

HOW TO BE FRAGMENTED?

A Master's Thesis

by

Ali Kerem Erođlu

Department of Philosophy

İhsan Dođramacı Bilkent University

Ankara

May 2018

ALİ KEREM EROĐLU

HOW TO BE FRAGMENTED?

Bilkent University 2018

To my father,
my mother
and my brother

HOW TO BE FRAGMENTED?

**The Graduate School of Economics and Social Sciences
of
İhsan Doğramacı Bilkent University**

by

ALİ KEREM EROĞLU

**In Partial Fullfillment of the Requirements for the Degree of
MASTER OF ARTS**

**THE DEPARTMENT OF PHILOSOPHY
İHSAN DOĞRAMACI BİLKENT UNIVERSITY**

ANKARA

May 2018

I certify that I have read this thesis and have found that it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Arts in Philosophy.

W.G. Wringe

Asst. Prof. Dr. William G. Wringe (Bilkent University)

Supervisor

I certify that I have read this thesis and have found that it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Arts in Philosophy.

[Signature]

Asst. Prof. Dr. Y. Sandy Berkovski (Bilkent University)

Examining Committee Member

I certify that I have read this thesis and have found that it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Arts in Philosophy.

L. Thorpe

Asst. Prof. Dr. Lucas Thorpe (Boğaziçi University)

Examining Committee Member

Approval of the Graduate School of Economics and Social Sciences.

[Signature]

Prof. Dr. Halime Demirkan

Director

ABSTRACT

HOW TO BE FRAGMENTED?

Erođlu, Ali Kerem

M.A., Department of Philosophy

Supervisor: Asst. Prof. Dr. William Giles Wringe

May 2018

Actual human agents have limited cognitive capacity. They might display deductive failure, contradictory beliefs and imperfect recall. These and other similar cases raise a problem for idealized models of belief and behavior. The purpose of this thesis is to provide a way to accommodate these cases through a model of belief formation and retrieval based on information access. I will argue that belief formation and retrieval are sensitive to the informational context within which they take place. Human agents deploy information relative to the set of possibilities they take to be relevant.

Keywords: Belief, Belief Logic, Context Set, Decision Making, Fragmentation

ÖZET

NASIL PARÇALI OLUNUR?

Erođlu, Ali Kerem

Yüksek Lisans, Felsefe Programı

Tez Danışmanı: Dr. Öğr. Üyesi William Giles Wringe

Mayıs 2018

İnsanların bilişsel kapasiteleri sınırlıdır. Çıkarımsal hatalar yapabilir, çelişkili inançlara sahip olabilir ve parçalı hatırlayabilirler. Bunlar ve benzeri durumlar idealize edilmiş inanç ve davranış modelleri için bir problem oluşturmaktadır. Bu tezin amacı bunlar gibi problem teşkil eden durumları enformasyon erişimi üzerinden temellendirilen bir inanç oluşum ve ulaşım modeli ile açıklamaktır. Tez, inanç oluşumunun ve mevcut inançlara ulaşımın enformasyonel bağlama hassas olarak gerçekleştiğini savunmakta ve insanların enformasyon kullanımının mevcut enformasyonel bağlam içinde alakalı olduğunu düşündükleri ihtimal kümelerine bağlı olduğunu ortaya koymaktadır.

Anahtar kelimeler: Bağlam kümesi, İnanç, İnanç Mantığı, Karar verme, Parçalılık

ACKNOWLEDGEMENTS

I would like to thank my thesis supervisor, Asst. Prof. Dr. Bill Wringe for his crucial contributions to the thesis. I also thank Asst. Prof. Dr. Istvan Aranyosi and Asst. Prof. Dr. Sandy Berkovski for the valuable discussions on the content and the organization of this work.

This thesis was initiated during my visit to the School of Philosophy in the Australian National University. I thank Prof. Dr. Andy Egan and Prof. Dr. Daniel Stoljar for helping me to bring my initial ideas to maturity. I thank Prof. Dr. Alan Hajek for reading the first draft of this work and giving ingenious feedback.

I thank my parents, for they did not display any sign of discontent when I decided to quit my previous degree to study philosophy.

I thank my brother, for being my lifelong friend.

TABLE OF CONTENTS

ABSTRACT.....	iii
ÖZET.....	iv
ACKNOWLEDGMENTS	v
TABLE OF CONTENTS	vi
LIST OF TABLES	vii
CHAPTER I: INTRODUCTION.....	1
1.1 Logic of Belief.....	2
1.2 Hintikka’s Doxastic Logic and AGM model of Belief	3
1.3 Logically Idealized Agents.....	5
CHAPTER II: THE PROBLEM.....	11
2.1 Three Cases.....	11
CHAPTER III: INDEXING THE CONDITIONS.....	15
3.1 Accuracy of Indexing.....	16
CHAPTER IV: FROM INDEXING TO BELIEF IN CONTEXT OR HOW ARE FRAGMENTED STATES POSSIBLE?.....	21
4.1 Informational Context	21
4.2 Back to the Hacker.....	24
4.3 Entangled Houses	26
4.4 Cases Reconsidered.....	28
CHAPTER V: WHAT CAN THE FRAGMENTED STATES BE?.....	32
5.1 Fragmented sets of beliefs.....	32
CHAPTER VI: CONCLUSION.....	37
REFERENCES.....	39

LIST OF TABLES

1. Indexing for the Freshman...	15
2. Indexing for the Swamped Professor...	17
3. Indexing for the Fortunate Hacker	18

“To explain irrationality, we must find a way to keep what is essential to the character of the mental – which requires preserving a background of rationality – while allowing forms of causality that depart from the norms of rationality.”

– Donald Davidson

CHAPTER I

INTRODUCTION

Belief has been one of the most central notions in the history of Western philosophy. This is primarily due to the role of beliefs in human cognition and behavior. We take beliefs to be the building blocks of human reasoning and behavior. Therefore, there has been an enormous literature on belief. Nonetheless, there are many ongoing controversies even about the most basic features of beliefs. Some of the main topics in the recent philosophical literature of belief has been shaped around metaphysics, (Stich, 1983; Lycan, 1986; Schwitzgebel, 2001, 2010; Crane, 2001, 2013) reasoning (Stalnaker, 1984; Stich, 1993) and decision making and behavior. (Egan, 2008; Elga & Rayo, 2015)

This thesis falls in between these fields. It might be considered as a preface to an overarching theory of belief-and of intentionality-that attempts to achieve a general understanding of belief states of everyday agents. Given the complex nature of providing such an account, this comparably short work can only accomplish a first step in a grand journey. Being in the first step, I will try to leave the metaphysics of belief, which I take to be the last step of the journey, out as far as it can be kept out, although chapter 5 will inevitably discuss metaphysical constraints resulting from the particular understanding of belief proposed in the thesis. Still, rather than focusing on the metaphysics of belief, what I will do is to start from a simple but vicious problem stemming from a traditional understanding of belief and use that problem to start formulating an adequate and proper account.

In the first chapter, I will discuss the traditional logic of belief and knowledge put forward by Hintikka and the AGM model of belief. In the second chapter, I will proceed to the big problem about this type of formalization and its various treatments

in the recent literature. This problem, that is generally called “the problem of logical omniscience” is a vicious problem for Hintikka type idealized logics of belief. Given that any traditional formalized system will need to assume that agents are ideal in terms of logical consequences of their beliefs, the problem of logical omniscience- and some other resulting problems- will continue to pose a serious threat for any logical understanding of belief and knowledge. In the third chapter, I will consider a very recent attempt to deal with the problem that is put forward by Elga and Rayo (2015). I will argue that their attempt does not achieve what they are trying to achieve.

In the fourth chapter, I will attempt to provide a model that works for any kind of traditional logic of belief that deals with the problem, which use traditional logical tools to make sense of belief and knowledge. My attempt to save traditional models of belief from the problem requires a novel understanding of belief that treats belief sets as fragments, rather than ascribing a unified, single set of beliefs to human agents. This understanding provides a proper description of the different aspects of reasoning and behavior of actual human agents. In the fifth section, I consider some objections in the literature to ascribe fragmented sets of beliefs to agents and how my model tackles them. In that section, some metaphysical constraints on the nature of fragments will also be discussed. The sixth chapter will include some conclusory remarks.

1.1. Logic of Belief

A rigorous study of belief is required for many reasons. A non-exhaustive list of such reasons consists of the questions;

What are the objects of beliefs? What is the relation of the agent holding the beliefs to these objects of belief? What is the place of beliefs in the agent’s cognitive processes such as reasoning and rational behavior? What is the nature of our practice of ascribing belief to others and ourselves? What is the relation between these two practices and the nature of intentionality?

Given these questions, a doxastic and/or epistemic logic should represent the objects of belief in some specific sense which provides a general framework to understand belief states, their bidirectional relation both to the supposedly represented external object and to the animate agent herself. It should also shed light on our practice of

belief ascription through the logical treatment of belief in cognition and behavior. In that sense, the simple objective of a logic of belief is to have a rigorous system to model the relation between the agent who believes and the content of the belief.

The first task in doing so is to decide how to formalize belief content. Most formalized models of belief including Hintikka's use of modal logic takes the content of beliefs as propositions with a particular understanding of propositions that is based on possible worlds. Although I am sympathetic to the view that belief content are propositions as sets of possible worlds, I won't be assuming that this is the case but I will leave the detailed discussion elsewhere. The second task of a logic of belief is to make clear the relation of believing. That is, a formal model of belief is supposed to explain the relationship between an agent who believes a particular proposition and that particular proposition. This requires an operation which represents "having" or "acquiring" beliefs. One of the first well developed attempts to achieve these two tasks with a formal language was undertaken by Hintikka (1962). Hintikka used modal logic to clarify the "knowing" and "believing" relation between an agent and a proposition.

Hintikka's approach has no doubt provided theorists with many advantages in understanding what Russell defines as "...the passage from one belief to another..." (1921) or what Stalnaker calls inquiry: "...the process of forming and revising beliefs." (1984) Nonetheless, this and other similar formalizations suffer the problem of logical omniscience and some other resulting problems like having contradictory beliefs or impartial recall.

The second section of this chapter is a brief discussion of these idealized models of belief and knowledge, the problems stemming from them and the need for a solution.

1.2. Hintikka's Doxastic Logic and AGM model of Belief

In Hintikka's system, belief and knowledge are represented as propositional attitudes in the following way;

$K_c A$ is to say "C knows that A" and

$B_c A$ is to say "C believes that A"

where semantic content is to be interpreted as functions from possible worlds to truth values. In less technical terms, when one engages in attitudes of knowing and

believing, one simply distinguishes between different possibilities which are compatible with some possible worlds and not compatible with others, including the world one is in. This much of Hintikka's logic suffice to make the point required for the purposes of this paper:

Given the standard commitments of Hintikka logic, in such representation of belief and knowledge, agents must know all logical truths and their beliefs must be deductively closed. This, however, is apparently not the case for actual agents. Human agents do not know all logical truths following from their beliefs. This creates an important problem where this way of representing belief is not descriptively accurate for agents' states.

Hintikka's model does not attempt to explain belief change. For my purposes in this thesis, I do not need to get into a logic that considers belief change either. I will introduce another model very briefly, however, as it is a good example for another set-theoretical system which comes with a built-in ideal agent. This model is called AGM model after its developers' names. (Alchourrón, Gardenfors & Mackinson, 1985)

The issue about ideal agents is quite the same for AGM model of belief/theory revision. AGM uses a more simplistic set theoretical notation for representing belief states of agents;

A ∈ T is to say that "A is in the agent's belief set T"

There are three ways of revising propositions in a set in AGM. First is expansion where one adds new propositions compatible with the rest of the propositions in the set. Second is revision where propositions are added and if it is required, some old propositions are dropped out or changed in order to ensure consistency. And third is contraction, which is dropping out some propositions. These can be simply represented as;

T + A is the resulting set of expanding the belief set T by A

T * A is the resulting set of revising the belief set T by A

T - A is the resulting set of contracting the belief set by A

AGM model also operates under the assumption of an ideal agent just as Hintikka's doxastic and epistemic logic does. There is this built in requirement for agents to

know all logical truths in their belief set where agents are represented as logically omniscient.

One important difference between Hintikka and AGM is that Hintikka doesn't attempt to give an account of belief change while the latter is particularly to represent that. In that sense, it might be said that Hintikka's logic is concerned with present belief states of agents while AGM model tries to give a formal language for the conditions under which belief change occur. Krister Segerberg (1999) puts it by saying that AGM tries to represent belief and belief change in an absolute sense whereas Hintikka's logic treats them by including the importance of the context and hence indexical.

1.3. Logically Idealized Agents

Both of these logics treat belief states as unified-complete sets of beliefs but this has unwanted consequences which can be stemming from many motivations. Evidently, actual human agents are not logically omniscient, that is, deductively closed. They might have contradictory beliefs and they fail to recall particular propositions they believe in particular circumstances. Traditional logic of belief fails to explain and/or represent such aspects of belief and hence descriptively inadequate. This is already accepted in much of the recent literature but the one of the most striking discussion is by Robert Stalnaker's "The problem of logical omniscience" (1980). Stalnaker discusses the motivations to assume that doxastic agents are idealized in formal models of belief and lists four motivations to justify the idealization.

1. Getting at underlying mechanisms of reasoning
2. Simplification
3. Normative Reasons
4. Inevitableness

According to him none of these can provide a ground for representing agents' having an omniscient knowledge as an idealization. One justification is provided by Levi (1991), where he makes it clear that this idealization is not about an agent's actual belief states but about the beliefs he would hold in an all-powerful condition. In the recent literature on dynamic belief modelling, this built in idealization is accepted following Levi. Stalnaker convincingly puts that, nonetheless, the gap between the formal language of belief/knowledge and its intended domain of application is not a

problem to be avoided by appealing to any of these four, including Levi's. The problem requires further attention and attempts of a solution "...to get clearer about the concepts of knowledge and belief, and about what work we should expect a logic of these concepts to do." (Stalnaker, 1980)

This seems to amount to saying that one should either take the project of formally representing belief not as a project that has to do with such actual states of belief bearers or one needs to deal with the problem. If one takes the former standpoint, one needs to do nothing. But she will be condemned to a talk of belief applicable only to ideal agents. I take the formal systems of belief to be indeed attempts to represent actual states of belief of agents. There are many who think similarly and who offered different ways to deal with that problem.

Before discussing the ways, however, it is better to have a clear statement of the problem that the formal idealized systems of belief and knowledge have to face.

(The problem) Idealized formal systems of belief cannot properly model actual belief states of human agents.

The problem arises from the following fact.

(The fact) Human agents can find themselves in cases (not exhaustively) involving contradictory beliefs, imperfect access to their beliefs and deductive failure.

Given the fact, two ways of dealing with the problem exists. One is to take action from the standpoint of the logician. If actual states of the agents create a problem that cannot be properly modelled by existing logical tools, then the solution is about thinking about ways in which human cognition conforms to these tools. This has two further ways. One is to stick with the traditional tools while trying to solve the problem and the other is to come up with different sorts of-nontraditional-logical perspective to solve it. Some forms of paraconsistent logics are examples of the latter and the fragmentation project that argues human agents should be represented as having fragmented sets of beliefs is the mere example of the former.¹

Regarding the relevant forms of paraconsistent logic, it suffices to say that they are not particularly aiming at representing belief. Rather, they try to provide logics that

¹ It is not necessary to get into the literature of paraconsistent logic for the purposes of this work. Interested reader might want to see (Rescher, 1980) and (da Costa, 1989).

deny that inconsistency leads to absurdity. One possible way to do that, for instance, is to postulate non-standard possible worlds that can be taken as the base of the mapping of the function. da Costa & French (1989) uses such a paraconsistent base to have room for consistency in a set. Similarly, Rescher and Brandom (1980) offers a logic based on non-standard possible worlds to solve the problem.

The thing to model about human agents with limited cognitive capacity, however, is not the inconsistency. On the contrary, the element of rationality should be maintained in a model of belief and behavior. It is clear that humans cannot be deductively closed because of their cognitive limits and not because of an internally peculiar characteristic of the nature of their reasoning. When their inconsistencies are revealed, people are immediately disposed to revise their beliefs and are not disposed to keep the inconsistency in a way that is peculiar to their cognition. Therefore, the problem cannot be solved with the use of non-traditional logical tools.

How can we find room for **the fact** in a formal system, then? As we have seen in the AGM model, theory change can be modelled as an ideal formal system consisting of sets of propositions in a dynamic way but the case of actual agents does not only require a dynamicity in terms of the change of the propositions and the system but also a dynamicity of the “sets of propositions” the system consists of. That is, for a given time t , agent A has access to a particular set of her beliefs S while at time t_1 , she has access to S_1 .² This is to say that representing belief states as belief-sets in the traditional way is not enough since it cannot capture the fact. The belief states are dynamic in a sense that neither AGM nor Hintikka can give an account for. This dynamicity of actual agents has to do with the constant change of the beliefs agents have access to. It is not sufficient to represent present belief states. Present beliefs are in constant change depending on the situations that promote deploying certain information. Hence a formal system of belief needs also to represent the dispositions one has to act in certain ways in the face of a situation requiring a certain kind of – receiving or retrieving- information deployment. In the light of these one can formulate the task of a formal model of belief to take a step to solve the problem as follows;

² These specific time slices typically need to be situations where agents are to use some of their beliefs to make an inference or to be guided for their actions etc. Elga and Rayo (2015) gives an informational interpretation of such time slices that I favor where they call such situations elicitation conditions. There is more to come on this in the second chapter.

(The task) Agents' doxastic dispositions needs to be represented.

In order to achieve this task, I shall introduce another fact that seems utterly evident. Situations that makes an agent deploy some information for a particular purpose are about one's relational properties with her environment.

(The fact-2) Agents' doxastic dispositions are determined by the particular situations that prompt agents to deploy information for a particular purpose.

Following the task and the fact-2, it is clear that one needs to be able represent the dynamics of the information deployment given a particular situation in order to represent the doxastic dispositions of a particular agent.

The task, then, requires representing the belief states of an agent in a particular environment, at a given time that provides a context to deploy some information. One needs a rigorous formulation of this situation where belief formation and/or retrieval (based on information deployment) occurs with respect to the relational properties between an agent and her environment.³ In a similar fashion, Segerberg writes;

While the world state is no problem, the state of the agent is more complicated. The agent has a certain belief-set; that is, the total set of propositional beliefs (about the world) that he holds. Each belief can be represented in U by a subset of U : the set of points u such that if u is the real world state, then the belief is true. Thus also the agent's belief-set can be represented as a certain subset of U , namely, as the intersection of the representations of his beliefs. But this belief-set cannot be the same as his belief state, for it seems obvious that two agents may hold identical beliefs about the world (as to what is the actual state of the world) and yet react very differently to new information. In other words, it is not enough to represent the agent's current set of beliefs about the world, we must also represent what might be called his doxastic dispositions - the

³ Such a rigorous treatment of attitudes whose representation depends on the "time slice" of an agent is named as "de se attitudes" by Lewis (1986). De se attitudes are typically taken to be important to distinguish between self-locating beliefs and other beliefs. For now, what is important for my purposes is that time slice specifies the world plus the individual who has the relevant belief state. This amounts to specifying every relevant feature of the agent A who believes proposition P at a given time t .

way or ways he is disposed to act in the face of new information; in fact, the way or ways he is disposed to act in the face of any sequence of new pieces of information. (1999, p. 142)

This is true for many respects but not sufficient for a proper model since Segerberg assumes that the aspect to represent by a belief model is merely the dispositions of agents with respect to new information. The problem, however, is not only about an agent's disposition in the face of new information but in the face of "any" task that requires deploying certain information. This information can be in the form of receiving new information, retrieving information from one's memory or extract information from an external item (e.g. a notebook).⁴ In all these cases, it is the role of beliefs, combined with desires, that determines one's dispositions to act in a certain way. Treating belief states of actual agents as belief sets, therefore, is not enough to sufficiently represent the relation between the believer and what is believed. That is, without representing such dispositions to act, it is not possible to represent agent's belief states properly.

It seems that the traditional logic with its traditional logical commitments suffices to represent the environment (standard non-centered possible world, W) with the possible worlds framework as Segerberg clearly puts but representing the individual and her changing "mind-set" is the main problem the theorist must face.

Therefore, the question to ask to achieve a well-functioning and sufficient logic of belief and a proper model of behavior is to figure out how that new information is processed and/or coded information in memory is retrieved in an informational sense.

In the next chapter, I will give concrete examples of actual agents, demonstrate how these examples pose the threat by giving rise to the problem and critically examine a recent attempt to achieve the task as it is proposed by Elga and Rayo (2015). As we shall see, they attempt to provide a descriptive framework for these cases through a method they call indexing: specifying what information is available relative to different purposes of an agent. (p.5) I will argue that, even though indexing is indeed required to properly represent such states, Elga and Rayo's method fails to achieve

⁴Elga & Rayo (2015) illustrates this with ingenious examples, which I will discuss in more detail in the second chapter.

its goal. The second section will conclude that their proposed framework based on indexing falls into triviality if it is to describe the target cases accurately. In the last section, I will offer a refined version of indexing through what I call “informational context”, which provides a fruitful model for these states and offers better descriptive accuracy.

CHAPTER II

THE PROBLEM

This chapter is a thorough discussion of the problem, its elaboration and its constraints on the notion of belief.

The following section will provide concrete examples of the cases that pose some threat for idealized models of belief and behavior. After giving some examples of failure of deductive closure, contradictory beliefs and imperfect recall I will discuss how these are to be understood in an information-based account of belief and behavior, in what ways they are related to each other and to what extent they challenge the idealized models of cognition.⁵

2.1 Three Cases

Consider the following cases.

Case 1: A philosophy freshman forms the following beliefs A: that Frege was older than Russell and B: that Russell was older than Wittgenstein, say, in an introductory class of analytic philosophy. By hypothetical syllogism, one expects her to believe C: that Frege was older than Wittgenstein. She, however, does not believe that C; when someone asks her a question comparing Frege's and Wittgenstein's age, she thinks that she does not know. As a matter of fact, she lacks a belief about Frege's and Wittgenstein's compared ages. This is an example of failure of deductive closure.

Case 2: After sometime the same freshman forms belief D: that Carnap is older than Wittgenstein, although she well knows from the analytic philosophy class E: that Carnap is younger than Wittgenstein. When she is asked a question comparing the ages of Carnap and Wittgenstein, she acts accordingly with E, if the conversation is

⁵ By an information based account of belief, I do not necessarily mean an informational-theoretic account content that one can find in Dretske (1981). I prefer using information and information deployment as helpful notions to explain accessing the content of belief.

about the analytic philosophy course and with D if she is discussing about how Vienna Circle interprets early Wittgenstein. This is an example of having contradictory beliefs.

Case 3: A flight attendant knows F: that Cheryl Piecewise sits at 57C. She forms belief F while she is skimming the passenger list before the flight. Later in the flight, for some reason, when a colleague asks her “What is the seat number of Cheryl Piecewise?” she answers “57C.” without a second thought. But when her colleague asks her “Who is sitting at 57C?” she fails to give an answer. This is an example of imperfect recall.

In all these cases, agents are prompted to deploy some information for a particular purpose. It is the salience or lack of salience of information, it seems, that makes these phenomena possible. For instance, when the freshman is to deploy information for the conversation involving analytic philosophy course, the information that Carnap is younger than Wittgenstein is made salient rather than the information that Wittgenstein is younger than Carnap. The question is, then, why some particular information is made salient given these circumstances rather than some other information or no information.

Once we understand such cases as a matter of information access, as Stalnaker argues, we need to spell out the accessibility conditions of information so as to answer this question.⁶ (Stalnaker, 1999) He writes:

...the problem is that we need to understand knowledge and belief as capacities and dispositions-states that involve the capacity to access information, and not just its storage-in order to distinguish what we actually know and believe, in the ordinary sense, from what we know and believe only implicitly. (p. 253)

If we are to accommodate these cases in an account of belief, we need to treat beliefs as dispositions to access information. In that picture, beliefs are capacities not only to store but also to retrieve information for a particular purpose.⁷

⁶ The reader might want to go back to Hintikka (1962) to see how can this accessibility relation be understood in terms of possible worlds framework.

⁷ An information based account has many merits against a merely psychological one which does not appeal to any intentional properties of an agent. A mere psychological account, for instance, can explain how underlying psychological mechanisms make an agent bring what information to mind but

The first reason this is required is the following: we do not ordinarily attend to most of our beliefs. Still, as long as we can retrieve them when needed, even occasionally, we can be said to believe them.⁸ It is clear, however, that we cannot retrieve them whenever we need them since actual human agents do not have infinite computational power. Notice that this does not need to be due to forgetfulness. It might happen, for instance, when someone asks me a question about early eras in the evolutionary scale. Having taken a senior course on evolution in my biology minor, I do have beliefs, probably many correct ones, about the early eras of evolutionary history. I did not forget them. Say, as a matter of fact, I had discussions about them just a couple of hours ago. These beliefs are fresh in my mind. I can retrieve them. It is just that I can't do it *now*. They are not accessible to me now, to my present purposes, under my present circumstances.

How does this work with abovementioned phenomena? What happened when I failed to retrieve these beliefs? Does that make me not know or fail to have beliefs about the species in the Mesozoic Era? Consider the third case. Does the attendant believe that Cheryl Piecewise's seat number is 57C? If the answer is "yes" with respect to the first question and "no" for the second, does this amount to say that she sometimes believes that, and sometimes does not?⁹

Second, beliefs guide our behavior. If I believe that I have class at location *l* and at time *t*, I act accordingly. Beliefs guide our actions, combined with our desires. Consider Lewis' example (Lewis, 1982). If I believe that Nassau Street runs North-South, combined with my desire (and purpose) to go South, I use Nassau Street to go to my destination at South. But if I also believe that Nassau Street runs East-West then I should not take Nassau Street. What belief are we to ascribe to me? Which

it cannot explain why this is the case. A comprehensive comparison between the two is out of the scope of this thesis.

⁸ This is sometimes referred as the distinction between occurrent and dispositional belief. See (Lycan, 1986) and (Schwitzgebel, 2010).

⁹ See (Schwitzgebel, 2001), where he argues that it is better to ascribe neither a belief nor a disbelief in similar cases.

belief do I take to guide my action? Am I going to my friend's place in South, or my favorite diner at West?¹⁰

Understanding belief states as dispositions not only to store but also to access information and to act in certain ways is important especially given these two reasons. If beliefs are dispositions of this kind, the problem of information access becomes a fairly subtle one; if one is to spell out accessibility conditions, one can only do it by considering the action the relevant belief would guide. Needing to appeal to every particular purpose to deploy information to explain belief and behavior is a difficulty for it is a dynamic process that depends on agents' always changing circumstances. We need, then, a model first to describe and second to explain the accessibility conditions of information. First, what information is accessible to a particular purpose is to be described. Second, why a body of information is accessible to that purpose rather than some other information or no information is to be spelled out.

¹⁰ Here I am not committing myself to a particular type of account of belief called dispositionalism, which, roughly, states that believing is fundamentally having a set of dispositions to behave in certain ways.

CHAPTER III

INDEXING THE CONDITIONS

Elga and Rayo (2015) attempt to accomplish the first task. They argue that these and similar other cases can only be accommodated through a method of indexing, whose function is to specify "...what information is available to an agent relative to various purposes." (p.5). Their strategy is to construct access tables which specify what information is available under which conditions. This helps to create a descriptive table of accessibility conditions which links what they call elicitation conditions to available/unavailable information. Elicitation conditions are situations which prompt an agent to deploy information for a particular purpose. (p. 5)

To illustrate, an access table for the freshman in the second case can be constructed as seen Table 1:

Table 1. Indexing for the Freshman

Elicitation Condition	Information available relative to that condition
Discussion on Wittgenstein's impact on Vienna Circle	Information that Carnap is older than Wittgenstein is available.
Conversation on Introduction to Analytic Philosophy Classes	Information that Wittgenstein is older than Carnap is available.

Notice that the whole point of constructing access tables is about representing states that cannot be normally represented, exemplified by the above cases. What Elga and Rayo offer is a framework where they put forward the conditions under which one has access to particular information. Indexing, by its appeal to elicitation conditions,

which are situations where one is to act with some purpose, attempts to describe particular actions for which certain information is available and particular actions for which certain information is not available, or different information is available. Just by itself, this might seem to be in vain at first sight; if Elga and Rayo draw a table of access conditions and leave it there, how could that be helpful in making sense of these states of the agents?

A driving motivation behind Elga and Rayo's account is the merits of using information access as a concept. Describing these cases in terms of information possession is valuable in explaining the behavior of the agent. When we have to explain how an agent got to the last question in "Who Wants to be a Millionaire?", for instance, we speak in terms of her access to the relevant information of the answers to the previous questions. We do not need more than that to explain her success. We would be able to say, based on our access table that she had accessed this particular information relative to the purpose of answering that question, and that would do to explain her success.

3.1 Accuracy of Indexing

Explaining and predicting agents' behaviors based on access tables, nonetheless, requires one's descriptive table to accurately capture the link between the purpose of the agent and the value of the access variable (i.e. which information is available/not available). Consider, again, the Cheryl Piecewise case. When the flight attendant has access to relevant information with the purpose of answering a question, the table must show that that information is made salient *relative to the purpose of answering that question*. Access tables are constructed to represent the relationship between the elicitation conditions and the available body of information. In the rest of this section, I argue that this relationship is not necessarily captured by Elga and Rayo's way of indexing. To show how, let me start with an unsuccessful case of information access and its demonstration in an access table.

Swamped Professor: Think of a busy school teacher trying to deploy information about the grade of one of her students from a list she previously prepared, in the last minute of entering the grades to the system. Her purpose is to deploy information about the grade and the list (sorted by surname alphabetically rather than by grade) is totally in accord with that purpose. That is, given her purpose, the relevant

information is available to the teacher. However, she fails to access the information by missing the relevant line over and over again, due to her being tired and in a hurry. The access table for the professor in this case would be:

Table 2. Indexing for the Swamped Professor

Elicitation Condition	Information available relative to that condition
Extracting the grade of a given student	No information available

One expects access tables to represent the relationship between elicitation conditions and the available information. In this case, however, even though the elicitation condition is well in accord with the information, the access table reports the exact opposite. The table fails to describe the case accurately and provides a misleading characterization of information access relative to a purpose.

One may accept this and nonetheless think, as Elga and Rayo seem to do, that descriptive accuracy is provided de facto in successful access cases since the information is simply available to that particular purpose in these. The very fact that I can answer a question assures the link between available information and my purpose to answer that question. If this is correct, then one can argue that the access table framework can still be valuable, at least for the successful access cases.

Consider, then, following successful case.

Fortunate Hacker¹¹: Think of a hacker from the ‘80’s who is given a phone number, the owner of which happens to be rich. By posing as a bank’s accountant on the phone and using the owner’s name, the hacker intends to earn the owner’s trust and learn the information needed to gain access to her bank account. For his mission, the hacker goes to a store and asks for a phonebook sorted by number. The wrong phonebook, sorted by name, however, is delivered to her without her noticing. She opens the phonebook at an arbitrary page while preparing to call the given number. Only then she realizes that she was given the wrong phonebook. Just before she

¹¹ This is an altered version of the hacker example introduced in (Elga & Rayo, 2015).

closes the cover in anger, nonetheless, her eyes catch the exact same number she needs in the page she happened to open the book at. The access table in this case would be:

Table 3. Indexing for the Fortunate Hacker

Elicitation Condition	Information available relative to that Condition
Extracting the name of the owner of a given number	The name of the owner of number X is Y.

The information's being accessible to the hacker, here, has nothing to do with the way information is present and its link to the elicitation condition. Moreover, it is clear that the way information is present (being sorted by name) is not in accord with the given purpose of the hacker. Again, the access table fails to describe the case since it fails to capture the relationship between the elicitation condition and information access.

In short, neither of these cases accurately describes the connection between the elicitation conditions and available information, while the purportedly singular job of indexing is to capture that connection descriptively.

There are two ways to tackle these criticisms. One is to say that indexing "generally" captures the relationship between purpose and available information. The fortunate hacker case has a dramatically low probability of actually happening.

I agree with that reply. Access tables are usually descriptively accurate. But this is not enough, since there is always a possibility that they are misleading. The competitor in "Who wants to be a millionaire?" might be successful just coincidentally. Information can be made salient not in relation to the conditions we take to be elicitation conditions but to some other aspect within the context. Besides, such a reply does not explain the elderly swamped professor. Such failures happen all the time due to hurry, anxiety, nervousness, tiredness etc. Such cases are extremely common.

The second reply is to add another criterion to this type of indexing: indexing works

under optimal conditions of information access (e.g. being calm and collected) where agents are not to make mistakes when they have access to information relative to their purpose. This, however, is again in conflict with the main motivation of indexing: representing the actual cases of information deployment. Agents are often in suboptimal conditions and this would have an important impact on their access to a body of information.¹²

At this point an important question arises: What are elicitation conditions exactly? If one can claim that elicitation conditions contain these aspects of the agent, then one can argue that indexing describes information access accurately. One might think that elicitation conditions contain such situational aspects; if, for instance, the swamped professor fails to deploy the information, we could take it that the relevant information is not available given that particular elicitation condition.

There are two problems with this reply.

The first is that this characterization is in conflict with what Elga and Rayo have in mind as elicitation conditions. Elga and Rayo define elicitation conditions as choice situations for an agent where they, "...prompt an agent to deploy information for a particular purpose." (p. 5) For our swamped professor, the choice situation is not her being tired but it is her being in the last minute to enter the scores with the list sorted by name before her eyes. The very reason Elga and Rayo introduce the term "elicitation conditions" is to distinguish the choice prompting conditions (like answering a question) from other conditions an agent is in especially since their aim is to represent the information available relative to that elicitation condition.

One can resist. One can argue that the elicitation condition for the professor of course involves the elements of haste and tiredness when it is given that she tries to deploy information at that particular time t . If one is going to adopt that strategy and state that elicitation conditions include such aspects of the agent, though, the consequences are disastrous. By going fine-grained enough in the characterization of elicitation conditions one can come up with an access table that would match any type of behavior. This is explicitly acknowledged by Elga and Rayo:

¹² Suboptimal conditions are not required for an agent to have fragmented states or imperfect recall. It is totally possible and common to have fragmented states under optimal conditions. See (Egan, 2008), for instance, for an argument about why being fragmented is better for belief revision.

We admit that access tables wouldn't be interesting if in order to use them, one needed to resort to the trick of appealing to arbitrarily specific elicitation conditions. But one needn't resort to that trick. (p. 17)

As we have seen, however, one really *does* need to resort to that trick if one is to make sure of the descriptive accuracy of the access tables. If one needs to resort to that trick, then the whole project of using access tables as a descriptive framework becomes trivial. And if one stays at a coarse-grained level, indexing can always be descriptively misleading and hence explanatorily faulty.

The framework of Elga and Rayo, it seems, is caught between two fires. For a type of indexing to work, then, there are at least two things to be satisfied:

- a. Indexing must be descriptively accurate; it must assure the link between available information and the conditions to which one wishes to index information access.
- b. Indexing must enable one to explain not only the behavior of an agent but also how their fragmented/confused states are possible.

Despite all the problems their account has, Elga and Rayo's approach capture the very core aspect of fragmentation promoting cases: if the way the information is present does not *fit*, so to speak, with an agent's purposes, then it is not immediately available to the agent. Hence, indexing needs a supplementary model of information access which explains what information *would be* available relative to various purposes of an agent. To satisfy the conditions above one needs to describe (i) how information is present to an agent, combined with the agent's situational characteristics that determine their accessibility relation, and model (ii) the processes taking place in the formation or retrieval of the relevant belief states. These two can be achieved by staying coarse-grained to satisfy **a** without falling into triviality and by modeling the mechanisms at work during information access to satisfy **b**.

In what follows, I attempt to introduce such a model.

CHAPTER IV

FROM INDEXING TO BELIEF IN CONTEXT: OR HOW ARE FRAGMENTED STATES POSSIBLE?

In this chapter, I argue that a refined indexing that is descriptively accurate has to do with the dynamic context in which one forms, revises, and acts in accord with beliefs based on receiving or retrieving information. My conclusion will be that the beliefs we form and retrieve are sensitive to the context in which we are to deploy information. I offer a model which captures the dynamics of information deployment which leads us to have a descriptively accurate and explanatorily more powerful indexing.

4.1 Informational Context

What kind of context will I be talking about? People use the term “context” for many different things. But the underlying idea is the same. It is the interrelated conditions in which something occurs. For instance, if we are to think that interrelated conditions in a conversation are important to understand what a speech act means, then we consider the conditions in which the speech act occurs. Here, we are to think that the interrelated conditions are important to understand information deployment since agents can access information in some interrelated conditions and not in others- and/or access different/contradicting information. Our context, therefore, is the interrelated conditions in which a sort of information reception or retrieval is to occur. Once our target context is determined, we need to provide a model of it to understand how the context works. This is the case in the philosophy of language literature when theorists aim to understand how the interrelated conditions in which a speech act occurs work.

My aim here is to offer such a model for the conditions in which a sort of information is to be deployed by an agent. In order to do that, however, I shall first

make a few remarks about the notion of context as it is used in the philosophy of language literature since I will use a model of conversational context as a template for my model. Since my model is based on the possible worlds framework, I will briefly discuss how this framework functions for a conception of context and common ground in the philosophy of language and then make some clarificatory remarks on what I will be assuming in the rest of the paper.

The possible worlds framework treats propositions as functions from possible worlds to truth values. In such treatment, as we have seen in the first chapter, engaging in a conversation, people distinguish between different possibilities. Similarly, belief content is also to make distinctions between possibilities (Stalnaker 1978, 1984). Given that, the Stalnakerian model of conversational context (Stalnaker 1978, 2002) roughly goes as follows: When parties start a conversation, they start with an initial common ground which reflects their presuppositions about the relevant discourse. Common ground consists of the propositions the speaker and the audience mutually believe or accept for the sake of conversation, and hence determines the dynamic discourse (Grice, 1981; Stalnaker, 1978, 2002). Any compatible assertion -except for the tautological assertions- made by a party makes a new distinction between possibilities, and hence introduces new information to the common ground, reducing the context set; the set of possible worlds compatible with the common ground. That is, an assertion in a context excludes some possible worlds from the context set as long as it is compatible with common ground. Notice that common ground determines whether an assertion successfully counts as new information and reduces the context set; if the content of the assertion is not compatible with the context set determined by the initial common ground, it fails to add to the common ground. If interlocutors thought of a different common ground (partly or totally) from what the present speaker has in mind, the speaker's assertion would not be successful in adding new information to the ground; its content might fail to express anything meaningful (for the conversation) or might express different things to the audience from what the speaker intends to assert. This model of conversational context is generally agreed upon since it provides a good way to model the dynamic relation between pragmatics and semantics, where the meaning of a speech act is

individuated dependent on the always changing context.¹³ Common ground, then, in such model, are crucial for successful conversation. And successful conversation means successful information exchange between speaker and audience.

Another thing to note to avoid possible confusions that might stem from the use of the term “contextualism” in the literature is that my account is not based on the claim that belief and/or knowledge *ascriptions* are context-dependent. Rather, having Stalnaker’s picture of conversational context as a template, my claim is that belief formation and retrieval *themselves* (as they are characterized by deployed information) is dependent on the context.

Given these, I shall start introducing my account with the notion of informational context (IC).

IC: The context in which a sort of information is to be deployed. In a slightly formal way; it is a <time, world, location> triple where a sort of information is to be deployed by the individual. This is the *positional character* of the IC. The positional character fixes all the relevant properties of the individual and the way information is present.¹⁴ Since the positional character fixes all the relevant properties of the individual who is to deploy information, it also covers the *presuppositional character* of itself, based on informational ground (IG).

IG: The agent’s presuppositions (beliefs, acceptances etc.) about what the relevant information source could/would provide, within a context, i.e. a given <time, world, location>.

The use of centered possible worlds to define the positional character of IC is quite similar to the use of the notion in Lewis (1980) and Egan (2009) to fix the properties of a speaker in conversational context. The simple idea behind applying this to informational context is that once you determine the time, world and place of the agent who is to deploy information for a particular purpose, you determine all the

¹³ See (Abbott 2008) and (Kecskes and Zhang 2009) for some of the objections and criticisms. See also (Hawthorne and Magidor 2009) for another criticism not directed at Stalnaker’s common ground, but his notion of context-set and how it works. See (Stalnaker 2014) for further developments of his picture of conversational context partly as a response to some of these criticisms. None of these criticisms are orthogonal to the way I utilize Stalnaker’s model to information deployment.

¹⁴ At that point, a discussion on possible worlds vs. centered worlds is omitted even though it is required, due to space concerns. See (Lewis 1979a) and (Egan 2009).

relevant facts about the agent - her intentions towards the source of information, her motivations etc. - that might have an impact on her resulting belief content. You also determine the IG at that given <time, world, location> triple.

Informational ground plays almost the same role as the common ground. It sets the tone for the upcoming information processing just as common ground fixes the domain of a given conversation. Much as common ground consists of the mutually accepted or believed propositions about the conversation and what the parties get from it, informational ground includes the agent's beliefs, suppositions and acceptances with regard to her relation to a salient information source. It fixes the domain of information process. Informational ground is determined by the agent's practical and situational characteristics in a context.¹⁵

IC, as we shall see, puts constraints on the possible mental representations an agent forms based on the present information.

4.2 Back to the hacker

Let's consider again the case of the fortunate hacker. The hacker needs to find a phonebook of the relevant neighborhood which is sorted by numbers rather than names. On the other hand, if she was given a name to defraud, she would need the exact opposite. Both phonebooks are stored with the information she needs; in fact, they consist of the same information. The information she needs, nonetheless, is accessible in one phonebook but is not accessible/requires too much time in the other. The straightforward reason for this is simply the conditions under which the same information is stored. For the reverse phonebook the purpose is to store the relevant information when the phone number is known and the name is to be found and vice versa. What information is available relative to her purpose is dependent on the way the storage occurs/information is present. The initial question must, then, concern the way a particular storage occurs/information is present.

In the Cheryl Piecewise case, for example, the informational context in which the attendant is to store the relevant information determines how the information is

¹⁵ One might follow here Stalnaker's agnosticism with regard to the nature of the presuppositions comprising the initial common ground. Informational ground, however, operates in a different fashion from common ground. A detailed discussion of it is indeed necessary, but out of the scope of this work.

present as belief content. Whether the storage takes place with an emphasis on the seat number or the name of the passenger is totally dependent on the attendant's purposes and plans about what to do with that information. Simply, present and potential purposes and the related interests of the attendant lead her to store the information with an emphasis on the seat number and not with the name. Her informational ground consists of the presuppositions determined by the practical purposes and interests which promote storage with the emphasis on the seat number rather than the name, say, since flight attendants generally use the seat number for service purposes (e.g. "57C has a problem").

It seems from this case that there are two ways of adding to informational ground. The first is the presuppositions about what the information source can provide as information per se, and the second is the presuppositions about how the agent would use the information. These two aspects intermingle in the informational ground. For example, the attendant has some presuppositions about what the passenger list could/would provide *as* mere information per se and what it could/would provide *as* information to be used for the present or potential purposes.

It is clear that the assertoric content of an utterance can be its semantic content. Stalnaker's model is for the times in a conversation where these two differ. It is the same, as we shall see, for beliefs too; belief content is normally individuated with respect to the information received in the actual world. This individuation occurs at times when the domain fixed by informational context does not rule out the possibilities corresponding to the fact. For instance, one's informational ground would not rule out the possibilities that this is Nassau street that the agent passes through but it can rule out the possibilities in which Nassau street goes North-South, given the particular encounter with the street. In such conditions, the possible worlds corresponding to the fact in the actual world are ruled out and hence are not accessible to the agent.

In other words, informational ground sometimes puts constraints on the possibilities an agent distinguishes between in such a way that it rules out some possibilities about what information source provides. When this constraint excludes some worlds in the context set, one might fail to form the relevant belief even though the

information would normally lead to that belief.¹⁶

4.3 Entangled Houses

In *Paradoxes of Rationality*, Davidson provides a great example for this, where he depicts an ordinary day he was on his way back home to his faculty housing from his office at Princeton:

It was a warm day, doors stood open... I walked in the door. I was not surprised to find my neighbor's wife in the house: she and my wife often visited. But I was slightly startled when, as I settled into a chair, she offered me a drink. Nevertheless, I accepted with gratitude. While she was in the kitchen, making the drink I noticed that the furniture had been rearranged, something my wife did from time to time. And then I realized the furniture had not only been rearranged, but much of it was new-or new to me. ...it slowly came to me that the room I was in...was a mirror-image of that room; stairs and fireplace had switched sides, as had the door to the kitchen. I had walked into the house next to mine. (2004, p.191)

Here, Davidson's <world, time, location> triple, i.e. his informational context puts constraints on the set of possibilities he can distinguish between by fixing all the relevant properties in the moment of receiving the information. World fixes the relevant world, time fixes the given time and location fixes the person who is to deploy information with her particular purpose. Importantly, it also fixes the way the information is present with respect to the agent since it fixes the location of the agent. In this case, it is Davidson, in his neighbor's house, in world *w*, at time *t*. His informational ground, also determined by his <world, time, location> triple, the presuppositions he has with regard to the information source, limits the context set in such a way that it excludes or not introduce as relevant the possible worlds in which this is his neighbor's house, leading to failure of belief formation. Although the information he keeps receiving should make him form the belief that this is his neighbor's house, his presuppositions given his situation, based on his purpose (even

¹⁶ "Normally" is used for the times where the fixed domain or context set does not rule out some possibilities about the way the world is.

trivial like taking a short nap after work) and interests, prevent him from forming the relevant belief.

In other words, the set of information that would normally lead to the belief that he was in his neighbor's house was not accessible to him as it normally is.¹⁷ Davidson puts this in a different manner:

...did I have contrary evidence? Not from my point of view, for I thought that my neighbor's wife was being exceptionally kind in offering a drink in my house; I thought my own wife had rearranged the furniture and even introduced some new furniture. I...did not make the possibly absurd mistake of supposing my evidence supported the hypothesis I was in my own house rather than in another. (2004, p.191)

That is, information source would normally lead to the belief that this is not his house and one would expect Davidson to form the relevant belief and leave the house, with, maybe, a little apology to his neighbor. But that information was not available to him given IC since it was ruled out in his context set determined by his IG. This is, for informational context, quite similar to the cases where a bifurcation between semantic content of the assertion and the meaning in context differs. Semantic content is sometimes not available as it would be available in another context; it is available to the interlocutors with the impact and the constraint of the particular context. What an information source can lead an agent to believe differs in the sense that the relevant belief is formed with the impact and the constraints of informational context. Informational ground, in that sense, determines what is accessible to the agent in a given context just as common ground determines what is accessible to the audience when a certain utterance is made in a given context.

What is more, if we are to ascribe a belief based on Davidson's actions, we typically ascribe him the belief that he was in his neighbor's house, making a short visit. He enters into the house, takes a seat, accepts his neighbor's kind drink offer. Everything is to be interpreted as he having the belief that he is in his neighbor's house paying a visit. This error is immensely problematic for any account of belief since the problem

¹⁷ See footnote 16.

in the chapter two arises here, again; how are we to explain his behavior based on his beliefs?

Before proceeding how my account helps to answer this question, I think there is another important thing to note. If Davidson did not encounter with accumulating contrary information, he would continue to believe that this was his house. For his informational ground would not be changed so as to re-adjust the context set to include some possible worlds in which this is not his house. Similar to the Lewis' case where the same street is to be treated as two different streets running different directions, Davidson would live happily without noticing the "gross factual error" he made.

Well, how does informational ground undergo a change? Accumulating contrary evidence incompatible with context set leads an agent to a kind of reconsideration of the information. Lewis (Lewis 1979b) makes it clear that within a well-run conversation presuppositions are created and destroyed all the time. A similar change takes place in the informational ground. Some presuppositions are destroyed and others added due to contrary or new information updating the context set.

4.4 Cases Reconsidered

Based on my earlier criticism on Elga and Rayo's indexing, this initial characterization of my model is supposed to help us with two things:

First, it has to be able to solve the problem of going fine grained enough to include failure and fortune within informational context, without the implication of triviality, in order to ensure descriptive accuracy.

Second, it has to be able say why that information is available given a particular IC, rather than some other relevant information or no information. This is to explain what information *would be* available relative to various purposes of an agent, which is equal to explain under what conditions these fragmented states are possible.

On the one hand, talk of IC, just by itself, ensures the descriptive accuracy without triviality. It is simply this: the positional (<world, time, location>) character of IC fixes all of the relevant properties related to information deployment without going fine-grained. In that sense, it fixes the actual conditions related to information access. By indexing to IC, we are ensuring descriptive accuracy by including any possible

suboptimal conditions with a coarse-grained characterization while doing at least the same amount of explanatory work as Elga and Rayo's type of indexing does. Recall that the problem with Elga and Rayo's framework was that information could have been made salient not in relation to the conditions we take to be elicitation conditions, but to some other aspect within the context. IC fixes the actual conditions including choice prompting conditions and the way the information exists and hence necessarily captures descriptive accuracy.

On the other hand, the model based on informational ground explains what information *would be* available relative to various purposes of an agent and how fragmented states are possible, given an IC. That model limits one's scope only to the cases where information availability is relative to the presuppositions contained in the informational ground and hence we do not need to be concerned with the cases of fortune or non-informational failure. Indeed, there is no way to propose a general model other than limiting the scope, since such fortune and failure can occur arbitrarily anytime. Once we limit our scope and set our goal to provide a general model for the information-related dynamic mechanisms (not others like fortune) within IC, we can narrow our scope to cases occurring under optimal conditions. Once we are to talk about information access relative to IC and IC's presuppositional character, we come to have a canonical explanation of how deductive failure, fragmented belief states and imperfect recall cases are possible.

The most straightforward explanation offered by this model is that of fragmented states, and I shall start with them.

Fragmented States and Deductive Failure: Fragmented states are simply due to the failure of informational context to cover both (or more) beliefs together. Take our freshmen case. In context C, where she thinks about Wittgenstein's impact on Vienna Circle, she believes D: that Carnap is older than Wittgenstein. This is where she does not have access to information E: that Wittgenstein is older than Carnap. Her informational ground in C excludes the possible worlds where E is true so that she could continue to believe that D. This is because she formed belief D while she was thinking or discussing Wittgenstein's impact on Vienna Circle, where her informational ground either excluded or didn't introduce as relevant the possible worlds where E is true. In context C' however, when she thinks about her undergraduate years, she believes E. Her informational ground in C', this time, either

excluded or didn't introduce as relevant the possible worlds where D is true. As long as IGs in C and C' are not updated in such a way that they make the relevant context sets start including possible worlds where E or F is true, her fragmented state stays in place.

Whenever she is within an IC in which she forms belief D, she has access to that belief and acts accordingly. Whenever she is within an IC in which she forms belief E, she has access to that belief and acts accordingly.

If one day, however, a friend of hers tells her something about that analytic philosophy class while they were discussing about Wittgenstein's impact on Vienna Circle, the relevant context was introduced, and hence she would have access to the relevant information through adding to the IG and updating the context set, which enables her to have an "aha!" moment, where she realizes she has contradictory beliefs.

For deductive failure, the case is the same. As long as the relevant context, where she is to consider beliefs A and B together, is not introduced as relevant she fails to deduce C. Just as in the abovementioned case of fragmented belief states, however, as soon as the context is introduced, that "aha!" moment takes place where she can compare the beliefs A and B and make the relevant logical deduction.

Let me now proceed to two imperfect recall cases.

Imperfect Recall-1: IG partly contains agents' presuppositions about the information source. Take Cheryl Piecewise case again. As noted above, the flight attendant's presuppositions are shaped in at least two different ways. One is her presuppositions about what the list of passengers can provide her per se, and the other way is her other presuppositions about how that information would be used by her with her present or potential purposes. The second type of presuppositions makes her pick a specific storage for the information: either with an emphasis on seat number or with an emphasis on name. This determines in turn to what question the relevant information becomes immediately available and to what question it does not.

The emphasis here is not a mysterious property of an agent. It is what we value more in the new information we receive, which is based on present or potential ways of

using it and hence contained in IG. If she values seat numbers more than the names, let us say, due to her profession, then she sees the relevant connection with the seat number easier than she would with the name.

Imperfect Recall-2: You are asked to come up with an English word the last two letters of which are “mt” (Elga and Rayo 2015, p. 3). You cannot answer. Your friend gives the answer “dreamt”, and you reply “I knew that already.” How is this case explained by the model introduced above? The informational ground for learning new vocabulary does not seem to include any beliefs or acceptances about potential uses of words’ last two letters. In other words, they are not introduced to the IC as relevant. As long as it remains not introduced as relevant, you might fail to answer the question, but if your IG for learning words introduced, for some reason, the last two letters as relevant, your context set would be updated accordingly, and you would be able to answer the question immediately with the words you know to be ending with “mt”. But, presumably, one’s IG regarding learning new words contains presuppositions about the purpose of use relative to meaning of the words. Therefore, any competent English speaker is most likely to immediately answer the question “What is the past participle of the word we use for seeing images usually in a form of narrative during sleep?” with “dreamt”.

CHAPTER V

WHAT CAN FRAGMENTED STATES BE?

Although the model I introduced motivates fragmented states in representing one's epistemic and doxastic states, one should be careful about what she means by a fragmented state. In this thesis, I avoid any metaphysical commitments regarding the possible nature of fragments of beliefs. My attempt is to provide a descriptively adequate and explanatorily fruitful model that enables one to represent belief and knowledge (and related dispositions and capacities) properly. In what follows, I will briefly discuss the possible ways the project of fragmentation can go and in what way I would like to limit my account. To do so, the objections in Aaron Norby's "Against Fragmentation" (2014) are valuable to consider.

5.1 Fragmented Sets of Beliefs

One natural way to understand fragmentation about beliefs is to say that fragments are "activated" or "triggered" when the relevant context is there for the agent. This is how fragmentation, in its brief literature, is described. (Lewis, 1980; Stalnaker, 1984; Egan, 2008) Norby attacks this way of understanding fragmentation. (2014) In this section, I will consider his attacks on what he calls "fragmentationalism" and discuss how my version deals with these criticisms.

Appealing to context sensitivity in information deployment seems to offer a solution to the puzzles about belief put forward in the first chapter. There are few things to say about this solution and its implications on a logic of belief and knowledge. One important thing is that since the model enables one to explain the dynamic process of information deployment during belief formation and retrieval, traditional logic has no problems in representing beliefs. This is to say that:

Every particular context has its own closed set of beliefs where there are no logically problematic phenomena such as contradictory propositions or deductive failure.

In more fine-grained epistemic terms, this model says more about dispositional and occurrent beliefs (or implicit and explicit beliefs) and clarifies important aspects of them. Given that information deployment is context-relative and it is the base of belief formation and retrieval, every particular circumstance has its own logically ideal set of beliefs. After the informational context goes away, there is no more reason to utilize that set of belief so agents drop them. In such picture, agents have the beliefs they form in certain contexts floating, so to speak, in their minds when the informational context no more exists. Information is there to bring to memory when the relevant context requires them. One, revolutionary, way to see that is to say that belief content is not subjected to any logical rule except for the logical constraints one has (e.g. a square circle), without the need to use them in a context that is posed by the particular IG. Logical rules dictate upon subjects' propositional contents only when the relevant context requires the utilization of a belief set.¹⁸ This context is always limited by the limits of one's memory, attention and other related cognitive capacities.

In that sense, human agents are fragmented. These fragments are not to be understood, however, as enduring cognitive entities once they are formed and "triggered" or "activated" under certain circumstances but "created" or "constructed" depending on the certain informational contexts an agent is in and "destroyed" or "dissociated" when the need goes away.

Another merit of such introduction of informational context is that informational ground is always open to further study under the light of the advancements of cognitive sciences and philosophy of mind. As I will conclude, my initial introducing of this model is only a first step in the study of possible elements consisting IG.

This formulation of fragmented states can easily reply to Norby's first rejection that fragments cannot be psychologically valuable or well-motivated. He simply argues that a "static" conception of fragmentationalism, where fragments reside in some mysterious way in one's mind and activated in a particular context cannot be true and has no value psychologically speaking. He appeals to the fact that in recent literature of cognitive psychology, findings suggest that contradictory beliefs can exist in the same psychological division. That is, such contradictory beliefs can exist in the same

¹⁸ A thorough discussion of how this requirement can be met requires another thesis long text.

division with regard to the cues or primes that trigger a retrieval. He writes;

Factors like the presence of a relevant cue, however, are ones that appear to be able to vary across items within psychological divisions, and so make it likely that beliefs within single divisions will give rise to Magellan-like phenomena. In general, the set of cues that elicit one memory will be distinct from those that elicit a second memory, even though there may be some cues that elicit both, and, critically, even though both memories may be stored in the same memory subsystem. (2014, p. 32)

In a static understanding of fragments, which is what is found in the classical literature of fragmentation, this is indeed a big problem. If recent findings in cognitive psychology suggest that actual cases of belief retrieval do in fact take place in that way, insisting on the view that static fragments being necessary to represent belief states becomes a vain position to take. In the model I introduced, however, there is no such claim. The fragments are created and destroyed, similar to the picture of common ground in David Lewis' *Scorekeeping in a language game* (1979b), in an ongoing process of information deployment relative to the changing context. This makes it possible for contradictory beliefs to be in the same psychological division or whatsoever. It does not make any commitments to the ways of coding information devoid of context. Rather, it provides an explanation of the ways in which coded information retrieved in particular contexts. This covers psychological cues and covers what Norby discusses in depth.

In that sense, the model introduced draws the limits of logical representation of belief and these limits are shaped by the informational context. If it must say something metaphysical about the fragments, the only thing it can say is that the nature of fragments will be clarified with the clarification of informational ground, which works hand in hand with recent findings in cognitive science.

The second attack Norby makes is to the view that fragmentation is a good way to describe beliefs as they relate to agent's capacities to act certain ways. In this understanding, "fragmentationalism" does not refer to any psychological base, but it

is only a good tool to describe and offer a general explanation for the conditions like Freshman.

The model I introduced solves this problem and all other relevant problems posing from Norby's initial attack including this second one. His attack is aimed at this static understanding of fragments that are activated depending on the context. Such static conception of fragments seems not to be able to say much to explain the ways in which an agent is disposed to behave given her belief. As long as an account to explain the dynamics of fragment activation is not provided, a theory involving fragmented states cannot explain the problem. What it can do, is only to refer to the same old problem and reformulate it.

As Norby writes, such fragmentationalism can claim at most that if you believe that p , then you will in some circumstances (those in which the belief or its fragment is activated) behave as if p . This is, however, is not enough to solve (the problem) and/or explain the place of beliefs as dispositions to act certain ways.

None of these attacks apply to this account of fragmentation with the model I propose based on IG. This model enables us to have a much stricter claim than Norby ascribes to the "static" fragmentationalists since this account makes it clear when and why one would behave as if p , given that she believes that p . This account also clarifies the place of belief as a propositional attitude and the difference between belief and different attitudes plays a role in behavior as it draws the line for the dynamic processes involving in belief formation.

Norby writes;

...it is not strong enough to distinguish the connection that belief has to reasoning and behavior from the connection that other, wildly different, states have. If one thinks, as many do, that it is part of the job of a theory of belief to pick out belief's role in reasoning/behavior, then this would be enough of a problem to justify a rejection of this version of the view. (2014, p. 35)

His point is that since the weaker claim that static notion of fragments can endorse does not distinguish between belief states and other similar states like acceptances, fragmentationalism cannot capture the role of belief in reasoning and behavior, which, in fact, was the aim of fragmentationalism in the first place. This is a completely fair point.

The account I provide, however, describes the role of all such attitudes for an agent's behavior. Recall that informational ground includes one's acceptances, presuppositions and beliefs together to limit the set of possible worlds compatible with them. It is never one single belief or acceptance that makes an agent behave in certain ways. Rather, it is informational ground and resulting informational context which makes one become disposed to behave that way. And other explicit acceptances such as actors have during a play, which, evidently, are not beliefs¹⁹, are not a problem for this account since this account can make the more powerful claim which makes it clear when and why a belief leads one to be disposed to behave in that particular way.

Having considered Norby's two main attacks to fragmentationalism, then, we can have a better understanding of my proposal of a new, dynamic notion of fragments of beliefs. The above-introduced account is able to make the stricter and more powerful claim:

If agent A believes that P, A behaves as if P given P's place in the present informational context resulting from A's informational ground.

P's place in IC is made explicitly clear by the model proposed in the previous chapter. Hence, both tasks of distinguishing beliefs from other attitudes like acceptance or pretense and explaining why and when would the particular behavior take place are accomplished.

¹⁹ This might not be as evident as it seems in the first sight, after all.

CHAPTER VI

CONCLUSION

Actual belief states are yet to be modelled properly, that is, without assumptions about logically idealized agents. In the first and second chapters, I formulated the task to achieve a proper model of non-ideal belief states. In the third chapter, I examined Elga and Rayo's access tables and concluded that access tables cannot do their purported job for two reasons; first they do not ensure descriptive accuracy, and, second, if they are modified to ensure the accuracy they become trivially constructed tables as their arbitrary characterization can match any type of behavior. Indeed, no merely descriptive attempt to represent the target cases can be successful as it is. It needs an additional model of the mechanisms at work during information deployment in order to stay at a coarse grained level. As a result, In the fourth chapter I introduced two characterizations of informational context, positional and presuppositional, where the former deals with concerns of descriptive accuracy of modelling actual belief states, and the latter offers a model of informational context based on informational ground that provides a general picture of the dynamic mechanisms at work.

Two things have been achieved. First, since the actual conditions are captured relative to the agent who is to deploy information by the positional character of IC, the threat of triviality by going fine-grained is evaded. Second, what information *would be* available relative to various purposes of an agent is explained through a general model of the agent's dynamic properties in an informational context.

Such understanding of belief can shed light on many important problems in the history of philosophy and provide a new perspective to the studies of the relation between thought, belief and rationality. Plus, it can feed empirically motivated studies in cognitive science that are about cognitive dissonance, implicit bias and self-deception since it offers a new and more elaborate conception of belief than the

way it is used as a term in folk psychology. And maybe more importantly, the proposed model can always renew itself (or abrogate itself when needed) with respect to the developments in cognitive science.

At that point I should state that I have proposed only an initial way of thinking about informational ground and its constraints on intentional content. What informational ground contains apart from the presuppositions regarding information source and environment, how it gets updated, how it can be modeled and other related questions require further promising research, which needs to take advantage of recent and traditional discussions in philosophy of mind, philosophy of language, and the developments in cognitive science.

A last thing to note is that, the view that informational context and informational ground are the basis of belief content are not merely about saving modal accounts of mental content that are based on the view that propositions are sets of worlds. It is true that the problem of logical omniscience poses a particular-and fatal-threat to the so called coarse grained theories of content. Nonetheless, the cases of irrationality and partial recall are creating problems for many important views about content in the history of philosophy from Frege (1948) to Davidson (2004). The most crucial outcome of this model is that it treats “beliefs” as not being subjected to logical rules without specifying the conditions to form new beliefs or make use of existing beliefs. Although that might seem to be relativizing belief content and belief ascription too much, it can capture too many things about mental content, that have not been captured by present accounts.

REFERENCES

- Abbott, B. (2008) Presuppositions and common ground. *Linguistics and Philosophy*, 31 (5): 528-538.
- Alchourrón, C. E., Gardenfors, P., & Makinson, D. (1985) On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50 (2): 510-530
- Crane, T. (2001) *The elements of mind*. Oxford University Press
- Crane, T. (2013) *The Objects of Thought*. Oxford University Press
- da Costa, N. & French, S. (1989) On the Logic of Belief. *Philosophy and Phenomenological Research*, 49 (3): 431-446
- Davidson, D. (2004) *Problems of Rationality*. Oxford University Press.
- Dretske, F. (1981) *Knowledge and the Flow of Information*. MIT Press.
- Egan, A. (2009) Billboards, Bombs and Shotgun Weddings. *Synthese*, 166 (2): 251-279.
- Egan, A. (2008) Seeing and believing: perception, belief formation and the divided mind. *Philosophical Studies*, 140 (1): 47-63.
- Elga, A., and Rayo, A. (2015). Fragmentation and information access. Available online. <http://www.princeton.edu/~adame/papers/fragment/fragmentation-and-information-access-2015-04-08.pdf>
- Frege, G. (1948) Sense and Reference. *The Philosophical Review* 57 (3): 209-230
- Greco, D. (2014) Iteration and Fragmentation. *Philosophy and Phenomenological Research*, 88 (1): 656-673.
- Grice, H. P. (1981) Presupposition and conversational implicature. In P. Cole (Ed.) *Radical Pragmatics*, (pp. 183–198). New York: Academic Press. Reprinted in H.P. Grice (Ed.) *Studies in the Ways of Words*, (pp. 269–282). Cambridge, MA: Harvard University Press (1989)
- Hawthorne, J., & Magidor, O. (2009) Assertion, Context, and Epistemic Accessibility. *Mind*, 118 (470): 377-397.
- Hintikka, J. (1962) *Knowledge and belief: an introduction to the logic of two notions*. Cornell University Press.
- Kecskes, I., & Zhang, F. (2009) On the Dynamic Relations Between Common Ground and Presupposition. In A. Capone et al. (Eds.) *Perspectives on Linguistic Pragmatics*, (pp. 375-396). Springer.
- Levi, I. (1977) Subjunctives Dispositions and Chances. *Synthese*, 34 (4): 423-455.

- Lewis, D. (1979a) Attitudes de dicto and de se. *The Philosophical Review*, 88 (4): 513-543.
- Lewis, D. (1980) Index, Context and Content. In S. Kanger, S. Öhman (Eds.), *Philosophy and grammar*. Dordrecht, Reidel. Reprinted in David Lewis, *Papers in philosophical logic*. (1998), (pp. 21–44). Cambridge, Cambridge University Press
- Lewis, D. (1982) Logic for Equivocators. *Nous*, 16 (3): 431-441.
- Lewis, D. (1979b) Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8 (1): 339-359.
- Lycan, W. G. (1986) Tacit Belief. In R. Bogdan (Ed.) *Belief: Form, Content, and Function*. Oxford University Press.
- Norby, A. (2014) Against Fragmentation. *Thought: A Journal of Philosophy*, 3 (1): 30-38.
- Pearce, D. (1995) Epistemic Operators and Knowledge-Based Reasoning: A Survey and Critical Comparison of some Recent Approaches in Philosophy and AI. In A. Lau, H. Wansing (Eds.) *Knowledge and Belief in Philosophy and Artificial Intelligence*, (pp. 1 – 27). Akademia Verlag
- Rescher, N. & Brandom R. (1980) *The Logic of Inconsistency*. Blackwell.
- Russell, B. *The analysis of mind*. Duke University Press.
- Schwitzgebel, E. (2001) In between believing. *The Philosophical Quarterly*, 51 (202): 76-82.
- Schwitzgebel, E. (2010) Acting contrary to our professed beliefs, or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, 91 (4): 531-553.
- Stalnaker, R. (1978) Assertion. in *Pragmatics*, ed. P. Cole, 315-332. New York: New York Academic Press.
- Stalnaker, R. (1999) *Context and Content*. New York: Oxford University Press.
- Stalnaker, R. (2014) *Context*. Oxford: Oxford University Press.
- Stalnaker, R. (1984) *Inquiry*. Cambridge: The MIT Press.
- Stalnaker, R. (1991) The problem of logical omniscience, 1. *Synthese* 89 (3): 425-440
- Seegerberg, K. (1999) Two traditions in the logic of belief: bringing them together. In H. J. Ohlbach, U. Reyle (Eds.) *Logic, Language and Reasoning*, (pp. 135-149). Springer.
- Stich, S. (1983) *From folk psychology to cognitive science: The case against belief*. The MIT Press.

Stich, S. (1993) *The Fragmentation of Reason*. New York: Oxford University Press.

Sutton, J, & Williamson, K. (2014) Embodied Remembering. In L. Shapiro. (Ed.)
The Routledge Handbook of Embodied Cognition, (pp. 315-325). Abingdon:
Routledge.