), 1–36.

Barwise, J., & Perry, J. (1981). Situations and Attitudes. *The Journal of Philosophy*, *78*(11), 668–691.

Barwise, J., & Perry, J. (1983). *Situations and attitudes* (Vol. 181). MIT Press Cambridge, MA.

Beaney, M. (2006). Wittgenstein on Language: From Simples to Samples. In B. C. Smith (Ed.), *The Oxford Handbook of Philosophy of Language*. Oxford University Press.

Bezuidenhout, A. (2002). Truth–Conditional Pragmatics. *Philosophical Perspectives*, *16*, 105–134.

Bezuidenhout, A. (2006, January). Semantics–Pragmatics Boundary. In K. Brown (Ed.), *Encyclopedia of Language & Linguistics (Second Edition)* (pp. 183–191). Elsevier.

Black, M. (1948). Introductory Note. *The Philosophical Review*, *57*(3), 207–208.

Boethius & Meiser, K. (1987). *Commentaries on Aristotle's De interpretatione: Anicii Manlii Severini Boetii commentarii in librum Peri hermeneias*. Garland.

Bowker, M. (2019). Saying a bundle: Meaning, intention, and underdetermination. *Synthese*, *196*(10), 4229–4252.

Bowker, M. (2022). Ineliminable underdetermination and context-shifting arguments. *Inquiry*, *65*(2), 215–236.

Braun, D., & Sider, T. (2007). Vague, So Untrue. *Noûs*, *41*(2), 133–156.

Buchanan, R. (2010). A Puzzle about Meaning and Communication. *Noûs*, *44*(2), 340–371.

Buchanan, R. (2012). Is Belief a Propositional Attitude? *Philosophers' Imprint*, *12*, 1–20.

Buchanan, R., & Ian Schiller, H. (2022). Pragmatic Particularism. *Philosophy and Phenomenological Research*, *105*(1), 62–78.

Burgersdijk, F. P. (1632). *Institutionum logicarum*. Leiden.

Caie, M. (2018). Semantic Indecision. *Philosophical Perspectives*, *32*(1), 108–143.

Cappelen, H., & Lepore, E. (1997). On an Alleged Connection Between Indirect Speech and the Theory of Meaning. *Mind and Language*, *12*(3-4), 278–296.

Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.

Cappelen, H., & Lepore, E. (2006). Shared Content. In E. Lepore & B. C. Smith (Eds.), *The Oxford Handbook to the Philosophy of Language* (pp. 1020–1055). Oxford University Press.

Carnap, R. (1947). *Meaning and Necessity*. University of Chicago Press.

Carnap, R. (1963). Replies and Systematic Expositions. In P. A. Schilpp (Ed.), *The Philosophy of Rudolf Carnap* (pp. 859–1013, Vol. 11). Open Court.

Carnap, R., & Bar-Hillel, Y. (1954). An Outline of a Theory of Semantic Information. *Journal of Symbolic Logic*, *19*(3), 230–232.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, R. (2013). Implicature, explicature, and truth-theoretic semantics. In M. Ezcurdia & R. J. Stainton (Eds.), *The Semantics–Pragmatics Boundary in Philosophy* (p. 261). Broadview Press.

Church, A. (1951). *A Formulation of the Logic of Sense and Denotation*. Liberal Ars Press.

Church, A. (1948). Max Black. A translation of Frege's Ueber Sinn und Bedeutung. Introductory note. *Journal of Symbolic Logic*, *13*(3), 152–153.

Ciardelli, I. (2016). Lifting conditionals to inquisitive semantics. *Semantics and Linguistic Theory*, 732–752.

Ciardelli, I. (2022). *Inquisitive Logic: Consequence and Inference in the Realm of Questions* (Vol. 60). Springer International Publishing.

Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2013). Inquisitive Semantics: A New Notion of Meaning. *Language and Linguistics Compass*, (9), 459–476.

Ciardelli, I., Groenendijk, J., Roelofsen, F., Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2018, November). *Inquisitive Semantics*. Oxford University Press.

Ciardelli, I., Roelofsen, F., & Theiler, N. (2017). Composing Alternatives. *Linguistics and Philosophy*, *40*(1), 1–36.

Couto, S. d. (1606). *In universam Dialecticam Aristotelis Stagiritae*.

Cross, C., & Roelofsen, F. (2024). Questions. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2024). Metaphysics Research Lab, Stanford University.

Davidson, D. (1967). Truth and Meaning. *Synthese*, *17*(1), 304–323.

Davidson, D. (1973). Radical Interpretation. *Dialectica*, *27*(1), 314–328.

Davidson, D. (1986). A Nice Derangement of Epitaphs. In E. Lepore (Ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (pp. 433–446). Blackwell.

Donnellan, K. S. (1970). Proper names and identifying descriptions. *Synthese*, *21*(3-4), 335–358.

Dretske, F. I. (1981). *Knowledge and the Flow of Information*. MIT Press.

Dummett, M. (1973). *Frege: Philosophy of Language*. Duckworth.

Dummett, M. (1996). Language and Communication. In M. Dummett (Ed.), *The Seas of Language*. Oxford University Press.

Dunbar, G. (2001). Towards a cognitive analysis of polysemy, ambiguity, and vagueness. *Cognitive Linguistics*, *12*(1), 1–14.

Eco, U. (1984). Signification and Denotation from Boethius to Ockham. *Franciscan Studies*, *44*, 1–29.

Fine, K. (1975). Vagueness, Truth and Logic. *Synthese*, *30*(3), 265–300.

Fine, K. (2012). Counterfactuals without possible worlds. *The Journal of Philosophy*, *109*(3), 221–246.

Fine, K. (2016). Angellic Content. *Journal of Philosophical Logic*, *45*(2), 199–226.

Frege, G. (1892). Uber Sinn Und Bedeutung. *Zeitschrift für Philosophie Und Philosophische Kritik*, *100*(1), 25–50.

Frege, G. (1948). Sense and Reference. *The Philosophical Review*, *57*(3), 209–230.

Frege, G. (1952). On Concept and Object. In P. T. Geach & M. Black (Eds.), *Translations From the Philosophical Writings of Gottlob Frege*. Blackwell.

Frege, G. (1953). *The Foundations of Arithmetic* (J. Austin, Trans.; 2nd ed.). Northwestern University Press.

Frege, G. (1977). *Logical Investigations* (P. T. Geach, Ed.). Blackwell.

Gauker, C. (1992). The Lockean Theory of Communication. *Noûs*, *26*(3), 303–324.

Gauker, C. (2002). *Words without Meaning*. MIT Press.

Gauker, C. (2012). Semantics and Pragmatics. In G. Russell & D. G. Fara (Eds.), *Routledge Companion to the Philosophy of Language*. Routledge.

Geach, P. T. (1969). The Perils of Pauline. *Review of Metaphysics*, *23*(2), 287–300.

Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.

Goodman, N. (1961). About. *Mind*, *70*, 1.

Graff, D. (2000). Shifting sands: An interest relative theory of vagueness. *Philosophical Topics*, *28*(1), 45–81.

Grice, H. P. (1968). Utterer's Meaning, Sentence-Meaning, and Word-Meaning. *Foundations of Language*, *4*(3), 225–242.

Grice, H. P. (1975). Logic and Conversation. In D. Davidson & G. Harman (Eds.), *The Logic of Grammar* (pp. 64–75). Dickenson.

Grice, H. P., & White, A. R. (1961). Symposium: The Causal Theory of Perception. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *35*, 121–168.

Grice, H. P. (1957). Meaning. *Philosophical Review*, *66*(3), 377–388.

Grice, H. P. (1967). Logic and conversation. In P. Grice (Ed.), *Studies in the Way of Words* (pp. 41–58, Vol. 3). Harvard University Press.

Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.

Grindrod, J., & Borg, E. (2019). Questions Under Discussion and the Semantics/Pragmatics Divide. *Philosophical Quarterly*, *69*(275), 418–426.

Groenendijk, J. (1999). The Logic of Interrogation: Classical Version. In J. Groenendijk (Ed.), *Proceedings From Semantics and Linguistic Theory 9* (pp. 109–126). Cornell University.

Groenendijk, J., & Roelofsen, F. (2009). *Inquisitive Semantics and Pragmatics*.

Groenendijk, J., & Stokhof, M. (1984a). On the Semantics of Questions and the Pragmatics of Answers. In F. Landman & F. Veltman (Eds.), *Varieties of Formal Semantics: Proceedings of the Fourth Amsterdam Colloquium* (pp. 143–170). Foris.

Groenendijk, J., & Stokhof, M. (1984b). *Studies on the semantics of questions and the pragmatics of answers* [Doctoral dissertation].

Groenendijk, J., & Stokhof, M. (1990). Dynamic montague grammar. In L. Pólos & L. Kálmán (Eds.), *Proceedings of the Second Symposion on Logic and Language*. Budapest: Eötvös Loránd Press.

Groenendijk, J., & Stokhof, M. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy*, *14*(1), 39–100.

Groenendijk, J., & Stokhof, M. (1997). Questions. In J. v. Benthem & A. t. Meulen (Eds.), *Handbook of Logic and Language*. Elsevier.

Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language*, *10*(1), 41–53.

Heck, R., & May, R. (2006). Frege's Contribution to Philosophy of Language. In B. C. Smith & E. Lepore (Eds.), *The Oxford Handbook of Philosophy of Language* (pp. 3–39). Oxford University Press.

Heim, I. (1982). *The Semantics of definite and Indefinite Noun Phrases*. University of Massachusetts Amherst.

Heim, I., Bäuerle, R., & Schwarze, C. (1983). File change semantics and the familiarity theory of definiteness. *Semantics Critical Concepts in Linguistics*, 108–135.

Hodgson, T. (2018). Meaning Underdetermines What Is Said, Therefore Utterances Express Many Propositions. *Dialectica*, *72*(2), 165–189.

Hoek, D. (2018). Conversational Exculpature. *Philosophical Review*, *127*(2), 151–196.

Kamp, H. (1981a). The Paradox of the Heap. *Journal of Symbolic Logic*, *49*(3), 225–277.

Kamp, H. (1981b). A Theory of Truth and Semantic Representation. In P. Portner & B. H. Partee (Eds.), *Formal Semantics - the Essential Readings* (pp. 189–222). Blackwell.

Kamp, H., & Reyle, U. (1993). *From logic to discourse*. Dordrecht: Kluwer.

Kaplan, D. (1964). *Foundations of Intensional Logic*. Monograph Collection.

Kaplan, D. (1968). Quantifying In. *Synthese*, *19*(1-2), 178–214.

Kaplan, D. (1989). Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and other Indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes From Kaplan* (pp. 481–563). Oxford University Press.

Karttunen, L. (1977). Syntax and Semantics of Questions. *Linguistics and Philosophy*, *1*(1), 3–44.

Kasabova, A., & Marinov, V. (2016). Aristotle on verbal communication: The first chapters of De Interpretatione. *European Journal for the Philosophy of Communication (Empedocles)*, *7*(2), 239–253.

Khoo, J., & Knobe, J. (2016). Moral Disagreement and Moral Semantics. *Noûs*, *52*, 109–143.

King, J. C. (2014). The Metasemantics of Contextual Sensitivity. In B. Sherman & A. Burgess (Eds.), *Metasemantics: New Essays on the Foundations of Meaning* (pp. 97–118). Oxford University Press.

Kleene, S. C. (1952). *Introduction to Metamathematics*. P. Noordhoff N.V.

Kretzmann, N. (1968). The Main Thesis of Locke's Semantic Theory. *The Philosophical Review*, *77*(2), 175–196.

Kretzmann, N. (1974). Aristotle on Spoken Sound Significant by Convention. In J. Corcoran (Ed.), *Ancient Logic and Its Modern Interpretations: Proceedings of the Buffalo Symposium on Modernist Interpretations of Ancient Logic, 21 and 22 April, 1972* (pp. 3–21). Springer Netherlands.

Kripke, S. (1959). A Completeness Theorem in Modal Logic. *The Journal of Symbolic Logic*, *24*(1), 1–14.

Kripke, S. (1963a). Semantical Analysis of Modal Logic I. Normal Propositional Calculi. *Zeitschrift fur mathematische Logik und Grundlagen der Mathematik*, *9*(5-6), 67–96.

Kripke, S. (1963b). Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, *16*, 83–94.

Kripke, S. (1972). *Naming and Necessity: Lectures Given to the Princeton University Philosophy Colloquium*. Cambridge, MA, USA: Harvard University Press.

Kripke, S. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Harvard University Press.

Kvanvig, J. L. (2011). Against Pragmatic Encroachment. *Logos and Episteme*, *2*(1), 77–85.

Landesman, C. (1976). Locke's Theory of Meaning. *Journal of the History of Philosophy*, *14*(1), 23–35.

Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review*, *123*(2), 131–171.

Leitgeb, H. (2017). *The Stability of Belief*. Oxford University Press.

Leitgeb, H. (2022). Ramsification and Semantic Indeterminacy. *Review of Symbolic Logic*, *16*(3), 900–950.

Leitgeb, H., & Carus, A. (2023). Rudolf Carnap. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2023). Metaphysics Research Lab, Stanford University.

Levinson, S. (1995). Three levels of meaning. In *Grammar and meaning: Essays in honour of Sir John Lyons* (pp. 90–115). Cambridge University Press.

Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT press.

Lewis, D. K. (1970). General semantics. *Synthese*, *22*(1-2), 18–67.

Lewis, D. K. (1975). Languages and Language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (pp. 3–35). University of Minnesota Press.

Lewis, D. K. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, *74*(4), 549–567.

Lewis, D. K. (1988). Statements partly about observation. *Philosophical Papers*, *17*(1), 1–31.

Lewis, D. K. (1969). *Convention: A Philosophical Study*. Wiley-Blackwell.

Locke, J. (1694). *An essay concerning humane understanding In four books. Written by John Locke, Gent.*

Lukasiewicz, J. (1930). Philosophische Bemerkungen zu mehrwertigen Systemen des Aussagenkalküls. *Comptes Rendus de la Société des Sciences et des Lettres de Varsovie*, *23*(3), 51–77.

MacFarlane, J. (2007). Semantic Minimalism and Nonindexical Contextualism. In G. Preyer & G. Peter (Eds.), *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics* (pp. 240–250). Oxford University Press.

MacFarlane, J. (2009). Nonindexical Contextualism. *Synthese*, *166*(2), 231–250.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.

MacFarlane, J. (2020a). Indeterminacy as Indecision, Lecture I: Vagueness and Communication. *Journal of Philosophy*, *117*(11/12), 593–616.

MacFarlane, J. (2020b). Indeterminacy as Indecision, Lecture II: Seeing through the Clouds. *Journal of Philosophy*, *117*(11/12), 617–642.

MacFarlane, J. (2020c). Indeterminacy as Indecision, Lecture III: Indeterminacy as Indecision. *Journal of Philosophy*, *117*(11/12), 643–667.

Malpas, J. (2023). Donald Davidson. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2023). Metaphysics Research Lab, Stanford University.

Masius, D. (1617). *Commentariorum in Porphyrium et in universam Aristotelis Logicam*. Mainz.

McCulloch, G. (1995). *The Mind and its World*. Routledge.

McDowell, J. (1989). One Strand in the Private Language Argument. *Grazer Philosophische Studien*, *33*(1), 285–303.

Mill, J. S. (1843). *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence, and the Methods of Scientific Investigation* (Vol. 1). Cambridge University Press.

Miller, A. (1998). *Philosophy of Language*. New York: Mcgill-Queen's University Press.

Montague, R. (1960). Logical necessity, physical necessity, ethics, and quantifiers. *Inquiry: An Interdisciplinary Journal of Philosophy*, *3*(1-4), 259–269.

Montague, R. (1968). Pragmatics. In R. Klibansky (Ed.), *Contemporary Philosophy: A Survey, Volume 1* (pp. 102–22, Vol. 1). La Nuova Italia Editrice.

Montague, R. (1970). Pragmatics and intensional logic. *Synthese*, *22*(1-2), 68–94.

Moravcsik, J. M. (1998). *Meaning, Creativity, and the Partial Inscrutability of the Human Mind*. Center for the Study of Language & Information.

Morris, C. W. (1938). Foundations of the Theory of Signs. In *International encyclopedia of unified science* (pp. 1–59). Chicago University Press.

of St. Thomas, J. (1930). *Cursus Philosophicus* (B. Reiser, Ed.). Turin.

Partee, B. H. (1973). Some Structural Analogies Between Tenses and Pronouns in English. *Journal of Philosophy*, *70*(18), 601–609.

Perry, J., & Blackburn, S. (1986). Thought Without Representation. *Aristotelian Society Supplementary Volume*, *60*(1), 137–166.

Picazo, C. (2022). Are utterance truth-conditions systematically determined? *Inquiry*, *65*(8), 1020–1041.

Putnam, H. (1975). The Meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science*, *7*, 131–193.

Quine, W. V. O. (1960). *Word and Object*. MIT Press.

Raffman, D. (1994). Vagueness Without Paradox. *The Philosophical Review*, *103*(1), 41–74.

Recanati, F. (1989). The Pragmatics of What is Said. *Mind and Language*, *4*(4), 295–329.

Recanati, F. (2001). What Is Said. *Synthese*, *128*(1/2), 75–91.

Recanati, F. (2003). *Literal Meaning*. Cambridge University Press.

Roberts, C. (2006). Context in Dynamic Interpretation. In *The Handbook of Pragmatics* (pp. 197–220). John Wiley & Sons, Ltd.

Roberts, C. (2012). Information Structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, *5*, 6:1–69.

Russell, B. (1905). On Denoting. *Mind*, *14*(56), 479–493.

Rutherford, D. E. (1965). *Introduction to Lattice Theory*. Hafner Publishing Company.

Ryle, G. (1949). *The Concept of Mind: 60Th Anniversary Edition*. Hutchinson & Co.

Ryle, G. (1953). Ordinary language. *Philosophical Review*, *62*(2), 167–186.

Schaffer, J., & Knobe, J. (2012). Contrastive Knowledge Surveyed. *Noûs*, *46*(4), 675–708.

Schiffer, S. (1999). The Epistemic Theory of Vagueness. *Philosophical Perspectives*, *13*, 481–503.

Schoubye, A. J., & Stokke, A. (2016). What is Said? *Nous*, *50*(4), 759–793.

Scott, D. (1970). Advice on Modal Logic. In K. Lambert (Ed.), *Philosophical Problems in Logic: Some Recent Developments* (pp. 143–173). Springer Netherlands.

Searle, J. R. (1978). Literal Meaning. *Erkenntnis*, *13*.

Searle, J. R. (1995). *The Construction of Social Reality*. Free Press.

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication*. University of illinois Press.

Shapiro, S. (2006). *Vagueness in Context*. Oxford University Press.

Simons, M. (2005). Presuppositions and Relevance. In Z. G. Szabo (Ed.), *Semantics Versus Pragmatics* (pp. 329–255). Oxford University Press.

Simons, M. (2013). Pragmatics of Speech Actions. In M. Sbisà & K. Turner (Eds.). De Gruyter Mouton.

Siu, C. (2020). Tipper is Ready but He is Not Strong Enough: Minimal Proposition, Question Under Discussion, and What is Said. *Philosophical Studies*, *177*(9), 2577–2584.

Skordos, D., Myers, A., & Barner, D. (2022). Quantifier spreading and the question under discussion. *Cognition*, *226*, 105059.

Smiglecius, M. (1618). *Logica*. Ingolstadt.

Smith, A. (2014). *Boethius: On Aristotle On Interpretation 1-3*. Bloomsbury Academic.

Soames, S. (1998). *Understanding Truth*. Oxford University Press USA.

Sorensen, R. A. (2001). *Vagueness and contradiction* (Vol. 71). Oxford University Press.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Blackwell.

Stalnaker, R. (1974). Pragmatic Presuppositions. In R. Stalnaker (Ed.), *Context and Content* (pp. 47–62). Oxford University Press.

Stalnaker, R. (1975). Indicative Conditionals. *Philosophia*, *5*(3), 269–286.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, *9*, 315–332.

Stalnaker, R. (1986). Possible Worlds and Situations. *Journal of Philosophical Logic*, *15*(1), 109–123.

Stalnaker, R. (1988). Belief attribution and context. In R. H. Grimm & D. D. Merrill (Eds.), *Contents of Thought* (pp. 140–156). University of Arizona Press.

Stalnaker, R. (1998). On the Representation of Context. *Journal of Logic, Language, and Information*, *7*(1), 3–19.

Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, *25*(5/6), 701–721.

Stalnaker, R. (2004). Assertion Revisited: On the Interpretation of Two-Dimensional Modal Semantics. *Philosophical Studies*, *118*(1/2), 299–322.

Stevenson, C. L. (1944). *Ethics and Language*. Yale University Press.

Strawson, P. F. (1950). On Referring. *Mind*, *59*(235), 320–344.

Strawson, P. F. (1952). *Introduction to Logical Theory*. Wiley.

Strawson, P. F. (1964). Intention and Convention in Speech Acts. *Philosophical Review*, *73*(4), 439–460.

Stroud, B. (2000). *Meaning, Understanding, and Practice: Philosophical Essays*. Oxford University Press.

Suikkanen, J. (2019). Contextualism, Moral Disagreement, and Proposition Clouds. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (pp. 47–69). Oxford University Press.

Suppes, P. (2000). What is a scientific theory. In *The Nature of Scientific Theory*. Routledge.

Tarski, A. (1935). Der Wahrheitsbegriff in den Formalisierten Sprachen. *Studia Philosophica*, *1*, 261–405.

Taylor, D. E., & Burgess, A. (2015). What in the World Is Semantic Indeterminacy? *Analytic Philosophy*, *56*(4), 298–317.

Travis, C. (1996). Meaning's role in truth. *Mind*, *105*(419), 451–466.

Travis, C. (2000). *Unshadowed Thought: Representation in Thought and Language*. Harvard University Press.

Uzgalis, W. (2022). John Locke. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2022). Metaphysics Research Lab, Stanford University.

Vallée, R. (2003). Context-Sensitivity Beyond Indexicality. *Dialogue: Canadian Philosophical Review / Revue canadienne de philosophie*, *42*(1), 79–106.

Vanderschraaf, P., & Sillari, G. (2023). Common Knowledge. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2023). Metaphysics Research Lab, Stanford University.

van Elswyk, P. (2020). What the metasemantics of know is not. *Linguistics and Philosophy*, *43*(1), 69–82.

von Fintel, K., & Gillies, A. S. (2011). 'Might' Made Right. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality* (pp. 108–130). Oxford University Press.

Wharton, T. (2008). "MeaningNN" and "showing": Gricean intentions and relevance-theoretic intentions. *Intercultural Pragmatics*, *5*(2).

Williamson, T. (1994). *Vagueness*. Routledge.

Wittgenstein, L. (1953). *Philosophical Investigations* (G. E. M. Anscombe, Ed.). Wiley-Blackwell.

Wittgenstein, L. (1991). *On certainty*. Wiley-Blackwell.

Wittgenstein, L., Colombo, G. C. M., & Russell, B. (1990). *Tractatus Logico-Philosophicus* (C. K. Ogden, Ed.). Routledge.

Yablo, S. (2014). *Aboutness*. Princeton University Press.

Zalta, E. N. (2024). Gottlob Frege. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Spring 2024). Metaphysics Research Lab, Stanford University.