

Social Ontology and Social Normativity

Brian Joshua Donohue

January 27, 2020

A Dissertation Submitted to
The Faculty of the Graduate School of Arts and Sciences
The University at Buffalo, The State University of New York
In Partial Fulfillment for the Degree of
Doctor of Philosophy

Department of Philosophy

ProQuest Number:27738506

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent on the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 27738506

Published by ProQuest LLC (2020). Copyright of the Dissertation is held by the Author.

All Rights Reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Acknowledgements

My gratitude belongs first and foremost to Barry Smith, whose decades of innovative research in ontology, and whose feedback, patience, and support throughout this process, made this dissertation possible. I am also especially indebted to Werner Ceusters, whose critical comments on earlier drafts saved this work from many serious errors. I also thank Ryan Muldoon, whose feedback toward the start of this process positively influenced the shape of the work. Any residual errors are my own.

In many respects, this dissertation represents the consummation of a journey that began a dozen or so years ago, when a gawky freshman enrolled in a course entitled “Philosophy of the Human Person” with Theresa Farnan, whose passion for philosophy and love of wisdom was infectious even at the ungodly hour of 8am. Over the next four years at Franciscan University, I had the pleasure of engaging many historical and contemporary sources in philosophy under the tutelage of John Crosby, Jonathan Sanford, Paul Symington, and John White. At the University at Buffalo, I was likewise privileged to study with many distinguished faculty members, especially (in addition to those already mentioned) David Braun, Jorge Gracia, Alex King, Lewis Powell, Neil Williams, and above all the late, great Jiyuan Yu. To the present committee, and to all my teachers over the years, I can offer you little as recompense except Aristotle’s own sentiments toward those with whom he studied philosophy: “Their worth cannot be measured against money, and they can get no honour which will balance their services, but still it is perhaps enough, as it is with the gods and with one’s parents, to give them what one can” (*Nicomachean Ethics* 1164b3-5).

Many of the ideas and arguments in this dissertation developed in no small part thanks to discussion with fellow graduate students at Franciscan University, especially Nathan Cornwell and Mark Erste, and at the University at Buffalo, especially John Beverley, Botan Dolun, Shane Hemmer, Bob Kelly, David Limbaugh, Stephen McAndrew, Jake Monaghan, Justin Murray, Neil Otte, Fumiaki Toyoshima, and Jon Vajda. I am also indebted to audiences who raised questions about an early draft of some of these ideas presented at the 2017 International Conference on Biomedical Ontologies in Newcastle-upon-Tyne, and to the anonymous reviewers who provided comments. The bulk of this dissertation was completed while working concurrently as an ontologist for CUBRC, where I had many thought-provoking discussions on social and deontic ontology, especially with Alex Cox, Mark Jensen, Ron Rudnicki, and Yonatan Schreiber. The last year of writing and revising was completed during a transition to Capital One, where I received much support and encouragement, especially from Bethany Sehon.

This dissertation would also not have been possible apart from the support of many beloved relatives, especially Kathleen Donohue, Patrick A. Donohue, Caroline Donohue, Steven Mosher, and Vera Mosher. Nor would it have been possible apart from my church families at Ss. Theodore and St. George in Buffalo and St. Patrick in Bealeton, who have embodied for me and my family the love, joy, and peace of Christ. I am also especially grateful for the personal support and counsel I received from Fr. John Boddecker, Fr. Jason Vansuch, Fr. Patrick Cardine, and Sung Kim-Kubiak.

Finally, I would like to express my profound love and gratitude to Evangeline and Finian, two burgeoning philosophers and natural ontologists, who daily amaze me with their affection, joy, humor, curiosity, and creativity; and to Hannah, best of friends and companions, the one who's stuck with me through thick and thin, and the one to whom I lovingly dedicate this work.

Contents

| | |
|--|-----------|
| Acknowledgements | ii |
| Contents | iv |
| Abstract | vi |
| | |
| 1 Introduction | 1 |
| On Ontology | 3 |
| Bracketing Morality | 5 |
| Looking Ahead | 7 |
| | |
| 2 Searle, Social Ontology, and Deontic Powers | 9 |
| Social Institutions as Deontic Phenomena | 9 |
| Institutional Facts | 11 |
| Institutions, Conventions, Norms | 13 |
| Constitutive Rules and Status Functions | 17 |
| Searle on Deontic Powers | 20 |
| Hohfeldian Elements and Relations | 24 |
| | |
| 3 The Grounding-Anchoring Framework | 30 |
| Speech Acts and Social Acts | 33 |
| Searle's Taxonomy of Speech Acts | 38 |
| Speech Acts and the Grounding Problem | 40 |
| A First Look at Grounding and Anchoring | 42 |
| Grounding and Anchoring among Other Relations | 46 |
| Anchoring, or Why Grounding Is Not Enough | 49 |
| | |
| 4 Groups and Grounding | 54 |
| Different Ways to Ground a Deontic Phenomenon | 55 |
| Groups and Collective Acceptance | 63 |
| Collective Intentionality | 65 |
| What Groups Might (Not) Be Like | 66 |
| Anchor Groups and Frame-Domain Groups | 69 |
| Nested and Competing Frame Principles | 73 |
| Systems of Frame Principles | 75 |
| | |
| 5 The Dispositions That Bind | 81 |
| Speech Acts and Collective Acceptance | 84 |
| Language and Implicit Commitment | 88 |
| Collective Belief | 93 |
| Regular Behavior and Monitoring | 95 |
| Dispositions and Anchor Groups | 98 |

| | |
|--|------------|
| Are Anchor Groups Plural Subjects? | 100 |
| Anchor Groups and Group Variables..... | 101 |
| Putting the Pieces Together..... | 103 |
| 6 Toward Deontic Ontology..... | 106 |
| Freestanding Y Terms | 108 |
| Reference, Reification, and Deontic Phenomena..... | 112 |
| Toward an Alternative Proposal | 115 |
| Representation and Dependence | 117 |
| Representation and Content..... | 121 |
| Representation, Direction of Fit, and Prescriptive Content | 124 |
| Addressee, Purpose, and Force | 129 |
| Roles | 135 |
| 7 The Ontological Status of Deontic Phenomena..... | 139 |
| Representation, Dispositions, and Frame Principles | 140 |
| Power, Authority, and Roles..... | 144 |
| Action Regulation..... | 148 |
| Deontic Relations? | 151 |
| Freestanding Y Cases Reconsidered..... | 154 |
| Conclusion..... | 157 |
| Bibliography | 159 |

Abstract

Many recent accounts of the ontology of groups, institutions, and practices have touched upon the *normative* or *deontic* dimensions of social reality (e.g., social obligations, claims, permissions, prohibitions, authority, and immunity), as distinct from any specifically moral values or obligations. For the most part, however, the ontology of such socio-deontic phenomena has not received the attention it deserves. In what sense might a social obligation or a claim exist? What is the ontological status of such an obligation (e.g., is it an entity in its own right)? And how do people come to have social obligations or permissions in the first place? In this dissertation, I argue that such social-deontic phenomena can be accounted for ontologically in terms of the existence of shared prescriptive representational content that is backed by collectively held dispositions to monitor for compliance, and to punish (sanction, blame, chide, look unfavorably upon) those who fail to comply.

Introduction

On Monday, Belinda borrows a power drill from Larry in order to hang a shelf in her kitchen, and she promises him that she will return it by Friday. Because Larry lent Belinda the drill, Belinda is licensed to use it. Larry has the ability to lend the drill because, once upon a time, he walked into a hardware store and inserted a thin piece of plastic into a machine, which successfully initiated a transfer of bits and bytes from one server to another. If Belinda still has the drill when Saturday comes, Larry would be within his rights to call her up and ask why she hadn't yet returned it. And so long as she has the drill, Belinda is responsible for making sure that the drill is kept in good condition. Belinda is decidedly not licensed to abscond with the drill, but even if she were to do so, this would not thereby entail that the drill is among her rightful possessions.

In one respect, this scenario is perfectly mundane, and yet “its metaphysical complexity is truly staggering” (Searle, 1995, p. 3). If we were to enumerate some of the facts comprising Belinda's borrowing the drill from Larry, we could identify those like the following:

1. Belinda's throat, larynx, tongue, teeth, etc., produced sound waves and caused them to travel through some portion of air.
2. Larry's outer ear collected these waves, which caused his eardrums to vibrate through the ossicles to the cochlea, which sent neural signals through the cochlear nerve to his brain, which interpreted these sounds.
3. In reply, Larry's throat, larynx, tongue, teeth, etc., likewise produced sound waves.

4. Belinda's outer ear likewise collected these sound waves, etc., and thus understood Larry's response.
5. Belinda communicated her acknowledgement of this response to Larry.

Taken together, these facts are supposed to make it the case that two further *deontic* facts obtain about Belinda and Larry, i.e., facts about their obligations, rights, claims, requirements, duties, responsibilities, prohibitions, permissions, licenses, authority, and immunities. These two further facts are:

6. Belinda is obliged to return the drill by Friday.
7. Larry has a claim upon Belinda that she return the drill by Friday.

But as Hume (1740/1975) famously remarked, it is puzzling why any facts about what "is" (e.g., that Belinda's larynx helps produce sound waves) should entail anything about what "ought to be" (e.g., Belinda's obligation to return the drill).

One might suggest that these new deontic facts about Belinda's rights and responsibilities, or about Larry's claim upon her, are identifiable with the mental states brought about by their linguistic exchange, e.g., Belinda's *intention* to return the drill by Friday and Larry's *expectation* to receive the drill. But the leap from facts about Belinda's intentions to facts about her obligations is far from trivial. Surely, merely formulating an intention does not place one under an obligation. If the next day, Belinda intends to buy a drill, she is not thereby obliged to buy one. Further, she might be obliged to return the drill even if she doesn't intend to. If she does abscond with the drill, intending never to return it, that does not absolve her of the obligation to return it. Similar considerations might be brought forward regarding Larry's expectation: the mere fact that an agent expects something does not mean he has a claim upon it.

All this is a way of expressing puzzlement at the fact that, in virtue of uttering a few words, Belinda incurs an obligation to return the drill to Larry by Friday—but how can that be so? Even if we can explain why facts about language, intentions, etc., give rise to further facts about obligations and claims, a second puzzle presents itself: what exactly *is* an obligation or a claim? As Reinach remarked, if obligations and claims are entities in some sense, then they are not identifiable with any ordinary, concrete physical objects, nor do they seem to be identifiable with any mental states (Reinach, 1913/1983). Reinach ultimately argues that claims and obligations are *sui generis*, quasi-abstract entities that have heretofore been neglected by philosophers. But it may be that there are no simple, discrete entities, which are the referent of our terms “obligation” and “claim”. There are other possibilities, e.g., that phrases like “has an obligation” or “has a claim” refer to complex configurations of facts about beliefs, desires, intentions, actions, and so forth, without there being some further discrete entity, “the obligation itself.” In any case, our second puzzle is simply that, at the outset, it is not at all clear where to situate the phenomena of claim and obligation within an ontological picture of the world. Accordingly, throughout this dissertation I will refer to claims, obligations, permissions, etc., as “deontic phenomena” (which is meant to be neutral regarding their ontological status) rather than “deontic entities” (which suggests they are necessarily entities in their own right).

On Ontology

Both of these puzzles could be described as *ontological* in nature. Traditionally, ‘ontology’ has referred to the field of philosophy concerned with a systematic account of what exists, or of the nature of being itself.¹ Thus, Aristotle’s (1984) exposition of categories, Ibn Sina’s (1973) distinction between

¹ The term ‘ontology’ has its roots in the seventeenth century, when some philosophers drew a distinction between ‘*ontologia*’ and ‘*metaphysica*’. According to Wolff, ontology is a subfield of metaphysics concerned with being or essence, as distinct from the subfields of cosmology and natural theology. Since the introduction of ‘ontology’ into philosophical parlance, other definitions of ‘ontology’ have been proposed. According to the familiar Quinean view, ontology is concerned with which beings exist (e.g., whether abstracta exist), or more precisely which beings a scientific theory is

what something is (*mahiyya*) and that which exists (*al-mawjud*), Leibniz's (1989) monadology, Heidegger's (1962) analysis of *Dasein*, and Quine's (1948) criterion for ontological commitment are each, in their own way, examples of traditional philosophical ontology.

In recent decades, however, the term 'ontology' has also been coopted to mean a formal representation of what exists, which can be encoded in a machine-processable language and used to integrate heterogeneous data sources, represent facts about categories of entity, and enable machine-reasoning over data. In this sense, an ontology of mammals might be used to integrate diverse bodies of research in mammalian biology, or as a way of encoding facts about mammals generally (e.g., that mammals are vertebrates), or as a way of reasoning over data about mammals (e.g., John is a mammal, and therefore John is a vertebrate).

These two senses of "ontology"—its traditional philosophical sense, and the more recent computational sense—are not wholly unrelated (Guarino and Giaretta, 1995; Gomez-Perez et al., 2004; Arp et al., 2015). That is not to say that computer scientists or informaticians as such typically express an interest in philosophical approaches to ontological questions. Likewise, philosophers are seldom concerned with the practicality of philosophical ontology. And yet the aims of philosophical ontology overlap in several ways with those of computational ontology. Both philosophical and computational ontology are concerned with producing a catalogue of entities, what they are, and how they are related to one another. Both are concerned with how one might formalize a theory of entities in logically respectable ways. And in fact, many computational ontologists have drawn upon work in philosophical ontology, recognizing that the philosopher's concern for working out a principled

committed to. Naturally, however, this view is anything but universal. It differs not only from Wolff's, but also from that of Husserl, for whom the task of ontology is (in the words of Varzi, 2010, p. 3) "to lay bare the formal structure of all there is, whatever it is." That is to say, for Husserl ontology is first and foremost concerned with the most general or fundamental characteristics of beings, and with the most general laws or axioms which govern their relationships. Thus understood, ontology not only addresses questions about which types of entity exist, but also attempts to articulate a general theory of identity, parthood, dependence, etc.

ontological theory can enhance the quality of computational ontologies (Gangemi et al., 2002; Hoekstra et al., 2007; Arp et al., 2015). The result of bringing philosophical ontology and computational ontology together in this way has been called *applied ontology* (Munn and Smith, 2008).

Ultimately, this dissertation is meant as an aid in support of applied ontology research. My primary task is explanatory: it proposes a way of resolving these puzzles through developing an account of where deontic phenomena fit within an account of reality. Its secondary task is constructive: to lay the foundation for a formal representation of these phenomena which can be useful for encoding, integrating, and reasoning over data about them. My main thesis is that socially normative or deontic phenomena like claims and obligations are not entities in their own right (as Reinach among others maintains), but can be accounted for ontologically in terms of shared prescriptive representational content backed by collective dispositions to sanction (blame, chide, look unfavorably upon) those who fail to comply with those prescriptions.²

Bracketing Morality

Although this dissertation aims to develop a theoretical account of the nature and grounding of deontic phenomena, it will not discuss moral obligations, moral rights, etc. One reason for this is that whether there are any moral obligations is an extremely controversial topic in its own right. There are certainly ongoing disputes about whether there are universally binding moral norms, whether (or under what conditions) we would have access to such norms, whether moral norms are (or are

² To clarify, then, my thesis would exclude the marginal case of a hermit who follows a religious code that includes self-punishment. If there are no divine or angelic beings holding him accountable, then he would not in fact be under any socially grounded obligation to, e.g., pray, meditate, fast, make prostrations, or anything of the sort. He might feel compelled to do so, but individual compulsion does not amount to an obligation. Nor does the belief that one is obligated make one obligated. Note that my thesis leaves open the possibility that he is under some other kind of obligation (e.g. a moral or religious obligation that is grounded independently of the collective dispositions in his society). My thesis even leaves open the possibility that there is a non-social kind of obligation that is totally private (e.g., I promise myself and then chide myself should I fail to keep it).

grounded in) facts about the natural world, whether moral language (e.g., “ought”) is descriptive or expressive, and so forth. But this dissertation will touch upon none of those issues. It is solely concerned with what might be called by contrast *social-deontic phenomena*, the deontic phenomena that (very roughly speaking) exist in virtue of social factors such as collective agreement or group behavior.

A helpful analogy should shed light on this decision to limit the scope of the dissertation in this manner. A major tenet of natural law theory is that: *an unjust law is no law at all*. On one way of reading this tenet, some X cannot be a law unless X is morally justified.⁴ Thus, a law permitting slavery would not merely be morally unjustified; it would fail to qualify ontologically as a law at all. Accordingly, the claim that slavery is permitted by law would literally speaking be false. In the nineteenth century, however, the legal theorist John Austin advanced the idea that it is possible to recognize the existence of a law independently of whether that law is morally justified. Thus, slavery might not be morally permissible, yet it could be a legal fact *that slavery is permitted*. As Austin put it:

The existence of law is one thing; its merit or demerit another. Whether it be or be not is one enquiry; whether it be or be not conformable to an assumed standard, is a different enquiry. A law, which actually exists, is a law, though we happen to dislike it, or though it vary from the text, by which we regulate our approbation and disapprobation. (Austin, 1832/1995, p. 157)

Austin’s view has come to be called *legal positivism*, since it concerns not the ultimate normative status of some law, but what laws there are as a matter of empirical sociology. In this vein, it would be fitting to describe my venture in this dissertation as one rooted in *social-deontic positivism*. I will assume that it is very well possible to have a socially grounded obligation, permission, or power, which, from the point of view of a universal morality, might be overridden or otherwise negated by loftier obligations, permissions, or powers. All the same, such social obligations, permissions or powers can be described

⁴ Another way would be: a law has binding force only if it is also morally justified. But to make my point more directly, I will not discuss this alternative reading.

as items of social ontology, as phenomena which exist in virtue of the behavior or dispositions of social groups, or of widely adopted social institutions. Accordingly, in this dissertation, I remain agnostic about all questions of ultimate moral justification, and focus instead upon deontic phenomena as types of social phenomena.

Looking Ahead

The overarching argument divides into two main parts. The first part comprises Chapters 2 through 5. In this part, my focus is on the question of what grounds social-deontic phenomena, i.e., in virtue of what do they exist? I begin by surveying some existing accounts in social and deontic ontology, criticizing elements of these accounts, and taking other elements as points of departure for my own account. Specifically, I look at John Searle’s influential account of social ontology and aim to clarify or revise some of its key notions, such as: *social institution*, *constitutive rule*, *status function*, and *deontic power*.

Chapter 3 examines the claim that speech acts (e.g., promising) are responsible for grounding deontic phenomena, which I argue is only half-true. Instead, I propose adopting Epstein’s (2015) *Grounding-Anchoring Framework* as a better way to approach what grounds deontic phenomena, and especially to get a clearer picture of the diverse ways deontic phenomena can be grounded. In the remainder of Chapter 3, then, I present the distinction between *grounding*, a relationship between a more fundamental reality and a less fundamental one, and *anchoring*, a relationship between one or more entities (paradigmatically, a social group) and a principle which defines grounding conditions. Thus, the fact that an entity is money might be *grounded* in facts about its origins from an official printer, whereas the principle that paper printed in such-and-such a way counts as money is *anchored* in the collective acceptance of that principle (what Searle would call a “constitutive rule”).

Chapter 4 is meant to highlight the idea that there is no one way of grounding deontic phenomena (e.g., in commissive speech acts), and that a given social group might be—or very likely

is—home to nested and competing sets of principles. In Chapter 5, I argue that the anchoring of such principles bottoms out in collective acceptance, which I characterize as a collective disposition to sanction defectors from the system of principles.

The second part of the argument comprises Chapters 6 and 7 and focuses on the question of where deontic phenomena fit within an ontological picture of the world. Specifically, in Chapter 6, I focus on Reinach’s proposal, more recently defended by Crosby (1983), Mulligan (1987a), and Smith (2003a; 2003b; 2008) that deontic phenomena like claims and obligations are entities that belong to a *sui generis* category of “quasi-abstract entities” or, as Smith sometimes calls them, “freestanding Y terms”. After presenting a worry about this approach, I begin to lay the groundwork for an alternative view by looking to the notions of *representation* and *prescriptive content*, and of the specified ways in which agents are socially empowered to act, especially insofar as these ways of acting are connected to *social (or institutional) roles*. In Chapter 7, I put these pieces together to argue that deontic phenomena can be accounted for ontologically in terms of what I will call *shared prescriptive content*, which is backed by the aforementioned collective dispositions to monitor for behavior in accordance with a frame principle (or system of frame principles). I conclude by showing how this view addresses what Smith calls “the problem of freestanding Y terms”.

Searle, Social Ontology, and Deontic Powers

In recent decades philosophers have shown growing interest in questions of *social ontology*, i.e., questions about the nature of social groups, roles, conventions, practices, rules, and institutions, and questions about what in virtue of which such phenomena emerge and persist. Such theorists of social ontology commonly note that the structure of social reality is inherently *deontic*, i.e., concerned with social rules governing behaviors, duties, obligations, norms, permissions, freedoms, privileges, licenses, claims upon others, power over others, and immunity from others' interference. In order to understand these phenomena better, however, it is imperative to situate them in the context of a general approach to the ontology of social reality. Accordingly, the task of this chapter is twofold: first, to situate my approach within recent literature on the nature and grounding of deontic phenomena in social contexts (especially John Searle's), and second to suggest at least one way these accounts fall short, and thus motivate further investigation.

Social Institutions as Deontic Phenomena

Much of the recent discussion surrounding the ontology of deontic phenomena has occurred within broader discussions of the ontology of social groups and institutions. A typical account of social ontology will focus on the ways in which individuals can come together to form a group, and as a group, to jointly construct a set of social institutions governing their common life. Thus, the reason

that money, or marriages, or sports teams, or governments exist is that individuals have somehow collectively agreed to their existence, or rather, to certain rules which establish the circumstances in which they exist. To paraphrase Searle (1995, p. 32), money is money because we believe it's money; the reason why select green slips of paper count as money is (again, very roughly speaking) that a collection of individuals have agreed that certain green slips of paper count as money.

A typical account of social ontology will also point out that social institutions¹ seem to have an inescapably deontic element. In fact, an institution is often described as constituted by a set of *rules* governing interactions in particular social contexts. For example, one might easily cash out the institution of American football in terms of the rules which define the game: when and how many players are allowed on the field, when and how they are allowed to move, and so forth. Moreover, with social institutions come social roles for the people and objects which participate in those institutions, and with those social roles come permissions, obligations, and powers. Thus, the institution of football comes packaged with roles for individuals (quarterback, halfback, fullback, etc.), but one striking feature of these roles is that they are defined in terms of what the role's bearer is supposed to do (authorized to do, allowed to do, meant to do) in the context of that game. The role is defined by an associated set of activity types, and the exact scope of those activities are heavily regulated and backed by threat of penalty. Similar considerations apply equally to cases of inanimate objects with social roles, e.g., chess pieces. A piece of wood or ceramic's role as the queen-piece defines a set of permissible action types which can be taken by the possessor of that piece in the context of the chess game. To have a social role, then, seems to be an inherently deontic phenomenon.

There is another sense in which social institutions have been thought to be inherently deontic, namely, the idea that social institutions arise through commitments undertaken jointly by the members

¹ 'Institution' in the sense of a practice, custom, or set of interlocking relationships (e.g., baseball, money, marriage), not 'institution' in the sense of a campus or facility.

of the relevant social group (cf. e.g., Searle 2010, p. 80). This, of course, raises the difficult question of whether such an act of commitment—e.g., an act of promising—is itself a social institution, and thus anthropogenic. This topic will be discussed in much greater detail beginning in Chapter 3.

Institutional Facts

What exactly is a social institution? Although the meaning of ‘social institution’ is somewhat ambiguous, the term is typically used to refer to an overlapping set of social phenomena, roughly: the roles, conventions, rules, and rituals endemic to a collection of individuals (Hayers, 1911; Miller, 2014). Thus, typical examples of social institutions include a language, a body of laws, a system of exchange, a family structure, a corporation, or a game. Accordingly, many philosophers and sociologists typically take social institutions to be a—if not *the*—focal point of any analysis of social reality (Rawls, 1972; Giddens, 1984; Turner, 1997; Searle, 1995; Tuomela, 2003). Note, however, that this rough characterization of social institutions does not clearly delineate between an institution and a culture or society as a whole (Miller, 2014). Accordingly, ‘social institution’ should be understood as a term of art, referring to those discrete roles, conventions, rules, and rituals which help constitute a society or culture as a whole. Nonetheless, it remains possible to throw enough light on the nature of social institutions to proceed with the development of an account of the nature and grounding of deontic phenomena.

One common way of characterizing social institutions is by differentiating *institutional facts* from *brute facts* (Anscombe, 1958). Here a brute fact is understood to be a fact which is discovered and described through research in “physics, chemistry, evolutionary biology, and the other so-called hard sciences” (Searle, 2010, p. 42; cf. also, Searle, 2005). By contrast, an institutional fact is understood as any fact which obtains, roughly, because collections of individuals have created or accepted that fact. For instance, the fact that a physical object moved from one point to another at a certain speed would

be a brute fact, whereas the fact that that physical object is a football player, or that in moving to a certain location on the field he scored a touchdown, would be an institutional fact.

Some authors distinguish further between institutional facts and the broader class of *social facts* (Durkheim, 1964; Gilbert, 1989; Searle, 1995). For Searle, a social fact is any fact that obtains in virtue of a group of agents adopting collective intentional attitudes toward some activity (Searle, 1995, p. 28; cf. also, Gilbert, 1989). In this sense, the fact that two people are taking a walk together is a social fact. An institutional fact, in Searle's sense, is more specific than this, in that it involves the collective acceptance of a special kind of rule which makes the institutional fact possible in the first place (Searle, 1995, p. 34). Thus, the fact that one person checkmates another is an institutional fact, because the activity of checkmating requires that both agents have accepted the rule *that threatening a player's immobilized king-piece counts as checkmating that player*. By contrast, two agents taking a walk together does not require the collective acceptance of any such rule, nor is there any rule which would, in virtue of its acceptance, make the activity of walking together possible. Thus, it is a social fact, but not an institutional fact.

From this point of view, a social institution can be understood as the ontological result of accepting a rule, or more broadly of a set of roles, conventions, rules, and rituals. In accepting such roles, rules, etc., a group of individuals introduces facts which go beyond brute facts, and create new roles for entities and new rules governing behavior. Thus, the institution of fiat money can be understood as a community's acceptance of a role for certain types of paper bills and coins, and of rules governing behavior with those bills and coins. Furthermore, when social institutions are understood in this way, as being grounded in relatively more brute facts, a principal aim of social ontology comes into view: explaining what is necessary and sufficient for an institution to arise from the beliefs, behaviors, attitudes, or dispositions of the groups whose institutions they are.

Institutions, Conventions, Norms

Given this initial characterization of social institutions, it might be wondered how they are related to the notions of *convention* and *social norm*. After all, both conventions and social norms might likewise be characterized in terms of a group's adoption of rules for behavior. Like 'social institution,' terms like 'convention' and 'social norm' have become terms of art. In what follows, then, I am aiming not to sharply distinguish each of them, but to outline how these sibling notions might be teased apart from one another.

Conventions have received considerable attention in recent philosophical literature, not only as subjects of interest in themselves, but also insofar as they are relevant to issues in the philosophy of language (Carnap, 1937/2002), mathematics (Poincaré, 1898), metaethics (Harman, 1996), and philosophy of law (Sugden, 1998). But like 'social institution,' the meaning of the term 'convention' is somewhat ambiguous (Rescorla, 2017). As a first pass, a convention might be described as an (explicit) agreement between two or more agents. But this characterization of convention fails to account for conventions which evolved gradually through patterns of behavior (e.g., most linguistic conventions), rather than generated through explicit agreement. Accordingly, some authors have instead characterized convention as "a general sense of common interest" (Hume, 1777/1975; Lewis, 1969): knowledge shared by agents which can be entirely tacit, but which is nevertheless effective in producing behavioral regularity that promotes their common interest. In this sense, a convention can be thought of as a solution to a coordination problem (Lewis, 1969). But this Humean approach to conventions has also been criticized on various grounds, e.g., that the approach entails incorrectly that conventions are almost always followed (Gilbert, 1989), or that conventions can exist in the absence of common knowledge (Binmore, 2008), or that the approach overlooks the significance of sanctioning against non-compliance with a convention (Sugden, 1986/2004).

Disagreements over the nature of convention aside, the notion of convention which emerges from this literature is one of collectively accepted rules. Thus, there is certainly overlap between the concepts of a social institution and a convention. That being said, there are a few preliminary ways we can trace something of a distinction between social institutions and conventions. First, following Weirich (1989), Searle (1995), and Hodgson (2006), we can describe some conventions as particular instances of an institutional rule (Hodgson, 2006, p. 2). For example, states in the U.S. share a common institution of assigning political offices by vote, but it is a matter of convention whether a state (or even a voting district) uses paper or a machine as the mechanism for casting votes. Second, some examples suggest that the notion of convention is much broader than that of an institution, insofar as a convention can be a mere behavioral tendency within a population. For example, it might be a convention to design interior house doors to open inward, but this convention is in no way a social institution.

Social norms have likewise been a subject of keen interest among sociologists and philosophers. Intuitively, a social norm is understood to be an uncodified rule governing some behavior in a society. But this does not help very much in differentiating between norms on the one hand and institutions or conventions on the other. For example, North defines a social institution as “the constraints individuals impose on themselves” (North, 1990, p. 5) and “the rules of the game in a society” (North, 1990, p. 3), but it is not clear how such constraints or rules are supposed to be any different from conventions or norms. Thus, there is again considerable conceptual overlap here between social norms, social institutions, and conventions. That being said, many authors have attempted to distinguish more sharply between them. Thus, Tuomela (1995) maintains that norms are regularities in behavior accompanied by expectations, whereas institutions are the product of explicit agreement and involve the threat of sanction. But this account falls short on at least two counts: first because a

norm can prescribe action, rather than describe behavioral regularities, and second because norms can likewise involve the threat of sanction.

How, then, should we understand the relationship between social institutions, conventions, and social norms? At least for the purposes of this dissertation, I propose that we can borrow a distinction from Searle (1969, 1995, 2010) between two types of rule, namely, *regulative rules* and *constitutive rules*. By ‘regulative rule,’ he means a rule which somehow regulates a pre-existing activity. For instance, the rule of etiquette to place forks on the left-hand side of a plate is regulative, since it regulates a pre-existing type of activity, namely, dining. By contrast, he defines a ‘constitutive rule’ as a rule which not only regulates an activity, but also brings a new form of activity into existence. For instance, the rule for checkmating not only regulates how chess is played, but creates the very possibility of checkmating.² Apart from such a rule, there is no such activity as checkmating.

My proposal, then, is that social institutions, conventions, and social norms can be differentiated based upon how they incorporate regulative and constitutive rules. A social institution is a system of constitutive rules because institutional facts involve going beyond brute, physical reality, to the imposition of a new role or significance upon an activity or an associated object. For example, bills and coins receive a new status as *money*, and small pieces of wood or ceramic receive a new status of *king piece* or *queen piece*. In each case, the social institution is not a system of rules regulating antecedent forms of behavior types, but rather rules which invent new forms of behavior: purchasing, checkmating, etc.

By contrast, conventions and social norms can be characterized as two types of rules which largely emerge to regulate pre-existing types of behavior.³ Then, following Bicchieri (2006), we can

² For criticisms of the distinction between regulative and constitutive rules, see: Randsell, 1971; Giddens, 1984; Garcia, 1987; Tuomela, 2002; and Hindriks, 2009.

³ That is not to say that conventions and social norms cannot involve, or even rely upon, institutions as well. Since we are enmeshed in a world of social-institutional reality, many of our conventions and norms will arise within the context of

distinguish between conventions and norms in the following way: whereas conventions regulate behavior in a way which aligns with self-interest, social norms regulate behavior in a way which places a check on self-interest.⁴ For example, placing forks to the left of the plate would count as a case of convention; it is a regulation of pre-existing dining behavior which promotes one's own interests, but not at the expense of others' interests. On the other hand, the rule that one should be honest is a social norm because it is designed to curb my self-interest, which would often be served by dishonesty. Thus, if we understand conventions (following Lewis) as representing the equilibria of standard coordination games (Bicchieri 2006, p. 2), we can analogously view social norms as the equilibria of "mixed-motive games," or games in which there is "conflict of interest but also a potential for gain" (Bicchieri, 2006, p. 3).

I must mention, however, that although I find this distinction helpful in making sense of how conventions differ from social norms, it might not always be clear whether a given regulating rule would count as a convention or a social norm, or perhaps a complex mixture of the two involving both coordinating and conflicting interests. Ultimately, my goal in this section is not to give a comprehensive account, but to suggest a way in which to understand the relationship between accounts of institutions, conventions, and norms. Moreover, I wish to reemphasize that social institutions will invariably involve various types of further regulations, and manifest themselves in various particular conventions and norms. An institution may bring into existence a new form of behavior (e.g., playing chess), but it can then be further regulated (e.g., the convention of which players goes first).

social institutions, and the new forms of behavior they enable. The point remains that conventions and social norms are more regulative in character than constitutive.

⁴ Or as Bicchieri (2006, p.38) puts it: "Because conventions do not run counter to selfish motives, but social norms often do, if only empirical expectations were fulfilled, one would have a reason to follow a convention, but he would be seriously tempted *not* to conform to a social norm. In both cases, the players are playing a coordination game without nonstrict Nash equilibria, but whereas a convention solves an original coordination game, a norm transforms (with a certain probability) an original mixed-motive game into a coordination game and at the same time helps players to select one equilibrium."

Constitutive Rules and Status Functions

Searle first presented his account of institutional facts and constitutive rules in *Speech Acts* (Searle, 1969) as a way to explain how meanings are attached to symbols (glyphs, sounds). He famously formulates such constitutive rules as: “X counts as Y in C,” where X is a schematic variable standing for a non-institutional phenomenon (e.g., an object, property, or process), Y for its institutional status (e.g., money, a marriage ceremony, or checkmating), and C for the context in which the rule is operative. Thus, the fact that the utterance “I promise” counts as a promise obtains in virtue of the constitutive rule: utterances of the form “I promise” (X) count as putting oneself under an obligation (Y) in the context of English communication (C).

In *The Construction of Social Reality* (Searle, 1995), however, Searle develops his account of constitutive rules further, specifically, by focusing on the nature and origin of Y terms, which he rebrands as *status functions*. By ‘status function,’ Searle means to capture two dimensions of Y terms. The first is that of some functionality externally imposed upon an object (the “X term”) by some agent or group of agents, as in the case of the screw-driving function imposed upon the materials used to create a screwdriver (Searle, 1995, p. 13). The second is that of a status or significance which goes beyond the intrinsic physical makeup of the corresponding X term. In the case of a screwdriver, its function, although externally imposed by agents, is entirely exhausted by its physical makeup. Put otherwise, nothing beyond its “brute physical structure” is needed for it to having its function (Seale, 1995, p. 69). By contrast, in the case of status functions, intrinsic physical properties are not sufficient for a thing to have a function (Searle, 1995, pp. 13-23; 2010, p. 7). For example, the physical characteristics of a pair of utterances “I do”—the passage of air through the larynx, the vibration of vocal cords, the movement of soundwaves—are not sufficient for that utterance to count as the act of entering into a marriage. Rather, some status function must be imposed on those utterances when

uttered within the context of the marriage ceremony in order for them to perform their designated function.

In Searle's view, such status functions exist in virtue of being represented as existing, and he states in a few places that their representation as existing is in part constitutive of their existence (Searle, 1995, pp. 62-63, 68-69). But status functions exist not only because agents have the capacity to represent entities in mind and language, but also because they can represent them as *symbols* of other entities. Thus, when Searle first introduces the notion of a status function, he does so by arguing that certain physical objects can "represent," "stand for," or "symbolize" other entities. Note that the assignment of other types of function does not require representation of symbols. For example, a function of driving screws can be assigned to a screwdriver simply in virtue of representing brute, physical features of the screwdriver and the relations between those features and one's own purposes and goals. By contrast, "[m]arks on [a] paper now have meaning in a way that a screwdriver, for example, does not have meaning, because the marks on the paper now stand for or represent objects and states of affairs independently of themselves" (Searle, 1995, p. 21). Thus, for an X to be a Y requires representing some physical phenomenon (X term) as a symbol (Y term) of some further reality.

The symbolic character of status functions is most readily apparent in examples of language, e.g., the assignment of the status function of meaning *that snow is white* upon the sequence of glyphs 'Snow is white'. But Searle explicitly states that status functions in general have a symbolic character. Consider, for instance, Searle's example of a deteriorating wall, which in ages past served as a physical barrier to intruders, but which now serves as a symbolic barrier (Searle, 1995, pp. 39-40). The wall no longer provides a physical obstacle to others, but it does *indicate* or *symbolize* the authority of the people whose border it is over a certain geospatial region.

This is why Searle also claims in multiple places that peculiarly social objects (dollar bills, footballs, wedding rings, etc.) are ontologically secondary to the processes in which agents engage with them (acts of purchasing, playing football, and marrying), and that the point of constitutive rules is ultimately the creation, not of new social objects, but of new forms of activity. Social objects exist only to play an ancillary role in their corresponding processes, as placeholders or facilitating instruments. Thus, Searle writes:

[Social] “objects” are really designed to serve agentive functions, and have little interest for us otherwise. What we think of as social *objects*, such as governments, money, and universities, are in fact just placeholders for patterns of *activities*... Such material objects as are involved in institutional reality, e.g., bits of paper, are objects like any others, but the imposition of status-functions on these objects creates a level of description of the object where it is an institutional object, e.g., a twenty dollar bill. The object is no different; rather, a new status with an accompanying function has been assigned to an old object... but that function is manifested only in actual transactions; hence, our interest is not in the object but in the processes and events where the functions are manifested. (Searle, 1995, p. 57)⁶

In this passage, Searle is arguing that social activities or processes are in some sense prior to social objects because such objects exist only in order to play a useful role in associated social activities. Without the activities of purchasing or paying debts, physical tokens of exchange would be pointless. Moreover, such social activities are expressions of what agents may do, or must do, or must not do, or have the collectively recognized power to do. Thus, such social objects exist in order to serve as useful markers of what Searle calls *deontic powers*, i.e., the ways in which agents are regulated in their behavior or empowered to behave. For example, a plastic card is collectively recognized as a driver’s license because it is a useful way of tracking which individuals are permitted to drive vehicles. Likewise,

⁶ Similarly, (Searle, 1995, p. 36): “It is tempting to think of *social objects* as independently existing entities on analogy with the objects studied by the natural sciences. It is tempting to think that a government or a dollar bill or a contract is an object or entity in the sense that a DNA molecule, a tectonic plate, or a planet is an object or entity. In the case of social objects, however, the grammar of the noun phrases conceals from us the fact that, in such cases, process is prior to product. Social objects are always, in some sense ... constituted by social acts... A twenty dollar bill, for example, is a standing possibility of paying for something.”

a slip of paper is collectively recognized as money because it is a useful way for agents to exercise their power to acquire new property rights and pay debts.

Searle on Deontic Powers

Searle believes that institutional facts, status functions, and collectively recognized constitutive rules are intertwined with what he calls deontic powers. By ‘deontic power,’ he has in mind attributes such as (1) being *obligated* or *prohibited*, which he calls “negative deontic powers,” since they in some sense remove possible courses of action, and (2) being *permitted* or having *authority*, which he calls “positive deontic powers,” since they in some sense create new possible courses of action (Searle, 1995, pp. 84-5). In Searle’s most recent account, status functions invariably confer such deontic powers (Searle, 2010, pp.8-9, 23-4). It is not difficult to see why status functions and deontic powers would intuitively be jointly assigned to X terms. For example, to assign some person the role of president of the United States seems to imply the assignment of some deontic powers, e.g., the right to veto legislation (Searle, 1998, p.139). In his earlier work, Searle had considered the ascription of an honor as the marginal case of a status function with no associated deontic power (Searle, 1995, p. 85), but more recently, he has shifted to viewing honor as a kind of deontic power, i.e., the right to be accorded respect (Searle, 2010, p.24).

Searle defends his characterization of deontic powers as *powers* on grounds that “the core notion of power” is that if *A* has power over *B*, then *A* has the ability to intentionally compel *B* to act (Searle, 2010, p. 151). Deontic powers fit this mold to the extent that they influence agents to act in certain ways (e.g., performing a promised action, or refraining from interfering with another agent’s activity). That being said, Searle does not think that deontic power is reducible to brute physical power, intimidation, or violence (Searle, 2010, p.147). While he thinks that some deontic forms of power (e.g., political power) will necessarily involve the exercise of coercive force, deontic powers themselves

(including the deontic powers within political institutions) require something beyond non-deontic forms of power.

The reason why, he argues, is that deontic forms of power introduce “desire-independent reasons for action” (Searle, 2010, p. 23). Ordinarily, the desirability of some state of affairs drives human action. A collectively recognized deontic power, however, serves as a new source of motivation and explanation for action, which does not require that the agent (overtly) desires to perform that action. Again, it is not difficult to see why Searle’s view might be regarded as intuitive: a person might do something unpleasant because she promised and is thus obliged, rather than because she finds the action itself or its result desirable. Accordingly, Searle disagrees with the view that action in compliance with deontic powers is motivated solely by an agent’s desires, including the desire not to be punished for failing to comply. Again, in many cases, deontic powers are reinforced by threats of punishment or by physical coercion, but there is more to an agent’s motivation for compliance with a deontic power than fear of sanction or punishment (Searle, 2010, p. 88).

Searle’s view is not merely that deontic powers are incidentally related to status functions or social institutions, but rather that deontic powers play a central role in the construction of social reality. “It is because status functions carry deontic powers,” he writes, “that they provide the glue that holds human civilization together” (Searle, 2010, p. 9). One reason why deontic powers play such a role is that much of the social world consists of relationships between agents and new forms of activity, and deontic powers are largely responsible for controlling (constraining, coordinating) those relationships and activities (Searle, 1995, p. 84; 1998, p. 131). For example, the institution of baseball introduces new forms of activity (striking out, hitting home runs, etc.), but the coordination of these activities depends upon the ascription of deontic powers to the relevant participants (e.g., by means of the three-strike rule). Hence, Searle goes as far as to claim that “collective acceptance is itself a mechanism for the creation of power” (Searle, 1998, p. 139), that “all of institutional reality is one way or another

about power” (Searle, 1998, p. 140), and that institutions are essentially “enabling structures that increase human power” (Searle, 2010, p. 105). Thus, money exists to augment an agent’s purchasing power, marriages to create spousal responsibilities and privileges, and governments to introduce codified regulations into communal life.

To see more clearly why deontic powers play such an essential role in the construction of social reality, consider again the thesis that status functions have a symbolic character, and more specifically that what a status function symbolizes is something deontic. Returning to his example of the deteriorated wall, Searle remarks explicitly that the wall “symbolize[s] a deontic status beyond the physics” (Searle, 1995, p. 74). Likewise, when Searle argues that the “primary target” of an analysis of social ontology are “not the social objects, such as money, governments, and universities,” but rather “the agents who operate on and within those objects,” then we see more clearly that “the great divide in the categorization of institutional reality is between what the agent *can* do [i.e., positive deontic powers] and what the agent *must* (and must not) do [i.e., negative deontic powers]” (Searle, 1995, p. 100). That is, social objects are of secondary importance to social ontology because they are merely instruments which facilitate agents’ exercise of their deontic powers. The powers come first, the social objects second.

However, I believe that Searle’s characterization of permission, obligations, prohibitions, and authority as *powers* is somewhat problematic.⁷ To see why, we need to consider again that, according to Searle, deontic powers are genuinely powers, and that “the core notion of power” is that if *A* has power over *B*, then *A* has the capacity to intentionally compel *B* to act (Searle, 2010, p.151). The problem is that this characterization of power does not sit well with the notion that an obligation is

⁷ It is important to note that Searle does not merely believe that obligations or rights are somehow related to deontic powers; he thinks they *are* deontic powers. For instance, he writes: “Typical names of deontic powers are ‘rights,’ ‘duties,’ ‘authorizations,’ ‘requirements,’ ‘permissions,’ and ‘certifications’” (Searle 2010, p. 123). And: “I have identified this [i.e. deontic] power by listing names of general types, such as obligations, authorizations, permissions, and requirements” (2010, p. 145).

literally a power, even a “negative power.” Searle insists that “[t]he notion of power is the notion of a *capacity*,” that “[every power] names a capacity, or ability” (Searle, 2010, p.145), and that the concept of power is captured in a formula such as “X has the power (is able, has the capacity) to do A” (Searle, 2010, p.146). But obligation, even if it presupposes some capabilities on the part of the obligor, is not a capacity to act, but a requirement to (not) act.

Likewise, Searle is emphatic that power in general “involves the ability of an agent of power to get subjects to do what the agent wants them to do whether the subjects want to do so or not” (Searle, 2010, p.147). If we take this passage at face value, then Searle’s view seems to be that any power would in some manner endow its possessor with the ability to make another person do something. But Searle also describes an obligation as a negative deontic power possessed by the person who is obligated (the *obligor*).⁸ This implies, unhappily, that an obligor, in virtue of his obligation, has the ability to get the person to whom he is obligated (the *obligee*) to do what he, the obligor, wants. But this is backward. It is not the obligor who exercises power over another in virtue of his obligation. It would be truer to say that the obligee is the one who exercises some power over the obligor by getting the obligor to do what the obligee wants. Indeed, this seems to be Searle’s intention in characterizing obligations as negative deontic powers when he claims that, “if I make a promise to you, then *you do indeed have a deontic power over me*” (Searle, 2010, p.148). What Searle seems to mean, then, is that “*A*’s obligation” does not refer to a power *A* has, but to a power *B* has over *A*.

The point here is not merely terminological. If Searle only called them deontic powers as a stipulated way of referring to them as a collection of entities, there would be no issue. But Searle goes further in claiming that everything he places under the header of “deontic power” *is a type of power*. And herein lies some sloppiness in Searle’s account: by Searle’s own lights, negative deontic powers, such

⁸ For example: “If I have a negative deontic power such as an obligation to pay you the money I owe...” (Searle 2010, p. 148).

as obligations, cannot count as powers. They would be more accurately characterized as *lacking a deontic power* or *being subject to a deontic power*. Thus, Belinda's obligation to return Larry's drill implies that she lacks the authorization to keep it. Likewise, Belinda's obligation to return Larry's drill undamaged implies that, if she returns the drill damaged, then Larry has the power to impose a compensatory obligation upon Belinda (e.g., to replace the damaged drill with an undamaged one). Therefore, insofar as Searle's account of deontic powers implies that obligations and prohibitions are genuinely powers, it is in that respect mistaken and in need of revision.

Hohfeldian Elements and Relations

One way to begin to go beyond Searle's account of deontic powers is if we draw a distinction between *first-order* and *second-order* deontic phenomena. The distinction I have in mind can be traced back to the legal theorist W.N. Hohfeld (1919), who distinguished between the kinds of elements that constitute a legal right. According to Hohfeld, the meaning of 'right' is ambiguous in legal discourse, referring sometimes to a freedom or liberty to perform an action, and at other times to another agent's obligations toward the subject of the right. The root of this ambiguity surrounding the meaning of 'right,' he argues, is that particular rights are, in fact, molecular combinations of deontic elements, namely: (1) *permissions* to act, (2) *claims* upon others, (3) *powers* over others' permissions, and (4) *immunities* from others' power.⁹

These four deontic elements are related in several significant ways. In the first place, whereas the first two elements (permissions, claims) are concerned with whether an agent is regulated or constrained with respect to an action, the last two elements (powers, immunities) are concerned with

⁹ The first of these elements, *permission*, is sometimes called *privilege* or *license* or *liberty* or *freedom*. Most authors treat these notions as equivalent, though Steiner (1994) insists on the term 'license,' and Thomson (1990) draws a distinction between 'liberty' and 'privilege'. For the purposes of this argument, I will take a permission to mean that the subject of the permission is allowed to perform a type of act under certain circumstances.

the generation, modification, transference, or elimination of permissions and claims, or of other powers or immunities. Thus, we can think of the first two elements as *first-order* deontic elements, or those deontic phenomena that have to do with regulating our actions directly, and the last two as *second-order* deontic elements, or those deontic phenomena that have to do (mostly) with modifying our first-order deontic phenomena.

A *permission* can be thought of as the lack of a prohibition upon one's action.¹⁰ For example, if Patty has the permission (freedom, license, privilege, authorization) to drink beer, then (to the extent that she is so permitted) she is not prohibited from doing so. Thus:

A has a permission to φ iff A is not obliged not to φ .

A *claim* of one agent (the *claimant*) over another (the *claimee*) can then be thought of in terms of the obligation or prohibition the claimee has with respect to the claimant. To be a claimee is to owe something to one's claimant. Thus, the notion of a claim introduces the notion of a requirement or obligation to do something:

A has a claim that B φ iff B is obliged to A to φ .

In short, then, permissions and claims are deontic phenomena that have to do with what agents *may* or *must* do.

A *power*, then, can be characterized as the ability to generate, modify, transfer, or eliminate some agent's permissions, claims (and thus, obligations), or powers. For example, if Patty has power over Matty as his employer or supervisor, then (to the extent that she has power over him) Patty can

¹⁰ This sort of "definition by negation" is common in explications of the semantics of deontic logic. Typically, some deontic notions is taken as primitive (either obligation or permission), and then used to define the related deontic notions. For example, $\Box A$ is defined as "It is obligatory that A ," which then allows one to define $\Diamond A$ as "It is permissible that A ," where ($\Diamond A \equiv \sim \Box \sim A$), or "It is permissible that A if and only if it is not the case that it is obligatory that it is not the case that A ". Of course, the same strategy is used to define the notions of necessity and possibility in non-deontic modal logics.

generate, or modify, or transfer, or eliminate a new permission for him. That is, she is capable of authorizing him to perform a task, or require him to perform a task (i.e., capable of generating a claim upon him that he perform that task). Note, however, that a power can also be exercised in the generation, modification, transference, or elimination of other powers. For example, if Patty promotes Matty to manager, then she exercises her power not merely by permitting or requiring Matty to perform some tasks, but also invests Matty with power over others. Thus:

A has power over *B* iff *A* has the ability to modify *B*'s permissions, claims, powers, or immunities.

Lastly, an *immunity* can be characterized as a kind of resistance to others' power. For example, if Matty's right to religious freedom entails that he is exempt from working on the sabbath, then we can say that he has immunity against Patty's attempts to take away his permission to a work-free Friday night or Saturday morning, or to impose a requirement that he work during those times. Thus:

B has immunity against *A* iff *A* does not have the ability to modify *B*'s permissions, claims, powers, or immunities.

In short, then, powers and immunities are deontic phenomena that have to do with whether, or to what extent, agents can modify first- or second-order deontic phenomena.

We can sum up the relationships between first- and second-order deontic elements as follows:

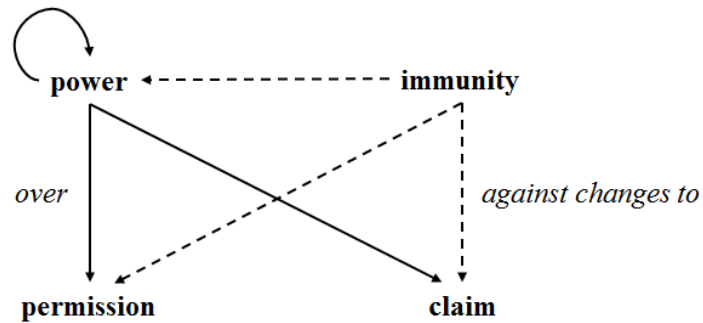


Figure 1 The four Hohfeldian elements

Note that to have a power is not merely to be licensed or unconstrained with respect to the exercise of that power. Certainly, if *A* has power over *B* to command *B* to perform an action, then to that extent *A* is permitted to order *B* to perform the action. But to say that *A* has power over *B* not merely to ascribe to *A* the permission to issue an order. It is, rather, to claim that *A*'s command would be binding upon *B*. Powers are more than permissions. For example, if Patty has the power to assign a task to her employee Matty, then she is certainly permitted to do so as well. But when Patty gains this power over Matty (e.g., when Matty signs an employment contract), what she gains is not a permission to do something she was previously forbidden from doing. If she had tried to assign a task to Matty prior to his employment, this would not be so much forbidden as ineffectual. Rather, when she gains the power to assign tasks to Matty, her commands are now binding upon him. She now has the ability to create obligations for him in the context of his employment.

The foregoing discussion of permission, claim, power, and immunity suggests that these four core deontic notions are interrelated with one another and with a few other key deontic notions (e.g., obligation). This, too, aligns with Hohfeld's analysis of core legal notions. In Hohfeld's view, these various deontic elements are variously related as *opposites* and *correlates*. For example, having a permission is the opposite of having a duty (obligation, requirement) and the correlative of someone else lacking a claim on the one who is permitted. We can depict these Hohfeldian relations in the

traditional way using two squares, one for the relationships between first-order elements, and the other for the relationships between second-order ones:

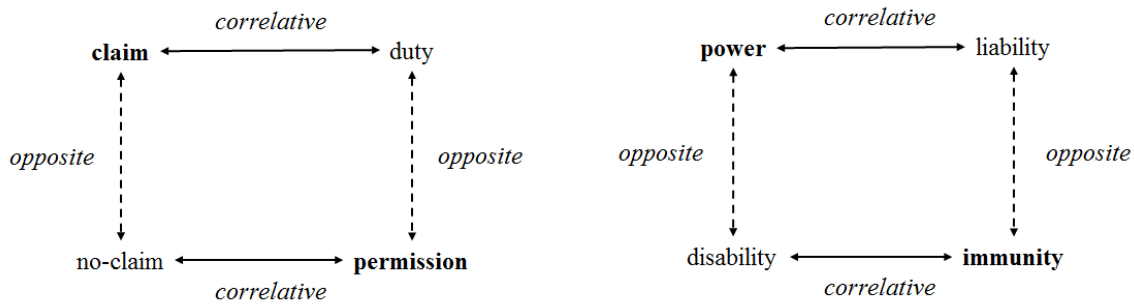


Figure 2 Hohfeldian relations of correlation and opposition

The relationships depicted above are somewhat complex and may require some revision.¹¹ But in the interest of explaining these relationships with a bit more clarity, consider the correlation between claims and duties. For Belinda to be the subject of Larry’s claim is for her to have an obligation to Larry, and to be prohibited from failing to perform the action. Likewise, her permission to use Larry’s drill is to be unconstrained with respect to that action. Thus, it is not difficult to see why, intuitively, claims correlate with duties. On this picture, to have a duty is for someone else (the claimant) to have a claim against the one who has the duty (the claimee), and likewise to have a claim against someone else (the claimee) is for that claimee to be obliged to the claimant. Moreover, since Belinda is obliged to Larry to return his drill by Friday, she is *not* permitted to keep it until Saturday. Thus, on the Hohfeldian picture, duty (obligation, requirement) can be characterized as the opposite of a permission (freedom, license, privilege, authorization).

¹¹ In particular, it is questionable whether every duty or obligation corresponds to a claim. Cf. Reinach (1913/1983, p. 12), who argues that there are some “absolute [i.e., not claim-relative] obligations”.

Deontic phenomena present at least two sets of metaphysical puzzles: what grounds their existence, and how to characterize the categories of deontic phenomena. Searle has developed a rich framework for social ontology which can serve as a point of departure for a closer study of deontic ontology. However, his account of deontic phenomena specifically as deontic powers runs into trouble. In the first place, as we have seen already, his account is imprecise in its characterization of obligations as negative deontic powers. But in light of Hohfeld's analysis, we can also see that Searle's may obscure an important distinction between first- and second-order deontic phenomena. So, given a generally Searlean approach to social ontology, how should we understand the ontology of deontic phenomena? The starting point, I believe, is to take a closer look at the notion of collective acceptance. Previously we saw that collective acceptance is invoked by Searle among others as the mechanism for grounding social institutions. But at the outset, Searle's account raises two main questions: What exactly is collective acceptance, and what does it mean for a constitutive rule to hold "in a context"? The following three chapters propose ways of answering these two questions.

The Grounding-Anchoring Framework

Consider, again, the case of Belinda borrowing Larry’s power drill. Why is Belinda obliged to return Larry’s drill? Larry would lend the drill only if Belinda promised to return it, and so she promised. But why does her promise place her under an obligation? The answer to this question is a matter of some dispute. In the first place, there is a broad point of disagreement over whether promising is in some sense artificial (conventional, anthropogenic) or natural (discovered, rather than invented, by human beings). Famously, Hume (1740/1975) argues not only that the practice of promising is a human convention, but also that the requirement to keep promises is a human convention, which derives its deontic force from the fact that as a whole the convention promotes the interests of each party by providing a means of assurance against those who might be tempted to break their resolutions.¹ Thus,

¹ At least, this is the traditional interpretation. Hume explains the origins of promising as follows: “Men being naturally selfish, or endow’d only with a confin’d generosity, they are not easily induc’d to perform any action for the interest of strangers, except with a view to some reciprocal advantage, which they had no hope of obtaining but by such a performance. Now as it frequently happens, that these mutual performances cannot be finish’d at the same instant, ’tis necessary, that one party be contented to remain in uncertainty, and depend upon the gratitude of the other for a return of kindness. But so much corruption is there among men, that, generally speaking, this becomes but a slender security; and as the benefactor is here suppos’d to bestow his favours with a view to self-interest, this both takes off from the obligation, and sets an example of selfishness, which is the true mother of ingratitude... In order, therefore, [to give greater assurance] there is a *certain form of words* invented... by which we bind ourselves to the performance of any action. This form of words constitutes what we call a *promise*, which is the sanction of the interested commerce of mankind... [Promises] are the conventions of men, which create a new motive, when experience has taught us, that human affairs wou’d be conducted much more for mutual advantage, were there certain symbols or signs instituted, by which we might give each other security of our conduct in any particular incident. After these signs are instituted, whoever uses them is immediately bound by his interest to execute his engagements, and must never expect to be trusted any more, if he refuse to perform what he promis’d.” (Hume, 1740/1975, 3.2.5.8-9)

Hume scoffs at the idea that the act of promising in itself could be responsible for generating, as if by magic, an obligation to keep the promise. Such a magical act he seemingly ridicules as

one of the most mysterious and incomprehensible operations that can possibly be imagin'd, and may even be compar'd to *transubstantiation*, or *holy orders*, where a certain form of words, along with a certain intention, changes entirely the nature of an external object, and even of a human creature. (Hume, 1740/1975, 3.2.5.14)

Searle (1969) likewise adopts the view that the practice of promising is anthropogenic, but then argues that, once the practice is established, there is no further need to explain why one is obligated to keep one's promises. The reason is that promising is a social institution, governed by a constitutive rule, roughly: "Saying 'I promise to φ ' counts as placing oneself under an obligation to φ ."² Thus, once a group collectively accepts this constitutive rule, the promissory act automatically and necessarily binds its speaker to the promised course of action (cf. Prichard, 1968). Reinach (1913/1983), however, seems to adamantly deny that promising is in any sense artificial or conventional. Rather, promising is a social act which any agent is capable of performing, and whose power to generate obligations is both intrinsic and discoverable through *a priori* reflection. Likewise, some legal theorists (Hart, 1955; Raz, 1975, 1977; Shiffrin, 2012) and ethical theorists (Owens, 2006, 2008) have adopted the view that promising is an inherent normative power grounded in our interests rather than in convention or agreement (Habib, 2018), or that the obligation to keep one's promises derives from moral principles governing what people owe each other (Scanlon, 1998).

From this cursory survey, we can glean two distinct questions about the ontology of deontic phenomena:

² The actual constitutive rule would obviously be more complicated. For example, it would need to specify that saying "I promise" generates an obligation only if the one to whom one is promising recognizes and agrees to the promise. Searle aims to enumerate such conditions in (Searle 1969, ch. 3).

- (1) In virtue of what do particular deontic phenomena come into existence (e.g., Belinda's obligation)?
- (2) In virtue of what do deontic practices (institutions, rules) exist (e.g., the act of promising)?

For Hume, the answer is that both exist in virtue of convention. For Searle, the first exists in virtue of something convention-like (or more precisely, because of an institution), but the latter exists because of implicit linguistic commitments. Thus, it should be evident that, in the view of many authors, deontic phenomena have their origin in some form of *agreement* or *commitment*. This is obvious in the theories of Hume and Searle. But even in Reinach's view, commitment figures prominently into his account of the genesis of obligation and claim through the social act of promising. It should be noted that some form of agreement or commitment has been invoked at various levels of analysis in theories of social ontology. Thus, some have argued or suggested that all obligations, permissions, etc., are grounded by commitment (e.g., an explicit or implicit act of commitment), whereas others have argued that social institutions are grounded by commitment (i.e., commitment to that institution), whereas others have argued that the existence of a social group which accepts an institution is grounded by commitment (i.e., commitment as a group to shared values or goals).³

In short, one common approach to grounding deontic phenomena is to appeal to commitment, and especially to commissive uses of language. Thus, the answer to question 1 is to appeal to acts of promising as the origin of social obligation. And in Searle's case, the answer to

³ For example, Gilbert (e.g., 2013, pp. 7, 85-88, 341-352) distinguishes mere aggregates of individuals from groups proper, whose existence, she argues, depends on individuals mutually committing to common norms, goals, actions, etc. To give a simple example of the view, two parties must jointly commit to enter into a marital contract in order to form a married couple. Borrowing Searle's terminology, we could say that such an act of joint commitment requires aggregates of individuals with the capacity for collective intentionality, and even an expression of readiness to form a group, but not the existence of a group with its own collective beliefs or intentions. So, on this view, it is not that the individuals must paradoxically be a group before they can commit to becoming a group, but rather that they must mutually commit to common goals, norms, actions, etc., before they form a group.

question 2 is similarly that all social institutions presumes commitment to performing speech acts correctly (e.g., in saying “I promise” to sincerely commit oneself).

I believe that this common, speech act-focused approach to grounding deontic phenomena is only half-true. But in order to argue this point, I will first need to introduce a somewhat different framework, which borrowing from Epstein (2015) I will call the *Grounding-Anchoring Framework*. The main task of this chapter is simply to present and defend this framework, which will then serve as a powerful tool in the following two chapters for answering the question of what grounds deontic phenomena. Before presenting that framework, however, I will take a closer look at the claim that speech acts are in some sense responsible for grounding.

Speech Acts and Social Acts

Many authors have remarked that the emergence of deontic phenomena is somehow related to the performance of special kinds of language act, which are responsible for constructing or maintaining social facts. Such acts clearly seem to have deontic ramifications. When Larry agrees to lend Belinda the power drill, she gains a permission to use it, as well as an obligation to return it by Friday; and likewise, Larry gains a claim upon Belinda that she return it by Friday. In this case, the language act itself—what are typically called *speech acts* or *social acts*—is apparently responsible for the new deontic fact. Belinda’s promising makes it thus that she is obligated.

The first philosopher to recognize the peculiar character of such speech acts was Thomas Reid (1710-1796), who drew a distinction between what he called “solitary acts” (e.g., thinking, judging, sensing, and feeling) and “social acts” (e.g., questioning, commanding, promising, and testifying), the latter of which are “as capable of being analysed as a proposition [i.e., an assertion] is.” (Reid, 1785/2002, p. 74; cf. also, Reid, 1788/2010; Smith, 1984; Schumann and Smith, 1990). In Reid’s view, social acts are not reducible to or composed of solitary acts because the former are necessarily

communicative or linguistic, and thus public rather than private. For instance, there is such a thing as private and tacit judgment, but there is no such thing as private and tacit testimony. Moreover, such social acts are essentially interpersonal. A command cannot be given unless somebody is the addressee of that command. Likewise, a promise cannot be made unless the addressee of the promise acknowledges and accepts the act of promising. Thus, Reid argues that there is a special class of acts, which are communicative or linguistic in some sense and which occur exclusively within social contexts.

C.S. Peirce (1839-1914) likewise highlighted the distinction between linguistic expressions whose purpose was to make assertions about the world and those whose purpose was otherwise. Peirce thought this distinction could be discerned if we considered the various ways in which we might approach a single proposition: “One and the same proposition may be affirmed, denied, judged, doubted, inwardly inquired into, put as a question, wished, asked for, effectively commanded, taught, or merely expressed, and does not thereby become a different proposition” (Peirce, 1904/1976, p. 248; cf. also, Brock, 1981). Thus, although it is tempting to express propositional content in the indicative mood (e.g., Socrates is wise), and thus suggest that the proposition in question is being asserted or affirmed, it is more accurate, Peirce argues, to represent propositions in a neutral form (e.g., *that* Socrates is wise or *Socratem sapientem esse*), since one and the same proposition can serve as content for a variety of propositional attitudes. Crucially, among the propositional attitudes Peirce lists, not all can be construed as solitary acts of the mind. Some, rather, are unmistakably social in Reid’s sense (e.g., inquiring, requesting, and teaching). Thus, Peirce too put his finger on one of the significant insights of subsequent speech act theory, namely, that (1) assertion is but one of language’s functions and (2) some ways of treating propositions or linguistic utterances are irreducibly social.

Neither Reid nor Peirce, however, expounded these observations into mature theories of speech acts. The next significant figure to touch this topic was Adolf Reinach (1883-1917) in his 1913

treatise *The A priori Foundations of the Civil Law (Die apriorische Grundlagen des bürgerlichen Rechts)* (Reinach, 1913/1983). The main purpose of Reinach's treatise was to argue that certain fundamental deontic entities exist independently of and prior to any positive law, that they exist as objective features of the world, that they are governed by certain necessary (or essential) laws or principles, and that both the deontic entities and the principles governing them are discoverable through *a priori* reflection. Reinach has in mind phenomena such as obligations, claims, and property, and principles such as the causal principle that an act of promising creates an obligation on the part of the promisor and a claim on the part of the promisee. In analyzing the act of promising, Reinach characterizes promising as a specifically and essentially social act, and is thus led to articulate a theory of social acts generally.

In the first place, then, Reinach holds that no positive law creates these categories of deontic entities (Reinach 1913/1983, pp. 2-4). Rather, there is first such a thing as, e.g., a right to one's property, and subsequently the regulation of property rights and exchanges through this or that property law in a given society. Thus, Reinach holds that deontic entities are objective, albeit metaphysically peculiar, features of the world. In seeking to explicate the nature of such deontic entities, Reinach argues that they emerge from special types of act called *social acts (die sozialen Akt)*. In particular, he focuses on the case of promising, which provides a clear example not only of a situation in which deontic entities are operative (namely, the obligation to keep one's promise), but also of a necessary law which governs deontic entities (namely, the law that an act of promising creates an obligation on the part of the promisor and a claim on the part of the promisee).

But to understand the nature of promising correctly, he argues, we must first understand the nature of social acts generally, of which promising is but one type. In doing so, Reinach seeks to provide an alternative to what he sees as the errors of his predecessors' theories of promising and deontic entities, namely, that promising is a mere convention invented by society (Hume), that deontic entities are projections of one's own psychological states (Lipps 1899, p. 152; 1909, pp. 208ff.), and

that deontic entities are reducible to other agents' expectations of how the obligor will act (Schuppe 1881, pp. 304ff.). Underlying all these theories, Reinach thinks, is the false view that an act of promising is a report of one's private mental states, akin to an expression of intention. Like subsequent pioneers of speech act theory, Reinach argues to the contrary that the purpose of many uses of language is not to make assertions about matters of fact (e.g., the fact that one intends to act), but rather to perform peculiar sorts of language-laden actions (e.g., making it the case that one is obliged to act by the utterance "I promise").

What, then, are social acts? Reinach (1913/1983, pp. 18-21) suggests that they have the following features. To start, they are marked by intentionality (Brentano, 1874/1995; Husserl, 1900/1984) and originate from the agent (Reinach calls them "spontaneous acts," *spontane Akte*), rather than occur passively within him. Moreover, they are necessarily directed at and addressed to somebody else. For example, in Reinach's view, an agent cannot literally make a promise to himself, or literally ask himself a question, or literally make a request of himself, or literally issue a command to himself. One might utter the words, "I promise myself I won't do that again," but Reinach maintains this would not be a genuine promise, but merely an expression of intention. Moreover, in at least the vast majority of cases, social acts must be expressed in some manner, e.g., in "mien, gestures, or words" (Reinach, 1913/1983, p. 20).⁴ Social acts must also necessarily be received in some manner by the addressee, or what later authors will call the *uptake* of a speech act. For example, an act of questioning has not occurred unless the addressee of the act has heard the question and understood it as a question. Of course, such uptake does not require that the addressee of the social act respond to the question or obey the command. Finally, Reinach thinks that, although it would be a mistake to think of social

⁴ Reinach considers a thought-experiment in which social agents have access to each others' private mental lives. In such a scenario, there would be no need to express oneself in language in order to perform a social act. This sort of scenario is considered (not so much hypothetically as eschatologically) by Augustine in *The City of God*, XXII, 29. Reinach also mentions prayer as an example of a potentially silent social act (Reinach, 1913/1983, pp. 20-21).

acts as reports about one's private mental acts, they nonetheless have an inner component (a private, mental, intentional state) that is coupled to the outer component (the mode of expression). For example, the act of requesting has as its inner component the desire that the requested state of affairs obtain. Divorced from this inner component, an act of requesting would become a "pseudo-request" (Reinach, 1913/1983, p. 22).⁵

However, Reinach's theory of social acts went largely unnoticed within Anglophone philosophy.⁶ Instead, J.L. Austin (1911-1960) and John Searle (b. 1932) were largely responsible for popularizing speech act theory among contemporary philosophers. At the heart of speech act theory is the now commonplace idea that there are some forms of language use whose purpose is not to describe the world as it is, but to change the way the world is, which Austin dubbed *performatives*. But the focus of Austin and Searle's early work in speech act theory was not primarily on the ontological status of deontic phenomena, but rather on the nature of language (e.g., its principal content vs. its effects on the addressee) and the various kinds of language use (e.g., descriptive vs. performative). That being said, both philosophers recognized the important connections between certain speech acts and deontic phenomena. For example, they noted that, assuming certain typical conditions hold, the utterance that a couple is now married engenders new spousal responsibilities and rights, and the promise to repay a debt by Friday entails an obligation to repay it by Friday.

⁵ Reinach notes in passing that it is unclear whether an insincere promise (or "pseudo-promise") creates an obligation. The same question arises in Austin (1962) and Searle (1969), who distinguish between infelicities which depreciate the act and those which nullify the act entirely.

⁶ There are likely two reasons for this. First, there was no English translation of Reinach's treatise until Crosby's in 1983. Second, the early realist phenomenology of Husserl, of which Reinach was a leading proponent, was abandoned by Husserl himself in 1913 and quickly overshadowed in the interwar and post-war years by Heidegger and his philosophical progeny. It is not known for certain whether Reinach exerted any influence on Austin's development of speech act theory. However, as Smith (1992, n. 10) notes, Austin's colleague, Gilbert Ryle, was in possession of an annotated copy of Reinach's complete works (*Gesammelte Schriften*).

Searle's Taxonomy of Speech Acts

Searle (1976; 2010) has provided a helpful five-category taxonomy of speech acts, which rests on a distinction between what he calls two “directions of fit” which an utterance might have. Thus, an utterance might be such that its content is meant to conform to the way the world actually is; then, we can say that the act has a “word-to-world” direction of fit (symbolized as ↓). Conversely, an utterance might be such that the world is meant to conform in some sense to the content of the utterance; then, we can say that the utterance has a “world-to-word” direction of fit (symbolized as ↑). Accordingly, Searle differentiates:

1. ↓ *Assertions* (or *Representatives*), whose purpose is to describe the way the world is, was, or will be. For example, the purpose of the utterance “San Diego is a city in California” is to represent in linguistic symbols the fact that San Diego is a city in California.
2. ↑ *Directives*, whose purpose is to induce the addressee to act in accordance with the Directive. For example, the purpose of the utterance “Go to San Diego” is to induce the addressee of the utterance to go to San Diego.
3. ↑ *Commissives*, whose purpose is to bind an agent to a future course of action. For example, the purpose of the utterance “I agree to go to San Diego” is to bind the speaker to the act of going to San Diego.

In some cases, however, Searle notes that an utterance has no direction of fit, but rather a presupposition (as opposed to an assertion) of some truth, symbolized as (P). For example, the utterance, “Congratulations on your promotion,” does not assert that the addressee has received a promotion, though it does obviously presuppose that she has been promoted. Thus, Searle also introduces the category:

4. (P) *Expressives*, whose purpose is to convey the speaker's feelings. For example, the purpose of the utterance “Thanks for sending me to San Diego” is to express the speaker's gratitude to the sender for sending him to San Diego. Note that it presupposes the speaker has been sent to San Diego.

Lastly, Searle introduces a category of speech act which has both direction of fit at once (↕):

5. \updownarrow *Declarations*, whose purpose is to effect some new institutional fact. For example, the purpose of the utterance “I hereby appoint you to head the San Diego branch” is to make it the case that the addressee now has the role of leading the San Diego branch.

A Declaration has both directions of fit at once because it aims simultaneously to transform the world to conform to its content (world-to-word) and to represent the world as being in that condition (word-to-world).

Searle (e.g., 1983, pp. 6-7) also notes that (most) types of speech act correlate to a “pre-linguistic” intentional state (e.g., belief or desire). Specifically, Assertions correspond with *beliefs*, Directives with *desires*, Commissives with *intentions*, and Expressives with *emotions*. For example, much as Assertions have a word-to-world direction of fit, beliefs have a mind-to-world direction of fit; and much as Directives have a world-to-word direction of fit, desires have a world-to-mind direction of fit. However, Searle maintains that Declarations have no such pre-linguistic analogue. Specifically, he argues that whereas it is possible to have beliefs, desires, intentions, and emotions without language, a Declaration presumes a linguistic capability, i.e., to regard some entities as symbols with meaning.

In light of Searle’s taxonomy of speech acts, we can ask the question of which types of speech act are of relevance for developing an account of the ontology of deontic phenomena. I believe three of Searle’s categories fit this bill: Directives, Commissives, and Declarations. Whereas Assertions describe how one takes the world to be, and Expressives convey how one feels about how one takes the world to be, Directives and Commissives are both tied to what an agent is being prescribed or required or allowed to do. Declarations are responsible for changing an agent, object, or process’s status in the social world, including whether they have been required or permitted to do something.

Unsurprisingly, in Searle’s view, language plays a crucial role in the emergence of social institutions. Specifically, he holds that the generation and maintenance of social institutions requires Declarations or Declarative-like representations. Recall that, for Searle, social institutions come into existence because of collectively accepted constitutive rules of the form X counts as Y in C, where Y

stands for some feature of X which goes beyond its intrinsic, physical makeup. But Searle also argues that social institutions come into being through either Declarations or else other forms of representation that share the same logical form as Declarations:

With the important exception of language itself, all of institutional reality, and therefore, in a sense, all of human civilization, is created by speech acts that have the same logical form as Declarations. Not all of them are, strictly speaking, Declarations, because sometimes we just linguistically treat or describe, or refer to, or talk about, or even think about an object in a way that creates a reality by representing that reality as created. These representations have the same double direction of fit as Declarations, but they are not strictly speaking Declarations because there is no Declarational speech act. (Searle 2010, pp. 12-13)

Thus, even in cases where there is no explicit Declarative speech act, Searle takes the representation of social institutions to be quasi-Declarative. A constitutive rule, then, can be thought of as a way of generalizing and perpetuating such Declarations. Thus, Searle concludes, all social institutions seem to require at least the capacity for Declarations or Declaration-like mental representations, and thus the capacity for language.⁷

Speech Acts and the Grounding Problem

The speech act-focused approach to grounding deontic phenomena faces one major challenge, which I will call the *Grounding Problem*. In short, proponents of the view that speech acts are responsible for the emergence of deontic phenomena face the difficulty of explaining why language acts, like uttering

⁷ To be more precise, Searle maintains that all social institutions require the *institution* of language, i.e., not just a capacity for language, but also a set of shared linguistic commitments. However, Searle qualifies his view by arguing that the institution of language cannot develop through Declarations because Declaration presupposes the institution of language. Rather, he proposes an evolutionary account, according to which humans first gain the capacity to represent the world using symbols, which is done idiosyncratically to convey “speaker meanings” (Searle, 2010, pp.73-75). Eventually, some speaker meanings are subsequently collectively adopted as conventional ways of communicating meanings to others. Thus, language begins as an individual capacity for representation, not as a social institution. It becomes an institution only if certain linguistic or symbolic ways of representing the world are collectively accepted as normative ways of conveying meaning. Thus, although the institution of language as a system of collectively accepted conventional meanings, and a commitment to certain linguistic practices, is required for other social institutions, it is not required for the institution of language itself.

“I promise,” have deontic effects. The choice here seems to be between a conventionalist approach, which maintains that speech acts’ deontic effects are a matter of agreement, and a non-conventionalist approach, which maintains that they are not. Now if one adopts the former approach, then one is faced with the question of why any original agreement or convention carries deontic force. One cannot presumably use a Commissive act to imbue Commissive acts with deontic force without vicious circularity.

On the other hand, if one adopts a non-conventionalist approach, then one is likewise faced with the same question: why, then, do some speech acts have deontic force? The answer given by the non-conventionalist tends to be disappointing: that the deontic force of speech acts is primitive—*they just do*. For example, this is the answer Reinach gives to this question when he argues that promises produce obligations simply because there are “essential apriori laws” governing the act of promising. There is no way of further explaining such laws. All we can do is recognize their evident truth, just as one might with simple mathematical formulae. As Reinach puts it:

Strictly speaking we are not proposing any *theory* of promising. For we are only putting forth the simple thesis that promising as such produces claim and obligation. One can try, and we have in fact tried, to bring out the intelligibility of this thesis by analysis and clarification. To try to *explain* it would be just like trying to explain the proposition, $1 \times 1 = 1$. (Reinach, 1913/1983, p. 46)

Both traditional conventionalist and non-conventionalist approaches face their own challenges in explaining what grounds the deontic force of speech acts. My claim here is not that one or another approach cannot successfully resolve this problem. Rather, I merely wish to point out this problem as a way of motivating an alternative approach, which I believe better accounts for the power of speech acts to generate, modify, or eliminate deontic phenomena, but which does not ground a system of obligations, claims, permissions, authority, etc., fully on speech acts or commitment, and thus does not encounter the same problem. The details of this proposal will be filled in over the next few

chapters. For now, I will outline a framework which I call the Grounding-Anchoring Framework, which will later help us rethink how to ground deontic phenomena.

A First Look at Grounding and Anchoring

As a way of approaching the question of what explains the existence of deontic phenomena and institutions, I will draw a distinction, borrowed from Epstein (2015), between two metaphysical relations: *grounding* and *anchoring*. An extended example should help throw light on the nature of, and distinctions between, these two relations. First, consider some institutional facts like the following:

- I1** John was married to Yoko.
- I2** Madoff was arrested.
- I3** Rizzo struck out in Game 3 of the NLCS.
- I4** Cruz purchased a beef sandwich at Charlie the Butcher's.

On the one hand, such facts are perfectly mundane. But what makes (or made) it the case that any of these facts (**I1-I4**) obtain (or obtained)? To answer this question, we might plausibly point to other facts, for example:

- G1** John and Yoko uttered certain formulae during a legally recognized ritual in the presence of a recognized officiant.
- G2** A warrant for Madoff's arrest was issued, and a law enforcement official approached Madoff and told him he was under arrest, and then proceeded to take Madoff into police custody.
- G3** After accumulating two strikes, Rizzo either swung his bat and missed the ball, or else failed to swing when the ball passed through the strike zone.

- G4** Cruz requested a Beef-on-Weck sandwich, and then transferred money in the amount of \$6.99 plus tax to Charlie's, and for this reason received the sandwich in return.

But noting that these other facts (**G1-G4**) also obtained only goes so far in answering the question of why (**I1-I4**) should obtain. Why should uttering certain formulae in a certain context make it the case that two people are now married? And what does the efficacy of uttering of these formulae have to do with the ritual action being legally recognized? Similarly, why should the issuing of a piece of paper, or the uttering of certain words, mean that a person is now under arrest? Why is it apparently permissible for the law enforcement official, but not just anybody, to physically detain Madoff? Why should swinging or failing to swing an elongated piece of wood make it the case that Rizzo has struck out? Why should this fact make it the case that he is now required to leave the batter's box and return to the dugout? And why would the exchange of strange piece of paper or coins make it the case that Cruz may now consume the beef sandwich prepared by Charlie's?

Here, Searle (1969, 1995, 2010) would argue that (**G1-G4**) ground (**I1-I4**) because the following principles (in Searle parlance, "constitutive rules") are operative, roughly:

- P1** Uttering certain formulae counts as getting married in the context of a certain socio-legal ritual.
- P2** Certain pieces of paper issued by judges count as warrants in the context of a certain legal community.
- P3** Swinging and missing with two prior strikes counts as striking out in the context of a certain game.
- P4** Certain pieces of paper and coins count as money in the context of a certain community.

This, of course, invites the further question: why should (**P1-P4**) be operative? Here, Searle's answer is very roughly that a group has *collectively accepted* these rules. Thus, (**P1-P4**) are operative because:

- A1** Some group g collectively accepts (**P1**).
- A2** Some group g collectively accepts (**P2**).
- A3** Some group g collectively accepts (**P3**).
- A4** Some group g collectively accepts (**P4**).

This approach, then, starts with a set of institutional facts (**I1-I4**) and asks the question: in virtue of what do these facts obtain? The answer is (**G1-G4**), where (**G1-G4**) satisfy some principle, namely, (**P1-P4**), which are operative in turn because of (**A1-A4**).

Following Epstein (2015), I will describe the relationships between facts like these in terms of *grounding*, *anchoring*, and *frame principles*. Using the example of John and Yoko, these relationships can be depicted as follows. (**I1**) is *grounded* by (**G1**), which holds because of the general *frame principle* (**P1**), which in turn is *anchored* by (**A1**). (Note that in the following diagram the principle (**P1**) has been slightly formalized.)




Figure 3 An example of grounding, frame principles, and anchoring

The importance of the distinction between grounding and anchoring will be further defended below. For the present, however, I will simply make a few remarks on the notion of a frame principle. To

start, a frame principle can be thought of as roughly equivalent to a Searlean constitutive rule. But there is an important difference between them. As I will demonstrate with some examples below, frame principles do a far better job of specifying the situations in which a set of grounding facts ground a set of institutional facts. The reason is that the notion of a *frame* as found in Epstein (2015) is meant to formally capture the idea that principles like (P1-P4) hold relative to a set of possibilities, times, and contexts—basically, to a set of possible worlds.⁸ A frame, then, can be thought of as a structure containing the worlds in which the frame principle is operative.⁹

Generally, then, I will say that I-facts are *grounded* by G-facts, P-facts are the *frame principles* of G-facts and I-facts, and that P-facts are *anchored* by A-facts. Or, to give a generic diagram:




Figure 4 The Ground-Anchoring Framework

To this point, I have provided an overview of the Grounding-Anchoring Framework. But what exactly the grounding relation, the anchoring relation or a frame principle is needs further exposition.

⁸ My intent is not to commit myself to the existence of possible worlds, e.g., as sets of propositions, or counterpart worlds, or what have you. Rather, I am appealing to possible worlds merely as a useful heuristic for throwing light on the grounding and anchoring relations.

⁹ This is, admittedly, ontologically imprecise. When I describe a frame, I am attempting to describe the scenarios (actual or possible) in which a frame principle holds, i.e., in which certain social institutions are operative. How this gets cashed out ontologically really depends on one’s views of what makes statements about possible scenarios true. For David Lewis (1986), such frames would end up being portions of reality at some concrete, (relatively) actual world. For more traditional analytic philosophers (e.g., Plantinga 1974), possibilities might be cashed out as *abstracta* such as propositions. Or statements about possibilities might be true in virtue of actual dispositions for those possibilities to be realized, which would mean in turn that talk of frames would be an oblique way of describing dispositions to engage in social institutions. However, simply in introducing this theoretical apparatus, I am trying to avoid presenting a particular account of what makes statements about possibilities true, and thus of what frames are. But I will return to this question in Chapter 7.

Grounding and Anchoring among Other Relations

In recent analytic metaphysics, *grounding* is typically characterized as a relation of *metaphysical explanation* or *reason*. To give a few standard examples: the complex entity Socrates exists *in virtue of* simpler entities (his parts) existing; the proposition *that Socrates is white* is true *because* Socrates is white; the fact that an act is loved by the gods *is explained by* the fact of its piety.¹⁰ The first thing to note about such examples is that the relationship they describe is not physical causation. Although statements about causation are also explanatory, the kind of explanation involved in grounding is non-causal. The statement *that the barn burned down because Mrs. O’Leary’s cow kicked over the lantern* provides an explanation in terms what of what causal chain of events had the effect of burning down the barn. But neither the relationship between the cow’s action and the barn’s charred state, nor the relationship between the fire and the barn’s charred state, is like that of the relationship between the fact that the proposition *that Socrates is white* is true and the fact that Socrates is white. In the latter case, there is no causal chain to speak of.

Moreover, the grounding relation should not be confused with the relation of counterfactual dependence. To see why, consider the example of Plato’s Euthyphro problem (*Euthyphro* 10a): “Is an act pious because it is loved by the gods, or is it loved by the gods because it is pious?” When Socrates presents Euthyphro with the problem, he is in effect asking him about what grounds what. Does the fact that the gods love an action ground the fact that it is pious, or does its piety ground the fact that the gods love it? Socrates’ question presents a dilemma for any theologically grounded ethics. If acts are pious in virtue of being loved by the gods, then these acts have no intrinsic piety. Thus, the proponent of a theologically grounded ethics would be pressed to hold that any act whatsoever (murder, theft, etc.) could be pious so long as it is loved by the gods, which suggests in turn that the

¹⁰ The following discussion is indebted to (Hendricksen 2014).

value of acts is wholly arbitrary. But if acts are loved by the gods because they are pious, then apparently the value of a pious act is not theologically grounded. Thus, the Euthyphro problem presses the proponent of theological ethics into a dilemma of admitting arbitrariness into their ethics, or of rejecting theologically grounded ethics.

However, the Euthyphro problem is ostensibly *not* about counterfactual dependence. We can see why if we try to express the problem as a pair of counterfactual statements:

E1 If some other action had been pious, then God would have loved it.

E2 If God had loved something else, that thing would have been pious.

As Hendricksen (2014) argues, however, the divine command theorist might accept both **(E1)** and **(E2)** without having fallen victim to the Euthyphro problem. The reason is that, on standard possible world semantics, **(E1)** and **(E2)** would be further analyzed as follows:

E3 All of the nearest worlds where God does not love an action X are worlds where X is not pious.

E4 All of the nearest worlds where an action X is not pious are worlds where God does not love X.

But the divine command theorist can accept that, at all nearby worlds, actions that God does not love are also those actions which are not pious. And he can also accept that, at all such nearby worlds, actions that are not pious are not loved by God. So the divine command theorist does not face a dilemma between **(E3)** and **(E4)**, and yet he does not thereby avoid the Euthyphro problem. The reason why is that **(E3)** and **(E4)** say nothing about which of these facts is more fundamental—i.e., the fact that God loves X, or the fact that X is pious. The Euthyphro problem, then, is a question of

what grounds what, not a question of counterfactual dependence. Therefore, the grounding relation cannot simply be equivalent to the relation of counterfactual dependence.

Lastly, the relations of supervenience and grounding are also distinct. Usually, the relation of supervenience is characterized as follows: a set of properties A supervenes upon the set of properties B if and only if there can be no change in A -properties without some change in B -properties. One reason why the supervenience relation is typically thought to be distinct from grounding is that, whereas grounding is usually taken to be a relation between facts, supervenience is usually taken to be a relation between sets of properties. That being said, there is some lack of clarity in the literature. Some metaphysicians do not restrict the grounding relation to facts, or the supervenience relation to properties. Fortunately, however, there are other ways of demonstrating the difference between the two relations. Suppose that the relata of both relations are facts, and thus that the restriction to sets of properties were lifted from the above characterization of supervenience. Then, every fact supervenes on itself because, trivially, for any fact, any change in that fact entails a change in that fact, such that there could be no change in A -facts without a change in A -facts. But not every fact is grounded in itself. For example, the fact that the proposition *that Socrates is white* is true is not grounded in itself, but rather in the fact that Socrates is white. Put otherwise, the supervenience relation is necessarily reflexive, whereas the grounding relation is not. Therefore, we also need to distinguish grounding from supervenience.

What, then, is grounding? I think it is best understood as the relation of metaphysical priority holding between facts. It is a relation concerning which facts are more fundamental than other facts. Thus, the relatively more brute fact that John and Yoko uttered “I do,” etc., is metaphysically prior to the institutional fact that John is married to Yoko. The one set of facts are the metaphysical reason or explanation for the other. Derivatively, then, we can speak more generally of one entity being grounded in another. For example, if the fact that John is married to Yoko is grounded by the fact

that they uttered marriage vows, then we can say likewise that their marriage is grounded by the marriage vows.

Anchoring, or Why Grounding Is Not Enough

Above, I suggested that the fact that John was married to Yoko was grounded by the fact that John and Yoko uttered certain phrases in a marriage ceremony, which was conducted by a recognized officiant. Of course, this suggestion was an oversimplification. Many more facts would need to obtain as well: that John and Yoko uttered the phrases in earnest (as opposed to pretending), that there were witnesses besides John, Yoko, and the officiant, and so forth. There is additionally the fact that *these* are the conditions that needed to be fulfilled in order for the further fact that John and Yoko were married to obtain. And there are further questions about why this last fact—i.e., the fact that *these* are the requisite conditions—should obtain. From the purview of the Grounding-Anchoring Framework, the answer to such latter questions means providing an account of which frame principles are anchored, and what exactly anchors them. Thus, when Searle (1969, 1995, 2010) argues that institutional facts obtain because of collectively accepted constitutive rules, he is providing an account of which frame principles (“constitutive rules”) are operative and how they are anchored (“collective acceptance”).

Now it might be thought that the relationship between a frame principle and its anchor is likewise that of grounding. In other words, just as the fact that John is married to Yoko is grounded by the fact that John and Yoko uttered certain formulae, etc., so too the frame principle, “For all x, y , if x and y utter certain formulae, etc.,” is grounded by the fact that the group g collectively accepts the principle. Thus, there would not be two different types of relations—grounding and anchoring—but rather simply two levels of grounding: one between the institutional facts and its grounding conditions, and the other between the frame principle and a group. Epstein (2015) calls this view *conjunctivism*,

since it amounts to the idea that the conjunction of these sets of facts collectively ground the fact that, e.g., John and Yoko are married.

But Epstein rejects conjunctivism, arguing that this relation of anchoring is distinct from that of grounding. I follow Epstein in this, and I should explain why. According to conjunctivism, the fact that John and Yoko are married is grounded by the conjunction of the following facts (again, simplified slightly for the sake of example):

CJ1 If x and y utter wedding vows, then that grounds the fact that x and y are married.

CJ2 John and Yoko uttered wedding vows.

Moreover, (**CJ1**) would be grounded in turn by some further fact, for example:

CJ3 People collectively accept (**CJ1**).

Put otherwise, according to conjunctivism, the fact that there is a certain operative principle should be numbered among the facts that ground the institutional fact that John and Yoko are married.

Epstein rejects this account of what grounds the institutional fact that John and Yoko are married because it gets the modal conditions for its grounding facts wrong. To see why, consider the fact that many social kinds can be instantiated across various actual or possible situations, including those where the facts which anchor the relevant frame principle do not obtain. For example, consider the social kind *kosher*, and the facts (**K1**) and (**K2**), where (**K1**) is grounded by (**K2**):¹¹

K1 Peggy the pig is not kosher.

K2 Peggy the pig has cloven hooves.

¹¹ Cf. Epstein, 2015, pp.115-124.

Why does (K2) ground (K1)? The answer lies in the Jewish dietary laws defined in Leviticus 11:3-8:

K3 For all x , if x has cloven hooves, then x is not kosher.

There are obviously other possible frame principles which would likewise make it the case that (K2) grounds (K1) (e.g., “For all x , if x is cooked without first draining x ’s blood...”), but let’s take (K3) as our example. Now we can ask: what makes it the case that (K3) is an operative constitutive rule? For simplicity’s sake, let’s suppose, following Searle, that the following further fact establishes the frame principle (K3):

K4 Group j collectively accepts (K3).

But now, we can see that it makes sense to ask whether an animal counts as kosher in actual or possible situations in which (K4) does not obtain. For example, if we suppose that Peggy the pig lived prior to the institution of kashrut dietary laws, we can still ask whether Peggy would be kosher, despite the fact that, *ex hypothesi*, there is no group collectively accepting the rule (K3). Therefore, (K4) cannot be among the grounding conditions for (K1). After all, (K1) can be fully grounded despite the fact that (K4) did not obtain. Therefore, conjunctivism is false.

Here, then, is precisely why *frame principles* and the *anchoring* relation make a difference in how we approach issues of the ontology of social institutions. A general principle that establishes the conditions under which certain institutional facts are grounded cannot merely be about the institutional facts that do obtain. Rather, such general principles are intended to cover all possible (i.e., including non-actual) situations. Or as Epstein puts it:

[S]ocial kinds can be instantiated across the universe of different situations, contexts, or worlds. When we set up the conditions for some social fact to obtain, we set up the grounding conditions for that universe. We set up the conditions for it to obtain even in situations,

contexts, or worlds where we do not exist. This means we have to generalize from a single situation, context, or world to a universe of possible worlds. A *frame* is a structure containing this universe, that is, a set of possible worlds in which the grounding conditions for social facts are fixed in a particular way. (Epstein, 2015, p. 78)

Thus, it is crucial to maintain a distinction between (1) which more fundamental facts are grounding which institutional facts relative a frame, which is a question of *grounding*, and (2) the frame itself, which is defined by a frame principle, which in turn is *anchored* by some further facts (e.g., collective acceptance). Whereas the grounding facts tell us about which institutional facts do obtain, a frame provides a space of possibilities for institutional facts to obtain.

One further remark about the Grounding-Anchoring Framework. In describing a frame principle—e.g., “For all x, y , if x and y utter certain formulae..., then x is married to y ”—it will be essential to clearly state the principle’s *domain of quantification*. A quantifier (e.g., “For all”) is best understood as a function relating a set of to a set of possible individuals (its domain) to a set of individuals satisfying the function (its range). To use a simple example, if the domain of quantification for the above frame principle is the set of possible human beings, then the only individuals that could be bound to the x and y variables are human beings. A parrot declaring “I do” would, thus, not fall within the principle’s domain of quantification. This will be an important detail to keep in mind later when I provide an account of which groups are affected by an anchored deontic frame principle.

The Grounding-Anchoring Framework has the virtue of clarifying two different sorts of metaphysical foundation that a deontic phenomenon might have, and thus of clarifying the difference between what grounds a social entity, what frame principles are responsible for establishing the relevant grounding conditions, the contexts in which those frame principles are operative, and what sets up those grounding conditions in the first place. These distinctions, although implicit to varying degrees in previous work in social and deontic ontology, are otherwise easy to overlook. As I hope to show in the following two chapters, the Grounding-Anchoring Framework is also a powerful tool for

analyzing what grounds deontic phenomena and anchors their associated frame principles—one which does a better job of explaining what role speech acts play in the emergence of deontic phenomena, but one which also does not attempt to found all social obligations, permissions, or authority on speech acts themselves.

Groups and Grounding

In surveying some typical accounts of social ontology, we saw a common pattern emerge: social institutions—together with their roles, powers, permissions, obligations, etc.—exist because individuals come together to collectively accept that social institution. While I believe that this picture of what grounds deontic phenomena is basically the correct one, this picture also leaves us with many questions about its details. Given that social institutions are established through collective acceptance, we can still ask: Are all social institutions accepted in the same way? By the same type of group? How are different social institutions related to one another? Are some social institutions based on other ones, and if so, then are those derivative ones likewise based on collective acceptance? Given that a social institution is collectively accepted, to whom does it apply? To only the group which has accepted it, or to those outside that group? What if there are two incompatible social institutions in a single social group?

The aim of this chapter is to shed some light on these questions by advancing a pluralist approach to what grounds deontic phenomena and what anchors their associated frame principles. According to this approach, there is no one way of grounding a deontic phenomenon. Rather, any particular case of grounding might depend on one or more operative frame principles. For example, an obligation might be grounded in an explicit agreement (e.g., an act of promising), but it might also be grounded in the decree of a preexisting authority. Moreover, in any given social context, there

might be diverse (even incompatible) anchored frame principles, depending on the groups or subgroups an individual belongs to. This is possible, in part, because the group anchoring a frame principle is not necessarily the same as the group to which a given frame principle is meant to apply.

Different Ways to Ground a Deontic Phenomenon

In light of the Grounding-Anchoring Framework, we can see that there is no one way in which deontic phenomena are grounded. Rather, the same institutional fact can be grounded by many different types of grounding facts. Consider four cases of agents having obligations, each of which is grounded in distinct ways. These four cases are not meant to be exhaustive, but rather to illustrate a few typical ways in which deontic facts are grounded: agreement, authority, rights violations, and normative expectations. In principle, there is very little restricting how a deontic fact might be grounded.

Case 1: Agreement-Based Grounding

Penny says to Andy, “I promise to pick you up from the airport,” and thus gains an obligation to do so. Given the Grounding-Anchoring Framework, we can describe the following features of this case. First, there is the institutional fact that Penny is obligated to pick up Andy from the airport. Second, this institutional fact is grounded by the fact that Penny said, “I promise to pick you up from the airport.” These grounding conditions are set up by the frame principle: For all x, y , if x says to y , “I promise to φ ,” then x is obliged to y to φ . (For simplicity’s sake, I will ignore any other grounding conditions that might need to be fulfilled.) Thus, we can depict this deontic scenario as follows:




Figure 5 A case of promising

Note that the point of disagreement between, e.g., Hume, Searle, and Reinach is not about institutional facts, grounding facts, or frame principles, but about what anchors that frame principle. All agree that promising in some sense grounds an obligation; they disagree about *why* a promise does. (Much more on this in the next section.)

This case of an obligation being grounded in promising can be understood more generally as a case of a deontic phenomenon being generated directly through agreement. There are obviously myriad cases of promise-like activity in which one agent commits to a course of action and thereby becomes obliged to do it, and in which others gain a claim against that agent. For example: volunteering to help a friend move, signing up to bring a dish to a potluck, and getting married. Note that, in such cases, the strength of the obligation or claim can vary widely depending on the nature of the agreement or agreed upon task. In many (if not all), we believe that obligations can be overridden by countervailing factors, perhaps competing obligations or permissions. To use a simple example, Penny’s obligation to pick up Andy from the airport might be overridden by a family medical emergency. In some such cases, the negative effects of not fulfilling the obligation can ground a compensatory obligation, i.e., an obligation to make up for failing to keep some initial obligation.

There are also cases of permissions, powers, and immunities, which are generated through agreement. For example: Larry's agreement to lend Belinda his power drill generates a permission for her to use it; signing a contract of employment generates a power in one's employers and supervisors over the employee; and a criminal can agree to testify in a criminal court in exchange for his immunity from prosecution.¹

Case 2: Authority-Based Grounding

Candy assigns Terry a task t , and thus Terry gains an obligation to perform t . Given the Grounding-Anchoring Framework, we can describe the following features of this case. First, there is the institutional fact that Terry is obligated to perform t . Second, this institutional fact is grounded by the fact that Candy communicated to Terry the directive that Terry t , e.g., by telling him in person, or sending him an email. These grounding conditions are set up by a frame principle, for example: For all x, y , if x is y 's boss and x commands y to ψ , then y is obligated to ψ . (Again, take this as a simplified example of a frame principle.) Thus, we can depict this deontic scenario as follows:

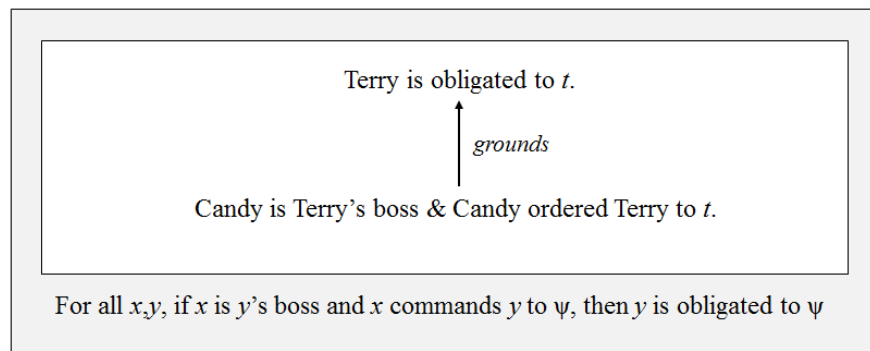


Figure 6 An obligation grounded by preexisting authority

¹ Gilbert's (2014, pp. 26-28) account of obligations arising from joint commitment would fall under this header, though her account attempts to generalize from cases of agreement to cases of "personal readiness" to perform an action.

In some respects, this case of obligation is significantly different from *Case 1*. Rather than commit explicitly to doing t , Terry finds himself in a situation in which another individual, Candy, has the power to generate obligations for him. In other words, within the context of their employment roles. Of course, this raises the question of why Candy has such power, i.e., in virtue of what she is Terry's boss. Thus, we can extend this second case as follows, noting that Candy's role as boss is grounded in further facts, which themselves as subject to frame principles. In other words, we can treat the extension of the second case as that of a frame defined in terms of multiple frame principles:

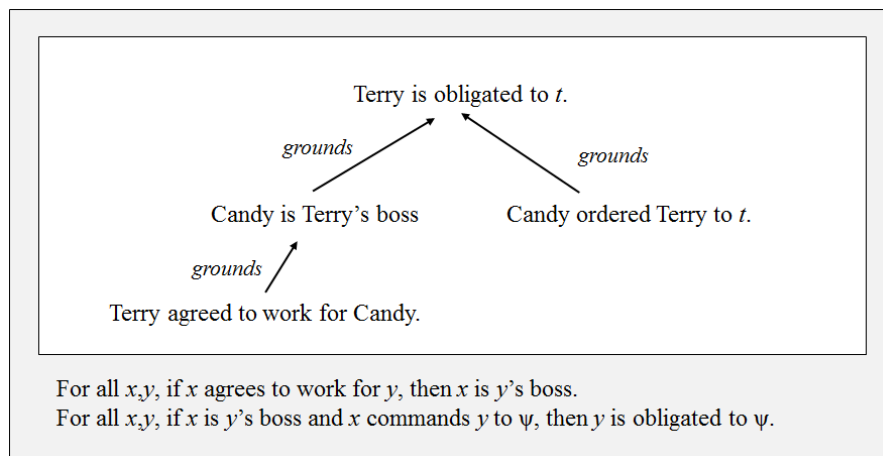


Figure 7 Authority grounded in turn by agreement

Moreover, the outer frame principle here might be nested within a further frame, e.g., one with a frame principle dictating the grounding conditions for agreeing to work. Moreover, note that in both cases I am omitting any discussion of what *anchors* either frame principle, since this will be the topic of the following major section.

There are also cases of permissions and immunities, and even of powers, which are generated because of an agent's existing powers to modify first- or second-order deontic phenomena, in

Hohfeld's sense. For example: the power of New York State to issue driver's licenses, or the power of Candy to put Terry in charge while she is away, and thus to impose tasks on other employees.

Interestingly, it is possible to construe agreement-based deontic phenomena as a species of authority-based deontic phenomena. For example, Penny's promise to pick up Andy might be understood as the exercise of her power over herself—i.e., her autonomy. Thus, just as Candy has the power to impose requirements on Terry, so too Peggy has power to impose requirements on herself. In so doing, she would also generate a claim for Andy, and thus disable herself from backing out of the agreement. But even so, describing promising in this way would be one way to unify what might look like two disparate cases of grounding obligations.

Case 3: Violation-Based Grounding

Henry hits Carrie's car, and thereby gain an obligation to compensate her for the damages. What should be striking about this case is that Henry's obligation is not based on agreement or pre-existing power relations. Rather, it is the negative consequences of his action that generate an obligation for himself. Given the Grounding-Anchoring Framework, we can describe this case as follows. First, there is the institutional fact that Henry is obligated to pay for damages to Carrie's car. Second, we know what grounds this institutional fact, namely, the fact that Henry damaged Carrie's car. It might not be intuitively obvious what the relevant frame principle is here, but we can suppose it is something simple like the following: For all x, y , if x damages y 's property, then x is obligated to make recompense to y .

Thus:




Figure 8 Violation-based grounding

One important lesson from *Case 3* is that, despite the emphasis which many theories of social ontology (rightly) place upon the role of speech acts in generating, modifying, or eliminating deontic elements, Henry's obligation is in no way grounded in a speech act. Rather, it is grounded in an act which violates one of Carrie's rights, namely, a claim against anyone damaging her property. In other words, the obligation arises out of the logical connections between preexisting rights and obligations, not out of a speech act. The frame principle in this case is, in fact, one rough example of how a claim-right and compensatory obligation might be interrelated.

What is central to *Case 3* is the violation of Carrie's claim against others not to damage her property. More generally, then, we can think of claim-violations as a broad category of grounding conditions for a deontic phenomenon, especially obligation and claim.

Case 4: Expectation-Based Grounding

In Sammy's society, there is the social norm that everyone who attends a dinner party shall send a thank-you note. So when Sammy attends Danny's dinner party, she is thus obligated to send a thank-you note. Given the Grounding-Anchoring Framework, we can describe the features of this scenario as follows. First, there is the institutional fact that Sammy is obligated to send the note. Second, this

institutional fact is grounded by the fact that Sammy attended Danny’s dinner party. These grounding conditions are set up by a frame principle, roughly: For all x,y , if x attends y ’s dinner party, then x is obligated to send a thank-you note to y . Thus, we can depict this scenario as follows:

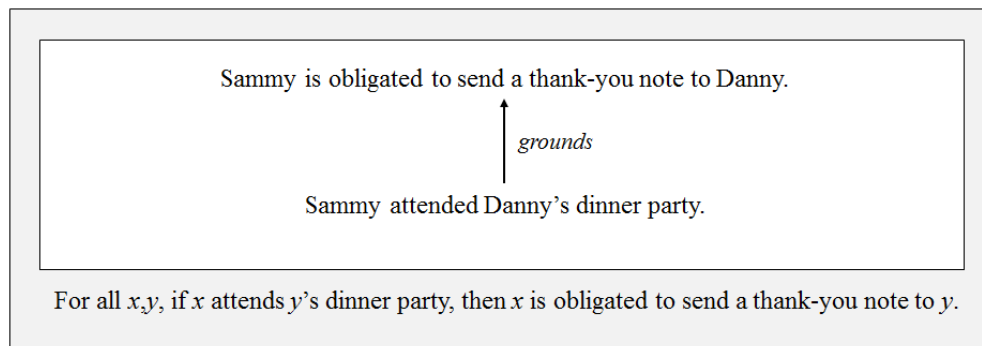


Figure 9 *Expectation-based grounding*

Again, we can note about *Case 4* that the obligation in question is not grounded in any speech act. Rather, the obligation is grounded in the fact that she attended the party, where that grounding fact satisfies an operative frame principle. Moreover, in this case, we have rough idea of what anchors that principle: Sammy’s society has what we might describe as a *normative expectation* that people send thank-you notes. By this, I mean (again, *very* roughly) that Sammy would be judged to have done something wrong in not sending a note, and that she might be sanctioned in some manner (even if only thought badly of or gossiped about). But this very rough idea will need to be refined considerably in the following section, which deals directly with deontic anchoring.

Such norm-like frame principles, which set up grounding conditions for obligations, are not uncommon. For example: the norm to not interrupt others in conversation, to not stand too close to others in conversation, or not to sit directly next to somebody if there is other seating available. In each case, the obligation is intuitively based on expectations about agents’ behavior. But it should also not be difficult to see how permissions, powers, and immunities might likewise be grounded in such

frame principles. For example: the freedom to eat what one wishes (within certain broad boundaries), the power to purchase what one wishes (again, within certain broad boundaries), and immunities against discrimination and harassment.

The upshot of the foregoing cases is twofold. The first is that we should be pluralists about what grounds deontic phenomena. Although speech acts are relevant for grounding many deontic phenomena (especially, agreement-based or authority-based obligations), there are also cases of deontic phenomena grounded not in speech acts, but, e.g., in expectations about agents' behavior. This claim should not be counterintuitive, since we have no good antecedent reason to believe that all deontic phenomena will be grounded in the same way. The second is that the Grounding-Anchoring Framework gives us a perspicuous way of accounting for the multiple ways that deontic facts can be grounded. Different frame principles set up different grounding conditions, and establish a space of possible worlds in which that frame principle holds. Thus differences in frames can thus be thought of as different perspectives on the same brute facts, which, given a particular frame principle, ground different institutional facts relative to that frame. These multiple ways of grounding thus reflect the fact that institutional reality is not monolithic, but rather a set of social facts tied specifically to a set of diverse principles grounding the ways in which that institutional reality can arise from a set of (ultimately) non-institutional facts.

This approach to what grounds deontic phenomena can be taken a step further, to the level of what anchors the frame principles which establish the grounding conditions for deontic phenomena. Any given society might be composed of diverse groups and subgroups, each of which might have established very different frame principles. So, the next step is to take a closer look at the makeup of groups and the different kinds of groups relevant to the anchoring of diverse frame principles.

Groups and Collective Acceptance

In a previous chapter, I remarked that it is commonplace in discussions of social ontology to hold that social institutions exist in virtue of a group coming together and accepting or committing to that institution. In Searle's view specifically, he describes this as the collective acceptance of constitutive rules—or what I have called frame principles. This view raises three broad questions about the collective acceptance of social institutions. First, what precisely does it mean to *accept* a social institution? Second, what does it take for an aggregate of individuals to form a *group*, capable of collectively accepting an institution? Third, to whom does an institution apply? For example, is the group which accepts the institution identical with the group to which the institution applies, or are there individuals or groups that do not accept the institution, but to which that institution nonetheless applies?

Intuitively, the collective acceptance of a constitutive rule will hold relative to a group context. Thus, to return to Searle's most typical example, certain slips of paper would count as money in the context of the United States economy because at least Americans—i.e., those who are actively engaged in using these slips of paper as a medium of exchange—accept these slips of paper as a medium of exchange. However intuitively appealing this picture of social grounding might be, it raises several problems. To start, it raises the problem of how much individual acceptance is required for collective acceptance. Intuitively, 100% acceptance would not be required. (A single dissenter to the institution of fiat money does not thereby overturn the institution.) At the same time, it would probably be too crude to mark 51% (or 50.1%, or 50.01%...) acceptance as collective.

Second, it raises the question of what is required for being a group. For example, according to Gilbert (1989, 2000, 2014), individuals constitute a group only if they have jointly committed to performing a course of action together. Alternatively, according to Tuomela (1995, 2005, 2007), individuals constitute a group only if they think of themselves as composing a “plural subject,” i.e.,

that they think of themselves as a literal corporate agent with its own distinct intentional states over and above those of its members. But we might wonder whether there is a looser notion of group required for grounding social institutions. For example, individuals engaged with common political institutions in a common geographic region might constitute a group, whether or not they have jointly committed to one another in pursuit of a common goal, or think of themselves as constituting a plural subject.

Lastly, it raises the question of the relationship between the group which accepts some social institution and the group that is subject to that institution. Previously I mentioned that a single dissenter to an institution does not thereby overthrow institution. But if that individual does not personally accept the institution, then why is he beholden to it? This concern can be brought out in another way. Suppose that the United States declares that it has annexed Canada, and that the 300+ million Americans accept this, whereas the 36 million Canadians do not. Despite the fact that nearly 90% of individuals in this collection accept that Canada is now part of the United States, one might object that this acceptance would not be sufficient for making Canada part of the United States. And in fact, there are plenty of real cases of territorial disputes, where competing claims to sovereignty render the institutional facts unclear. But note that this same issue can arise within the context of a single society, in which one might find, not a single homogeneous group, but overlapping groups with diverse beliefs, attitudes, and commitments. In this case too, we might wonder whether 90% acceptance within a society should affect the social or deontic status of the 10%. A society as a whole can certainly be considered a type of group, but within that group there will be subgroups (e.g., racial, ethnic, or religious minorities) with distinct and incompatible aims, values, beliefs, etc.

Collective Intentionality

Discussions of a group's acceptance of a social institution have also focused heavily on *collective intentionality*, or the capacity of individuals to adopt intentional states (belief, desire, intention, etc.) which have not just themselves as subjects ("I-intentionality"), but rather the group of which one is a member ("we-intentionality"). Now there are at least two main questions here regarding collective intentionality. First, there is the question of whether collective intentional states imply that there is a distinct collective agent or consciousness over and above individual group members (ontological holism) or not (ontological individualism).² Not many philosophers have endorsed ontological holism, although Popper (1945) criticizes Hegel on precisely these grounds. Most have, instead, opted for ontological individualism. For example, although Searle (1995) describes collective intentionality as the capacity for animals to "share intentional states" (p.23), he subsequently clarifies that he is not proposing that there is a literal "collective consciousness" or "plural subject," over and above individual subjects, which is the bearer of the collective intentional state. Needless to say, most contemporary philosophers and sociologists seem to adhere to an ontological individualist position.

Second, there is some debate over whether collective intentionality is reducible to individual intentionality (explanatory individualism) or not (explanatory holism). Gilbert (1989, 1996) and Tuomela (1995, 2002, 2005, 2007) staunchly deny that collective intentionality is so reducible. By contrast, Searle thinks that some cases of collective intentionality are reducible, whereas others are not. Specifically, he takes cases of coordinated, cooperative behavior to be irreducible (Searle, 1995, pp.23-36; 2010, pp.47-8). For example, a violinist's playing in an orchestra is a case of cooperative collective intentionality, since that violinist's intention is not merely *that I play*, or even *that I play and I believe that you play*, but rather *that we play (together, cooperatively)*. In other cases, however, Searle thinks

² For a thorough discussion of the history of this debate, see (Epstein, 2015), especially ch.1-4. The terminology in this section—ontological vs. explanatory, individualism vs. holism—is borrowed from Epstein.

that collective intentionality is so explanatorily reducible—specifically, in cases of the collective acceptance of constitutive rules. For example, he claims that A and B 's collective acceptance of a constitutive rule for checkmate is reducible to A 's recognition of the rule plus A 's belief that B recognizes the rule, and to B 's recognition of the rule plus B 's belief that A recognizes the rule. Note that, in Searle's view, collective acceptance of a constitutive rule does not imply approval of the rule (Searle, 2010, p. 8). For example, people might collectively recognize that a type of act is against the rules of a game, but not agree that it should be against the rules.

Note too that, in these recent discussions of collective intentionality, many authors have suggested that collective intentionality plays a role not only in the acceptance of a social institution, but in the constitution of a group itself. This point might be easily overlooked. The claim is not merely that a collection of individuals must share an intentional state in order for a social institution to get off the ground; rather, their sharing of that intentional state is the glue which holds the group together. But is that true of all groups? Is that even true of all groups which share intentional states? More generally, it highlights the question of what is required for constituting a group, as distinct from the question of the role that group's collective intentionality might play in the acceptance of a social institution.

What Groups Might (Not) Be Like

I will call a *group variable* any attribute which a group may have, which will be relevant to the question of what is required to constitute a group which collectively accepts an institution. Specifically, I will be concerned with identifying which group variables might be relevant to the acceptance or application of a system of rights, obligations, permissions, powers, and so forth. To begin, then, consider the following examples of collections of individuals, each of which are groups in their own way:

- people born in the month of January

- the commuters entering the 135 St. Metro Station on January 15th
- American citizens
- American residents (i.e., people who in fact reside in the U.S.)
- a family
- the Zulu people
- the New England Patriots roster
- the nine justices of the Supreme Court of the United States
- the Board of the Microsoft Corporation

One thing that should be striking about this list is its heterogeneity; there is not just one kind of collection of individuals which can be called a *group*. However, in sketching out some of the differences between different examples of groups, I believe we can gain a better understanding of the variables that come with different kinds of group and criteria for group membership. One simple example should suffice: for some of these examples of groups, membership requires an individual to know (at some time) that he is in the group, and to opt into the group, whereas in others, group membership requires no knowledge or choice to opt in. One cannot accidentally and unwillingly be a member of the New England Patriots, but any child born in America is an American citizen whether they know it or not, and even though they did not opt into the group.

What variables, then, might be identified? I propose the following three:

1. **Collective Goals:** Whether a group is such that *as a group* it has a goal, and if so, then whether or to what extent group members individually adopt the collective goal.
2. **Institutionality:** Whether a group's existence depends upon the existence of a social institution and its attendant institutional roles.
3. **Membership Origins:** How an individual becomes a member of the group (e.g., whether a member must intentionally opt into the group).

Some authors have focused specifically on groups in which members must share not just any collective intentional states, but specifically a collective goal. Thus, Gilbert's (1989) favorite example of a group

is two individuals who agree to take a walk together, who thus form a group specifically insofar as they jointly commit to achieving some goal together. However, we should note the distinction between a group having a goal *qua* group, and individual members of that group adopting or internalizing that goal. For example, there is no question that the employees of the Microsoft Corporation have a goal *qua* group, and moreover that the individual work by its employees are intended to promote the achievement of that goal. At the same time, it may be that many of Microsoft's employees do not adopt this collective goal as an individual goal irrespective of their group membership. Contrast this with the case of the New England Patriots, whose members are far more likely to adopt on an individual level the goal of the team as a whole (i.e., to win the game).

Some authors have also focused on the ways in which some groups are bound up with the existence of social institutions, whereas others are not. For example, Gilbert's two walking companions form a group, but the formation of this group does not occur within the context of some social institution. (In Searle's terminology, this is a *social fact*, but not an *institutional fact*.) By contrast, the existence of the New England Patriots as a group hinges upon the existence of the institution of football, and the existence of the Supreme Court hinges upon the existence of a host of American political institutions. For this reason, many groups of the latter kind are also defined in large part by the roles group members assume, and the relationships between those roles. With social institutions come the invention of institutional roles (in Searle: *status functions*), such as quarterback and Chief Justice, which help define both the group itself and the structural relationships between members. For example, a football team is defined in part by the various roles or offices occupied by the group members (e.g., quarterback, wide receiver), the activities characteristic of those roles (e.g., passing the ball), and the relationships between those roles (e.g., the power to call an audible). I note in passing that these structural relations are deontic in that they invariably involve what a group member *may* or *must* do within the context of that group.

Lastly, some authors have differentiated groups based on membership origins. One clear case is when group members must consciously and willingly opt into the group, whether formally (e.g., meeting official citizenship requirements) or informally (e.g., self-identification). By contrast, other cases of group membership include some members without requiring a member to consciously or willingly become a member of the group. For example, a child born in the United States becomes a member of the group of American citizens, although this group membership is unconscious (at least for a while) and non-voluntary. Another interesting set of cases concerns whether recognition from other group members would be required.

There is no one way of constituting a group; rather, what constitutes a group will depend upon what kind of group is being so constituted. Thus, addressing the question of the way in which a group is responsible for establishing the social institution, as well as of what group the institution applies to, will require a foray into the question of what those groups are like. Does the group have collective intentionality? Is this group itself reliant upon a social institution? Can one be a member of this group involuntarily?

Anchor Groups and Frame-Domain Groups

Up to this point, I have been focusing on some of the attributes that will be relevant for a group to accept a frame principle. But we can also ask what is required to constitute a group to which a frame principle applies. More formally, I propose a distinction between two different collections of individuals: the group accepting the frame principle (which I will call the *anchor group*) and the group which is subject to the frame principle, or to whom the principle applies (which I will call the *frame-domain group*). In some cases, these groups might be identical: all and only members of the anchor group are members of the frame-domain group. But it is also entirely possible that the group anchoring the frame principle is considerably narrower or broader than the group to whom the principle applies.

Earlier I mentioned that collective acceptance of a frame principle—whatever it might involve—is intuitively relativized to a group context, but that the details of this suggestion are not entirely clear. I believe that we can go a long way toward clarifying this point by looking to the distinction between an *anchor group* and a *frame-domain group*. Consider, for example, the anchoring of the following frame principle:

RF For all x , if x has a combination of ten, jack, queen, king, and ace, all of the same suit, then x has a Royal Flush.

We can depict the anchoring of the frame principle **RF**, together with an example of a grounding fact which satisfies the principle, as follows:

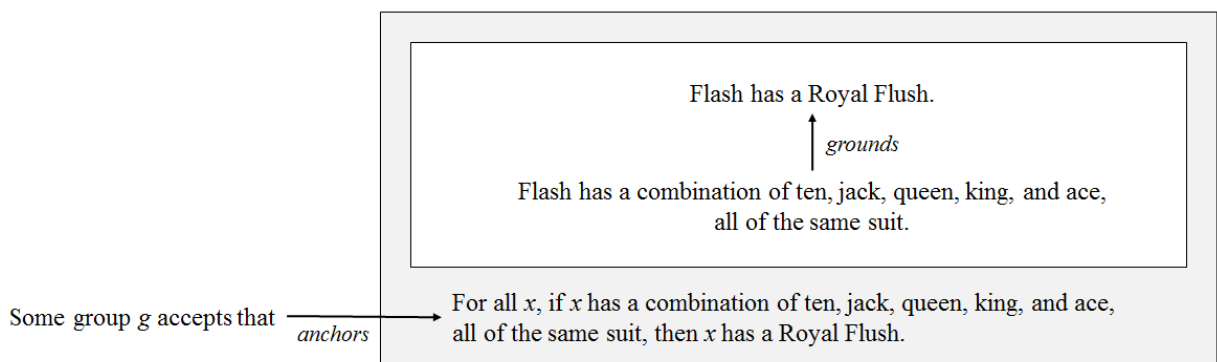


Figure 10 Grounding in a game of poker

In this example, we can identify and differentiate two relevant kinds of group. First, there is the group g which anchors the frame principle **RF**. This group might be, for example, an organization responsible for codifying rules for poker tournaments, or simply a group of friends who got together and agreed to play a card game with that rule. But what is important for our purposes is that this group g is not identical with the group of individuals to whom the rule applies. To see why this is the case, consider the example of the player Flash, who has never played poker before, and who has only a

vague conception of some of the rules (e.g., the rule of folding). Due to his beginner's luck, Flash ends up with a hand comprised of a ten, jack, queen, king, and ace, all of the same suit. Thus, Flash has a Royal Flush. However, Flash does not know that he has a Royal Flush, because he does not know the rule expressed by **RF**. Flash is not a member of the group g responsible for anchoring the rule, but the rule does apply to him, simply because he is playing poker. Thus, we can differentiate the following two groups: (1) the group of people who know about, and who collectively recognize, the standard rules of poker, and (2) the group of poker players. In a case like poker, it is highly likely that many (or even most) members of the latter group are also members of the former group, but (as the case of Flash demonstrates) there can be exceptions.

More generally, we can differentiate between an *anchor group*, which is the group which collectively accepts a (set of) frame principle(s), and a *frame-domain group*, which is the group to whom the accepted frame principle is meant to apply. In the case of **RF** above, we can suppose that the frame principle is anchored by the general population of poker players, who all recognize that a hand of that type counts as a Royal Flush. But we can also ask whom **RF** applies to. This frame principle begins “for all x ,” but what does that “for all” mean? What are the possible values of “ x ” in **RF**? In other words, we can ask about the *domain* of the quantifier found in **RF**. Obviously the scope of the quantifier is not meant to include just any entity, but rather only people engaged in playing a game of poker (and not bridge, blackjack, or rummy). What the case of Flash is meant to illustrate, then, is a case in which an individual belongs to the group of poker players (and thus falls within the scope of the quantifier in **RF**), but not to the group of those individuals who collectively accept **RF**. That being said, it is highly plausible in the case of poker that the rule would not be anchored unless the vast majority of poker players did collectively accept **RF**.³

³ But there is also the case of the card game Mao, in which newcomers are intentionally kept in the dark about the rules. New players must rather learn the rules through trial and error (and penalization for violating rules of which they were

It is possible, however, for an anchor group to be considerably narrower than a corresponding frame-domain group to which an anchored principle applies. The reason is that, in principle, there is little restricting what the domain of an anchored frame principle is. Consider, for example, the social kind *Gentile*, which is governed by the following traditional halakhic frame principle (which, note, is an example of what looks like a non-deontic frame principle, but which has deontic implications):

GN For all x , if x does not have a matrilineal Jewish ancestor, then x is a Gentile.

For simplicity's sake, let us assume that all 14 million Jewish people worldwide collectively accept **GN**, but that the vast majority of the remaining roughly 7 billion people worldwide are not cognizant of **GN**. Then, the anchor group would be the Jewish people as a whole. But what would the frame-domain group be? To whom does the anchor group intend the principle to apply? The intended domain is clearly *all* human beings.⁴ Thus, in this case, not only is the anchor group distinct from the frame-domain group, but it is also considerably smaller in size than the frame-domain group.

Note, too, that in this case it would not matter if the vast majority of Gentiles were aware of, or whether they accepted, **GN**. In the first place, the principle **GN** would be anchored even if all Gentiles were ignorant of it. More importantly, the acceptance of **GN** by the Jewish community and by Gentiles would not be viewed as holding equal weight within Jewish law. More directly, whether a non-Jew accepts a rule for who counts as Jewish would be thought of as irrelevant to the question of whether the rule is anchored. Even then, Jewish law might in reality further restrict membership in the anchor group to Jewish scholars or teachers. In other words, in the case of a principle like **GN**, additional conditions would need to be satisfied in order to count as a member of the anchor group.

unaware). This is an even clearer case of individuals falling within the frame-domain group without anchoring the corresponding frame principle. The rules are, rather, anchored by the total community of previous Mao players.

⁴ Recall that frame principles open up spaces of possible worlds. Thus, the frame principle **GN** is meant to apply to all human beings past, present, and future, and would even apply to counterfactual scenarios, e.g., if John had been born of a Jewish mother, he would have been Jewish.

This would, in turn, depend on further frame principles which establish the grounding conditions for being an official scholar or teacher. Unsurprisingly, **GN** would turn out to be a frame principle that is nested in other frame principles.

Nested and Competing Frame Principles

This distinction between anchor groups and frame-domain groups can also shed some light on how to account for social institutions which are anchored among diverse groups. Although I will not explore this question in depth, I will suggest that a diverse social population could be represented within the Grounding-Anchoring Framework in terms of a set of distinct frames, each of which is anchored by its own respective subgroup which accepts the corresponding frame principle, but whose frame-domain group might be considerably broader (e.g., the society as a whole). Moreover, insofar as a diverse social group is composed of distinct subgroups, we can also represent those subgroups as parts of a larger group.

The coordination between diverse subgroups within a population, then, could be understood in terms of deciding which frame principles are anchored by the larger group as a whole. Consider, for example, a society composed of subgroups with opposed ideas of the morality of drug use, and which advocate for two different frame principles, roughly:

M1 For all x , if x contains marijuana, then it is permissible to ingest x .

M2 For all x , if x contains marijuana, then ingesting x is prohibited.

In the first place, (**M1**) and (**M2**) can be thought of as competing frame principles which are anchored by particular subgroups (pro- and anti-marijuana camps, respectively). The larger group might come to accept one over the other, and this social acceptance might even receive legal ratification. In this case, it would become possible to say that the group as a whole collectively accepts just one of these

principles (e.g., **M1**). However, the other frame principle (e.g., **M2**) might continue to be anchored by the subgroup within that larger society (e.g., the anti-marijuana segment of the population, a subculture, a family). The main point here is that the Grounding-Anchoring Framework leaves room for non-monolithic social groups, which as a whole anchor a frame principle, but which might break down into distinct subgroups which anchor frame principles at odds with the frame principle anchored by the broader group.

Note, too, that this leaves open the possibility that the frame principle anchored by the group as a whole is a compromise between competing subgroup frame principles. For example, suppose two subgroups disagree over the following two frame principles, representing relatively extreme positions on the permissibility of gun ownership:

GO1 For all x , if x is a firearm, then it should never be permissible to own x .

GO2 For all x , if x is a firearm, then it should always be permissible to own x .

The resolution of this disagreement might be for a society to anchor (**GO1**) or (**GO2**), but it might be to anchor some third frame principle which is a compromise between them, i.e., that gun ownership is permissible under certain conditions C (e.g., after a certain age, with a background check and a waiting period):

GO3 For all x , if x is a firearm, then owning x is permissible only if conditions C are met.

This likewise highlights the way in which, from the perspective of the Grounding-Anchoring Framework, we can at least sketch an approach to representing the anchoring of frame principles within diverse populations.

More importantly, however, what the Grounding-Anchoring Framework gives us is an account of what it means for a frame principle to hold *relative to a context*. Searle's examples relied on a

fairly informal notion of the context in which an entity (X term) assumes a status function (Y term). But more formally, we can say that a frame principle is anchored by a particular group and that it holds for the intended domain of the principle's universal quantifier, with the caveat that there might be some competing anchored frame principle, which might trump that other frame principle. For example, if (GO3) is anchored in a society and legally ratified, it might still be true that there is a subgroup within that society which anchors the principle (GO1). The views of the group anchoring (GO1) would find themselves at odds with present legal permissions regarding gun ownership; however, they might still lobby for much tighter restrictions on gun ownership (e.g., the repeal of the Second Amendment), and personally sanction those who do own guns (in conversation, through campaigns, etc.). In sum, then, (GO1) would be anchored by the subgroup, but the intended domain of the frame principle is the society as a whole, but (GO3) would be anchored by the society as a whole (and backed by legal ratification), and would thus tend to trump (GO1) (e.g., via the protection of gun ownership by courts and law enforcement).

Systems of Frame Principles

To this point, I have presented simplified cases of what would be required for anchoring the frame principle, roughly, that anchoring is a matter of direct acceptance of a frame principle by an anchor group. This simplification, however useful for introducing the main idea, overlooks the complexity involved in accepting an institution, e.g., fiat money. As Searle rightly notes, collective acceptance is, in reality, acceptance not of individual constitutive rules, but of a "system of constitutive rules" (Searle, 1995, p. 140). For example, the acceptance of the institution of fiat money in the U.S. is bound up with the acceptance of the authority of the U.S. government to issue legal tender, which itself is bound up with the acceptance of a complex system of laws, courts, and offices, together with their attendant powers.

Some of the rules in such a system will be those responsible for the creation of new rules or the modification or elimination of existing ones. This is a point traditionally associated with Hart (1961), whose theory of law draws a distinction between a *primary rule*, which he characterizes as a rule that directly regulates behavior by permitting or requiring an action, and a *secondary rule*, which he characterizes as rules for generating, modifying, eliminating, and enforcing primary rules. Thus, an example of a primary rule might be a speed limit, whereas an example of a secondary rule might be a rule determining how speed limits are set. Note that this distinction bears some resemblance to the Hohfeldian distinction introduced in Chapter 1 between first-order deontic phenomena (i.e., regulations on action) and second-order deontic phenomena (i.e., powers toward, or immunity from, changes to the regulations upon one's actions). Thus, although for illustrative purposes, I will often refer to relatively simplistic cases of anchoring, the reality of anchoring would more accurately be depicted in terms of an anchor group's acceptance of a frame principle (or set of frame principles), in which other frame principles might be nested. Specifically, a group might collectively accept a secondary rule, which establishes the conditions for when certain facts can anchor a primary rule.

To borrow an example from Epstein (2015, p. 96), consider the primary rule *L* which claims that whoever kills with deliberately premeditated malice aforethought has committed a first-degree murder. What anchors *L*? It is not accepted directly through collective acceptance. Rather, *L* is anchored by the fact that there was a statute enacted by the legislature, which was signed by the governor, and which produced *L*, and was backed by judicial decisions, jury instructions, the Massachusetts General Laws, and so forth. But why should *these* facts anchor *L*? Because of a secondary rule *R*, which sets up the conditions under which the corresponding primary rules (e.g., *L*) emerge. Of course, *R* must be anchored in turn. We can depict this case as follows:

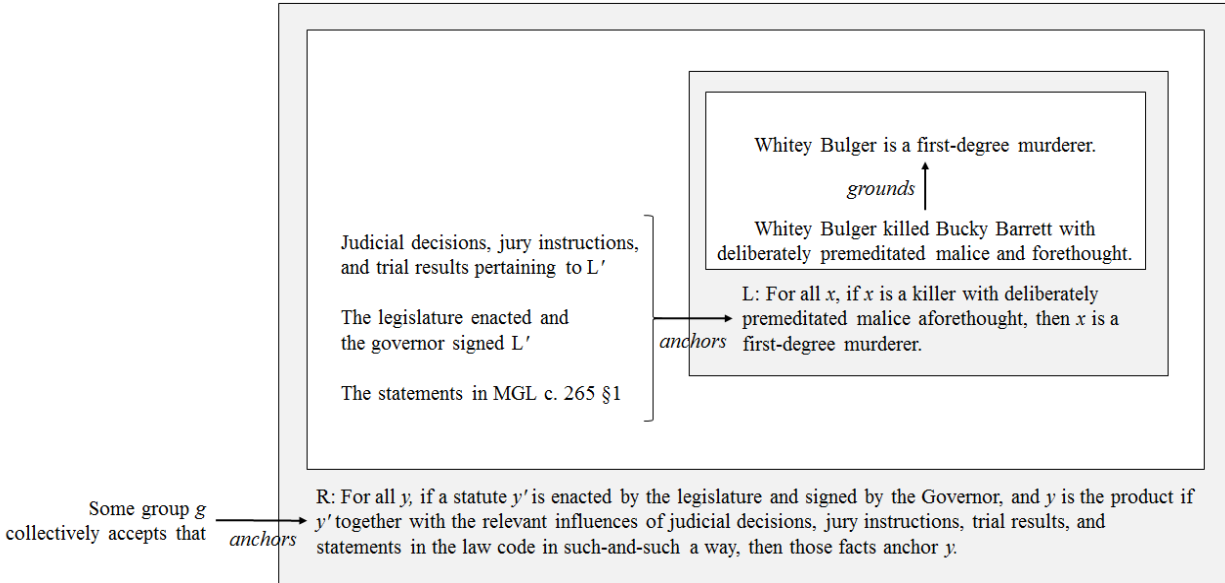


Figure 11 A system of nested legal frame principles

As I mentioned a moment ago, this Hartian distinction between primary and secondary rules mirrors in many respects the distinction I drew in Chapter 1 between first- and second-order deontic phenomena. Accordingly, it should not be too difficult to see how a similar approach might be taken toward how to anchor different kinds of deontic frame principles. For example, suppose that Manny the manager institutes a rule that whoever is more than 10 minutes late to the beginning of their shift will receive a warning. This rule could readily be expressed as a frame principle (**LW**) (for Late to Work):

LW For all x , if x is more than 10 minutes late to work, then x shall receive a warning.

Now we can ask further why Manny’s decision to issue the rule (**LW**) holds—in other words, what anchors (**LW**). Here, we can turn to some grounding facts (e.g., Manny’s being manager being grounded in Manny’s being hired as manager through Declaration), but also to a broader frame principle, for example (**MR**) (for ManagerR):

MR For all x, y , if x is a manager, and x declares that y is a rule, then those facts anchor y .

Now we can ask what, in turn, anchors the frame principle (**MR**). It might be direct acceptance, but more likely (**MR**) itself is anchored by the fact that an official document outlining the powers of company managers declares that (**MR**) holds. This anchoring fact would hold in virtue of a further frame principle (**CD**) (for Corporate Document):

CD For all x , if x is an official corporate document and x declares that some frame principle y holds, then those facts anchor y .

And, of course, this frame principle would need to be traced further still until it bottomed out in the acceptance of a certain complex legal system and perhaps of a system of language itself. The details are not important for present. What is important is to identify one important way in which the Grounding-Anchoring Framework accounts for a complex system of interlocking social institutions, namely, through nested frame principles, which bottom out in some form of collective acceptance. This is an important point. In what follows, I will assume that although a nested frame principle can be anchored in a Declaration—e.g., as (**LW**) is anchored in part due to (**MR**), and (**MR**) in part due to (**CD**)—such nested frame principles can always in principle be traced back to collective acceptance. Accordingly, the subsequent analysis of collective acceptance will be that of what *ultimately* is required to anchor a deontic frame principle, with the caveat that particular frame principles can be anchored in intermediate ways, e.g., through Declaration.

Thus, although for the sake of perspicuity I will continue to deal with relatively simple examples, my intention is that the account I develop below will accommodate more complex, real-world cases of nested frame principles. More immediately, my aim is to lay the groundwork for an account of how second-order deontic phenomena (e.g., deontic powers) can in certain cases be

anchored differently from first-order ones (e.g., obligations), though both will ultimately be anchored in collective acceptance.

This approach is intended to cohere well with the notion of “massively shared agency” found in Shapiro (2014). As Barry Smith (2013, pp. 175-6) summarizes the idea:

[W]hen philosophers, such as Searle, Gilbert or Bratman talk about collective action, they focus on small, local, interpersonal exchange between cooperating agents—agents who are aware of and largely agree about each other’s beliefs and desires—and on actions such as meeting for dinner or going for a walk... [But this approach] falls short in failing to give an account of how the authority of the law becomes operative among civilians who have made no such commitment. Shapiro’s intriguing solution to this problem rests on the modularizing power of the hierarchical division of authority, but shows how such hierarchical division creates a secondary effect in relation those ordinary civilians who fall outside the hierarchy but are still subject to its authority... Here, the *intentions* of civilian participants are irrelevant to the generation of the salient legal authority; what matters are the intentions of other legal authorities. The civilian has a subordinate legal status—is subject to the law—not because he has explicitly accepted this subordination, but because the participants in the legal system accept, are committed to, their shared plans to treat him as a subject.

I agree with Shapiro that accounts like Searle’s and Gilbert’s have suffered from focusing too much on small, interpersonal social exchanges and agreements. One indication of this difficulty in their approach is that, when confronted with larger-scale deontic phenomena in the social world (e.g., the layers of authority and the procedures for investing others with authority), many authors are pressed to maintain that the commitment involved is somehow implicit, unconscious, or behavioral. Thus, as we have seen already, Searle (2010) attempts to ground social deontology in the implicit commitment to using accepted linguistic conventions correctly. Gilbert (2014, pp. 26, 33) likewise suggests that, although explicit agreement is not required in these cases, there is nonetheless a joint commitment among agents in that they are all supposedly seeking to achieve a common social goal.

Like Shapiro, I believe we need to introduce such tiering or enmeshing within our deontology. I have characterized this in terms of frame principles nested inside other frame principles, where

within broader frame principles that are collectively accepted it becomes possible to anchor further frame principles. Thus, we can differentiate between a society that, as an anchor group, collectively accepts its political institutions as a whole (including the broad procedures that are followed to invest someone with political authority, e.g., voting), but which may not directly accept a particular law established by those authorities. This is because what we often find in the social world are not simple, disjointed deontic frame principles, anchored by agents consciously committing to abide by those frame principles. Rather, we find frame principles enmeshed within much larger, existing social frames, and anchored in only an indirect way by broader acceptance within a society.

In this chapter, I have argued that there are diverse ways deontic phenomena can be grounded. I have also attempted to highlight some of the complications that arise in answering the question of which systems of frame principles are anchored in a given society. On this latter point, my main suggestion has been that we adopt a non-uniform approach to such anchoring. My other main point has been that one needs to recognize that real societies adopt complex, nested systems of such frame principles, and that a given society may be comprised of sub-groups anchoring their own competing frame principles. Accordingly, the question of what grounds deontic phenomena in one's own context will be a decidedly empirical matter; it is a question which hinges on first carefully examining the system of frame principles in which one finds oneself. My goal has not been to make claims about the specific systems of frame principles found in a given society, but rather to sketch the structure of social-institutional grounding. In the following chapter, I take this proposal a step further by arguing that, regardless of the differences one might find among systems of anchored frame principles, all such systems depend in some way on *collective acceptance*.

The Dispositions That Bind

In Searle's (1995, 2010) view, institutional facts in general obtain because of collectively accepted constitutive rules. But given the Grounding-Anchoring Framework that I introduced in the previous chapters, how should we understand Searle's view? For starters, we can see that the notion of a constitutive rule corresponds to that of a frame principle: where Searle would say the slips of paper in his pocket count as money because of a constitutive rule governing when something counts as money, the Grounding-Anchoring Framework would say there is a frame principle governing when something counts as money. Indeed, when Epstein (2015) introduces the notion of a frame principle, he remarks explicitly that this is meant as a refined version of Searle's constitutive rules.¹ Similarly, then, we can say that the notion of collective acceptance establishing a constitutive rule corresponds to that of anchoring a frame principle insofar as, from the vantage point of the Grounding-Anchoring Framework, collective acceptance would be one species of anchoring. Thus, within the Grounding-Anchoring Framework, we can say that, for Searle, collective acceptance would be what *anchors* a frame principle.

¹ Epstein (2015, p. 77): "In talking about [grounding] principles, I am shifting away from Searle's term 'constitutive rule' altogether. As we will see, there are many different sorts of principles that give the grounding conditions for social facts. Many of them serve quite different purposes from the ones Searle discusses. The term 'constitutive rule' is so closely associated with Searle's formula *X counts as Y in C*, that it would be very confusing to retain his term for this much more general notion. Also, the term 'constitutive rule' was never a particularly appropriate one. Like the Holy Roman Empire, which was neither holy, nor Roman, nor an empire, constitutive rules are neither constitutive nor are they rules. Instead, I will call these general principles 'frame principles.'"

Accordingly, in what follows, I will mostly follow Epstein in shifting away from Searle's terminology. When discussing and evaluating Searle's views, I will often situate them within the Grounding-Anchoring Framework. Thus, again following Epstein, I take Searle's account of constitutive rules and collective acceptance to be an account of what Epstein and I would call frame principles and anchoring, respectively. Collective acceptance could, then, be described as Searle's answer to the question of what anchors the frame principles for social institutions.

But this naturally raises two questions. First, what exactly is collective acceptance? To begin, Searle clarifies that by "collective acceptance," he does not mean "collective *approval*":

Acceptance, as I construe it, goes all the way from enthusiastic endorsement to grudging acknowledgment, even the acknowledgment that one is simply helpless to do anything about, or reject, the institutions in which one finds oneself. So... to avoid this misunderstanding, I will use "recognition" or sometimes the disjunction "recognition or acceptance." The point is that status functions can only work to the extent that they are collectively recognized. I want to emphasize again that "recognition" does not imply "approval." Hatred, apathy, and even despair are consistent with the recognition of that which one hates, is apathetic toward, and despairs of changing. (Searle, 2010, p.8; cf. 1997, pp. 453-4).

What, then, does it mean to accept a frame principle (constitutive rule)? Unfortunately, Searle's answer to this question is less straightforward than it might first appear. Consider the various ways Searle explains what it means to collectively accept a frame principle. First, he famously claims that "something is money only because we believe it is money" (Searle, 1995, pp. 3, 33), which suggests that belief is a necessary, if not sufficient, conditions for collective acceptance of the relevant frame principle. Elsewhere, however, he suggests that something is money because we engage in buying and selling with it, independently of whether we are conscious of the relevant frame principle:

[T]he process of the creation of institutional facts may proceed without the participants being conscious that it is happening according to this form... They need not think, "We are

collectively imposing a value on something that we do not regard as valuable because of its purely physical features,” even though that is exactly what they are doing. (Searle, 1995, p. 47)

Searle is arguing here that new frame principles can be collectively accepted through repeated behavior (or through acquiring the disposition to behave) rather than through a conscious act of acceptance; the “creation of institutional facts” just is what collective acceptance is supposed to accomplish. But if such collective acceptance needn’t be conscious, then it isn’t clear why collective belief would be a necessary condition for collective acceptance. Thus, these claims suggest that collective acceptance might be a matter of belief, or behavior, or disposition to behave, or perhaps some combination of each, but they do not provide a sufficiently clear picture of what is supposed to be necessary or sufficient conditions for collective acceptance.

The second question is: what exactly is required for an account of anchoring *deontic* frame principles, or the deontic dimensions of frame principles? Given the Grounding-Anchoring Framework, Searle’s remarks on collective acceptance would be construed to be about what anchors social institutions in general. But do we also need an account of what anchors frame principles specifically for deontic phenomena? As it turns out, Searle’s account of deontic phenomena (or what he calls “deontic powers”) builds directly upon his theory of what anchors frame principles (or what he calls “collective acceptance”). As I explained in a previous chapter, Searle takes social institutions to be inherently deontic, because all social institutions are matters of what agents are permitted, obliged, or empowered to do, or immune from doing, within the context of these institutions. Status functions in particular come packaged with what Searle calls “deontic powers”. Thus, the frame principles governing deontic phenomena would likewise be anchored by collective acceptance.

However, Searle’s remarks also raise the question of whether speech acts play a role in the anchoring of deontic frame principles. Searle regularly emphasizes the importance of Declarations in the emergence of social institutions. He even goes so far as to argue that all social institutions (besides

language itself) are brought about through Declarations. Thus, in Searle's view, not only do some speech acts ground particular deontic powers (e.g., an act of promising grounding an obligation or claim), but they also are responsible for establishing the grounding conditions for them. Therefore, one might wonder how exactly speech acts fit into a theory of what anchors deontic frame principles, and especially if all deontic frame principles are ultimately anchored by Declarations.

The goal of this chapter is to defend an account of what is involved in the collective acceptance of a deontic frame principle. In my view, such a frame principle is anchored if a group is collectively disposed to monitor for compliance with the frame principle and to sanction those who defect from that frame principle. Before defending this view, however, I will take a closer look at Searle's language-based account of collective acceptance and offer some criticisms.

Speech Acts and Collective Acceptance

Beginning in (1969), Searle argues that constitutive rules of the form *X counts as Y in C*, which impose status functions on entities, are operative in virtue of collective acceptance. Thus, he writes:

[T]he application of [a] constitutive rule introduces the following features: The Y term has to assign a new status that the object does not already have just in virtue of satisfying the X term; and there has to be collective agreement, or at least acceptance, both in the imposition of that status on the stuff referred to by the X term and about the function that goes with that status. (Searle, 1995, p. 44)

In other words, in order for some entity X to have a status function Y in C, there must be some collectively accepted constitutive rule to the effect that Xs count as Ys in C. We can formulate this a bit more formally as follows:

X1 *X counts as Y in C* iff there is some x such that x is a group and x accepts that *X counts as Y in C*.

Now Searle also aims to spell out the invariably deontic character of social institutions, and argues that each constitutive rule of the form *X counts as Y in C* is equivalent to some constitutive rule of the form *S has power (S does A)*. Thus, (X1) is meant to be equivalent to:

X2 *S has power (S does A)* iff there is some *x* such that *x* is a group and *x* accepts that *S has power (S does A)*

and:

X3 *X counts as Y in C* iff there is some *x* such that *x* is a group and *x* accepts that *S has power (S does A)*³

Recall that, in Searle's view, all deontic phenomena (obligations, permissions, etc.) are a form of what he calls "positive" or "negative" power. Thus, the assignment of any status function—the Y term in the constitutive rule—is equivalent to the assignment of a power to an agent. Now this asserted equivalence between (X1-X3) might seem puzzling, considering that Searle also maintains that (what he calls) deontic powers are not identical with status functions. Fortunately, Searle clarifies his view on the relationship between them: when a group collectively assigns a status function, the intentional content of that assignment is the imposition, direct or indirect, of some deontic power upon an agent.

He writes:

[T]he content of the collective intentionality in the imposition of the status-function will typically be that some human subject, singular or plural, has some power, positive or negative, conditional or categorical. This will be directly the case where the status is imposed on an agent, as in, e.g., Jones is President, and indirectly the case where the status is imposed on an object, as in, e.g., this is a five dollar bill. (Searle, 1995, p. 98)

³ Searle is clear that he takes these two be equivalent.

Thus, in Searle's account, when a group collectively assigns a status function through having accepted a constitutive rule, it assigns some agent a deontic power, either directly or indirectly. The distinction between *direct* and *indirect* is obviously significant here. Searle means that a deontic power is either assigned directly to an agent (e.g., John has the power to lend his possessions to another) or else indirectly to an object in virtue of which an agent exercises that power (e.g., John possesses a dollar bill, with which he can purchase goods) (Searle, 1995, p. 104). This is why the acceptance of a rule of the form *X counts as Y in C* is equivalent to either the direct or indirect assignment of a power. For example, a group might accept a constitutive rule governing when someone becomes someone else's boss (i.e., where *boss* is a status function). Then, this constitutive rule can be cashed out in terms of which deontic powers come packaged with the status function of *boss*. In assigning Basil the status function of boss, we also assign to him hiring-firing power, and so on. Likewise, acceptance of the constitutive rule that certain paper bills count as money can be cashed out in terms of the powers an agent gains when he possesses those paper bills. Thus, Searle's view is that the acceptance of a constitutive rule, which in turn allows for the assignment of status functions, will necessarily have a deontic dimension, in that any assignment of status function is equivalent to the assignment of some deontic status.

At the same time, Searle clarifies that status functions are not identical to any deontic powers. One reason is that an entity can retain the same status function while changing with respect to some (or even perhaps to all) of its associated powers, regulations, etc. For example, if a constitutional amendment abrogates the power of presidential pardon, the President does not thereby lose his status function. At the same time, it is impossible to assign a status function without assigning *some* deontic status. Another reason is that status functions and some deontic powers might belong to very different ontological categories. For example, although Jake's status function of employee might imply that he has an obligation to show up for work on time, and although being obliged to show up for work on

time might be partly constitutive of his status function, the status function cannot be identical to the obligation itself. Jake is declared to be an employee; he is obviously not declared to be an obligation, even if that Declaration carries with it certain deontic ramifications.

Since Searle characterizes all deontic phenomena as deontic powers, it is not entirely clear whether the assignment of status function could be tantamount to the assignment of a permission or obligation to an agent, rather than a power proper. But it seems entirely plausible that this could be the case. For example, for Jake to be assigned the status function of employee is also for him to be assigned the obligation to show up for work and the permission to be on the premises of his place of work. Therefore, in what follows, I will assume that Searle's account can be generalized: collective acceptance anchors the constitutive rules which establish grounding conditions for what we previously described, following Hohfeld, as both first- and second-order deontic phenomena.

Searle remarks that formulas (X1-X3) have the same "logical form" as a Declaration, i.e., a speech act which represents some institutional fact as obtaining, and which in its performance also makes it the case that the institutional fact obtains. Thus, the Declaration, "I now pronounce you man and wife," is supposed to both represent a couple as married *and* make it the case that they are now married. This Declaration, Searle argues, has the same structure as the constitutive rule: "Uttering 'I know pronounce you man and wife' counts as effecting a marriage in the context of a marriage ceremony." As I argued previously, it is unsurprising that some deontic phenomena seem to come into existence because of Declaration—in other words, that they are grounded in Declarations. But do speech acts like Declaration play a similar role in *anchoring* deontic frame principles? For example, is a group's jointly declaring that some frame principle holds necessary or sufficient for anchoring that frame principle?

Searle rightly notes that it would be a mistake to hold that Declarative speech acts are necessary for anchoring social institutions in general:

The institution and the institutional facts within the institution require continued recognition or acceptance because they exist only as long as they are so recognized or accepted. One mark of recognition or acceptance is continued usage of the institution and institutional facts, and this requires the usage of the corresponding vocabulary. Acceptance need not take the form of an explicit speech act and can range all the way from enthusiastic endorsement to grudging acquiescence. This is why I frequently use the expression “recognition or acceptance” to mark the fact that I do not mean to imply approval by the specification of the attitude that institutions and institutional facts require for their continued existence. Recognition in the form of speech acts functions like Declarations, even when the speech acts are not in the form of a Declaration. By continuing to use the vocabulary of the status functions we reinforce the status functions. Unlike shirts and shoes, institutions do not wear out with continued usage. On the contrary, the continued usage of such institutions as marriage, private property, and money reinforces the institutions, but the “usage” requires talk, and that talk functions to maintain and reinforce both the institution and institutional facts within the institutions. (Searle, 2010, pp. 103-4)

What Searle suggests here is that the use of speech acts such as Declaration is, practically speaking, necessary for the continued existence of a social institution. However, he also claims here that what is really central to the emergence and persistence of a social institution is “recognition or acceptance,” which he describes in terms of “continued usage.” A speech act, then, can be thought of as a tool for facilitating collective acceptance, but which is ultimately not responsible for anchoring. That is, a speech act will only be effective if that type of speech act is itself already collectively accepted. Thus, a Declaration can be thought of as an intermediate anchor for a frame principle, but not its ultimate anchor, which will always be collective acceptance. This leads us back, at last, to the question of how exactly collective acceptance works in anchoring deontic reality. To understand Searle’s view on this question, we need to understand his suggestion that language is the primordial social institution, and that linguistic conventions come packaged with deontic force.

Language and Implicit Commitment

Searle maintains that, even though there are various specific ways a deontic phenomenon might be grounded, all of these cases bottom out in facts about a group's linguistic commitments. But I will argue that this view faces two problems: (1) it seems to entail a vicious regress, and (2) it rests upon the premise that collective acceptance entails private commitment, which is false.

Specifically, in Searle's view, participation in a linguistic community (i.e., a community which communicates through conventionally accepted linguistic expressions) entails commitment to using these expressions to mean what that community collectively takes them to mean. For example, to make assertions using conventional linguistic expressions is to have committed to being sincere in one's assertions because the conventional linguistic expression is collectively accepted (Searle, 2010, p. 80). Thus, he writes:

[O]nce we have an explicit language in which explicit speech acts can be performed according to the conventions of the language, we already have a deontology. We already have commitments, in the full public sense that combines irreversibility and obligation. [...] [I]t is inevitable that you will have deontology because there is no way you can make explicit speech acts performed according to the conventions of a language without creating commitments. This is true not just for statements but for all speech acts. [...] All types of speech acts contain an element of commitment. Most utterances are not literally promises; but the type of commitment, including both irreversibility and obligation, for which promising is the paradigm, affects all other sorts of speech acts: orders, thanks, apologies, and so on. (Searle 2010, pp. 82-83)

In other words, conventional language carries deontic commitments with it because agents have committed to conventional language, and thereby incur obligations to make honest assertions, to keep one's promises, and so forth.⁴ In other words, each agent in the group has made a public commitment to the effect that:

C "For any expression '*P*', I commit to meaning *P* by '*P*'."

⁴ Searle is clear that he takes insincere speech acts to still count as instances of that type of speech act. Thus, a false promise, though not paradigmatic, counts nonetheless as a promise. Cf. (Searle 1969, ch. 3).

Thus, in the case of a speech act of assertion, the expression ‘*P*’ would mean that one represents *P* as a fact, or in the case of promising, the expression ‘*P*’ would mean that one represents oneself as carrying out the content of *P*. Moreover, the commitment underlying (C) must be public rather than private, since on Searle’s account only public commitments make one liable to others (Searle, 2010, pp. 81-82). However, Searle likely holds that this public commitment is also implicit, since someone would rarely (if ever) explicitly commit to adopting the conventional meanings of expressions. Thus, Searle’s view seems to be that such public commitments are implicit in the adoption of conventionally accepted linguistic expressions.

Searle’s explanation of why my statement, “I promise to do *A*” obliges me to do *A* is that I have already committed to using “I promise, etc.” to mean that I thereby place myself under an obligation to do the promised action. But this raises the important question: how exactly did I commit to meaning that I promise by uttering “I promise”? As Searle acknowledges, it is unlikely that people will have explicitly committed to this by publicly promising to use language in this communally accepted way. And even if they did so explicitly, this would lead to a regress: I would need to have already committed to the commonly accepted meanings of English words in order to use English to make such a commitment.

So we are left with the alternative that the *use* of words with commonly accepted meanings is sufficient to commit me to using them in accordance with their commonly accepted meanings. But this seems mistaken. A group might collectively recognize that ‘*P*’ conventionally means *P* without it being the case that anybody in their linguistic community is committed to using ‘*P*’ to mean *P*. Searle himself remarks elsewhere that collective acceptance in no way entails collective approval (Searle, 2010, p. 8). It is certainly possible that a group collectively accepts or recognizes that, in their socio-linguistic context, ‘*P*’ is used to mean *P*, without giving any positive endorsement to the fact that ‘*P*’ is

used to mean *P*, and without holding anybody liable for failing to use '*P*' to mean *P*. Accordingly, collectively recognizing that '*P*' is a conventional way of expressing *P* would not imply that anyone in that context has also committed to using '*P*' to mean *P*. A society which failed to have such expectations for all of its collectively accepted conventional meanings might not run very efficiently, but it would not thereby be impossible.

Consider a fictional case. Eddie knows that, within his socio-linguistic context, the assertion "X is a person" means that X is a self-conscious being without severe mental impairments. Eddie, however, personally rejects this meaning of the key expression 'person' as discriminatory. Eddie's society has collectively accepted a conventional meaning for "X is a person," but Eddie, although he adopts most other conventionally accepted meanings, has made no such individual commitment regarding that assertion. It is difficult to see, as Searle maintains, why the mere fact that Eddie's linguistic community recognizes a conventional meaning for "X is a person," that Eddie has thereby personally committed to the conventional meaning for "X is a person". Nor is it plausible that, simply because he regularly participates in linguistic behavior with that community, Eddie has thereby committed to using all of that community's conventional linguistic devices in accordance with their conventional meanings. Note that, in this scenario, it is entirely possible that Eddie is liable to his linguistic community for misleading others with his idiosyncratic use of "X is a person" which includes those with severe mental impairments, but the reason for his liability cannot be that he made a commitment to using the expression in accordance with the conventional meaning. He has made no such commitment, and thus Searle's claim that social institutions are grounded in such commitments is false.

Searle's view suggests that individuals have made commitments when it is highly implausible that they have. However, we might try to rehabilitate Searle's account to some degree by focusing on one aspect of it that certainly does seem relevant to the emergence of deontic phenomena within the

social world, namely, that such deontic phenomena have something to do with wider group expectations regarding people's behavior. One reason why Searle's view carries some plausibility is that basic social coordination (greeting, conversing, buying, selling, etc.) obviously depends on a largely shared understanding of the meanings of words and phrases. Accordingly, we expect the people around us to use words and phrases according to their collectively recognized meanings; we rely on their 'yes' to mean yes, and their 'no' to mean no. In that regard, it is perfectly natural to see linguistic communication as having a normative dimension: there is a correct usage of language, and in some cases an incorrect usage is blameworthy (e.g., using misleading language to harm another person). There will of course be cases where the meaning of a word or phrase is controversial among the language-users (e.g., 'person', 'god', 'gender'), but in general we have a shared understanding of our language.

Where Searle's account falls short is in supposing that this generally shared understanding amounts to a commitment, and thus to an implicit deontology. What his account does help highlight, however, is the way in which the deontic (or normative, or rule-following) dimension of linguistic behavior is sustained by a wider community's expectations that language-users adopt collectively recognized meanings. In the following section, I wish to elaborate on this suggestion in order to present a different account of what ultimately anchors the deontic element of institutional reality, namely, a collective disposition to monitor for compliance with deontic frame principles. But what does it mean to be collectively disposed to monitor? This can be explained by looking at two components, roughly: (1) shared dispositional beliefs regarding the frame principle (or system of frame principles), and (2) shared dispositions to respond negatively to violations of the frame principle.

Collective Belief

It might be thought that collectively accepting a deontic frame principle is a matter of group sharing a belief. For example, for a group to collectively accept that “I promise to *A*” places its speaker under an obligation is to believe that it does. And as I have already noted, Searle informally describes the anchoring of the institution of money in doxastic terms: money is money because we *believe* it is money (Searle 1995, p. 32). More specifically, one might think that collective belief is responsible for anchoring deontic frame principles, e.g., saying “I promise” places one under an obligation because we all believe that it does.

In the first place, one might call into question whether such a collective belief would be *sufficient* to anchor a deontic frame principle. According to the well-known “Humean theory of motivation” (M. Smith 1987), having a belief does not entail that one is motivated or disposed to behave in some corresponding way. Thus, we can at least imagine a situation in which a group collectively believes that all promisors are obligated, but fails to act in accordance with the belief, e.g., if promises were regularly broken with impunity. Then, although the group believes that promises come with obligations, they would behave as though they do not. If, in addition to this (inert) belief, the group needs to be disposed to behave in accordance with the content of the belief, then the belief alone does not seem sufficient for grounding the obligation. Therefore, if it is possible to have a belief without any corresponding desiderative or dispositional state (e.g., motivation to act on the belief), then collective belief would not be sufficient for anchoring a deontic frame principle.

That being said, many authors (Durkheim, 1893/1984; Hart, 1961; Gilbert, 1989, 1996, 2000, 2014; Tuomela, 1995, 2002, 2007, 2013) number collective belief among the necessity conditions for establishing what I have been calling a frame principle (explicitly deontic or otherwise). For this view to be false, there would need to be a case in which a frame principle is anchored, but in which collective belief in the content of that frame principle is not required. To assess whether this view is correct, we

should first invoke the traditional philosophical distinction between *occurrent* and *dispositional* beliefs, where to have an occurrent belief is to be in an overt belief state, and to have a dispositional belief is to be disposed to have such an occurrent belief.⁵ For example, Bill might believe that Cajun catfish tastes good, even if he seldom thinks about Cajun catfish. Thus, Bill would have a dispositional belief regarding Cajun catfish, which would be realized in occurrent thoughts about how Cajun catfish tastes; for the most part, the belief would merely remain, as the saying goes, “in the back of his mind.”

With this distinction in hand, we can see that collective occurrent belief would not be necessary to anchor a social institution. Suppose, for example, that the frame principle in question was roughly **FM**, where Q stands for all the qualities the paper is supposed to have (green color, etc.):

FM For all x , if x is a slip of paper with qualities Q and x is printed by the U.S. Treasury, then x is fiat money.

Few, if any, people occurrently believe the content of **FM** on a regular basis. Their failure to occurrently believe **FM** in no way impugns its anchoring. Therefore, collective occurrent belief does not seem to be necessary for establishing a frame principle.

Collective dispositional belief, however, appears to be a much more plausible candidate for a necessary condition on the anchoring of a frame principle. Even if fiat money users in general do not reflect on what is required for an entity to be fiat money, it is likely that in the back of their minds they believe that something is fiat money if it is a particular kind of piece of paper (or coin) issued by the government. Therefore, it is plausible that whichever group is responsible for anchoring **FM** at least dispositionally believes the content of **FM**.

⁵ By ‘disposition,’ here and throughout, I mean an intrinsic property belonging an entity which can be manifested or realized under a certain set of circumstances, e.g., the disposition of a glass to shatter when struck by sufficient force. Cf. (Martin 1994; Lewis 1997; Bird 1998; Mumford 1998; Molnar 2003). Cf. especially, (Arp, et al., 2015, pp. 101-2).

Again, however, this is a simplified version of the real story, wherein the institution of fiat money is embedded within a complex socio-legal system of frame principles. In fact, the frame principle **FM** would likely be established through a procedure, which in turn would hold in virtue of another frame principle, and so forth, until the system as a whole could be said to hold ultimately though indirectly in virtue of collective acceptance. Social ontology proceeds on the assumption that institutions can be founded in some sense on groups of people, what they believe, how they behave, or how they are disposed to behave. Thus, the purpose of this section to clarify in what way belief might be relevant to the ultimate anchoring of deontic frame principles. The answer that suggests itself is that a group's dispositional beliefs appear relevant, insofar as the collective acceptance—direct or indirect—of a frame principle seems to involve being disposed to affirm that the content of the frame principle is true.

Regular Behavior and Monitoring

Collective acceptance of a system of frame principles might also be thought of as a matter of regular patterns of behavior in a society. Thus, an institution would be anchored by the group's behaving as though the frame principle holds. For example, the institution of money would be anchored by the fact that people behave as though slips of paper are money, e.g., by using them to purchase things and accepting them as valuable and legal tender. Hume seems to have endorsed a view akin to this when he claims that political power, whose origin lies in the offering of incentives to comply with artificial rules of justice, is rooted in the creation of regular patterns of behavior throughout a society.⁶

⁶ Hume (1740/1975, 3.2.7.6): "Here then is the origin of civil government and allegiance. Men are not able radically to cure, either in themselves or others, that narrowness of soul, which makes them prefer the present to the remote. They cannot change their natures. All they can do is to change their situation, and render the observance of justice the immediate interest of some particular persons, and its violation their more remote. *These persons, then, are not only induc'd to observe those rules in their own conduct, but also to constrain others to a like regularity, and inforce the dictates of equity thro' the whole society.* And if it be necessary, they may also interest others more immediately in the execution of justice, and create a number of officers, civil and military, to assist them in their government" (emphasis mine).

Again, however, we can draw a few distinctions. In the first place, we would not want to tie anchoring too closely to what agents are doing at any particular time (particular cases of purchasing or selling), but rather to how people behave with regularity (even if not with frequency), tend to behave, or are disposed to behave under given circumstances. In other words, much like in the case of dispositional and occurrent beliefs, we want to leave open the possibility that a type of behavior can lie dormant for a time, and yet still contribute to anchoring. Moreover, we can draw a distinction between two more specific claims about behavior and anchoring: first that collective behavior is in some sense necessary for anchoring, and second that collective behavior is in some sense sufficient for anchoring.

In my view, the first of these comes quite close to stating a necessary condition, though it needs refinement. To start, what does it mean for a group to *regularly* behave as though a frame principle holds? One possibility is that regular behavior might be interpreted to mean everyone has a *disposition to comply* with the frame principle. That is to say, a frame principle would be anchored only if the group will comply with the frame principle under certain conditions. For example, **FM** would be anchored because people will tend to behave in accordance with the frame principle provided that certain ordinary conditions are satisfied, e.g., they are conscious, they possess money, and they are engaged in a purchasing activity.

But what would it mean, in turn, for a group to be disposed to behave as though a frame principle is operative? It need not be the case that each member always or even frequently complies with the frame principle with respect to their own actions. We can imagine a society of hypocrites where nobody keeps their own promises, yet expects everyone else to do so. In that scenario, each hypocrite would still, on the whole, behave as though the frame principle is operative (e.g., they would readily shame their neighbor for breaking a promise), though they would show no signs of being disposed to keep their own promises. Conversely, a regular pattern of behavior does not entail that

some corresponding deontic frame principle holds. In plainer terms, it might be the case that everybody does what they promise to do, and yet no one believes that everyone is obliged to do what they promise to do, or sanctions anyone who fails to do what he assures others he will do.

Rather, to be disposed to behave as though the frame principle holds seems to mean that group members are inclined to monitor for compliance with the frame principle. For example, the frame principle for promising would be anchored by the fact that members of the anchor group are disposed to monitor whether others are making good on their promises. But all by itself, such a disposition to monitor would not be sufficient, since a group could be collectively disposed to monitor for compliance and yet not be disposed to respond in any way to failures in compliance. Thus, our society of hypocrites could be disposed to merely monitor whether their fellow citizens keep their promises; the promise-breakings might be noted but not receive sanction. And in this case, it is doubtful that the frame principle governing promising would truly be anchored. Central to the concept of promising is that others—most directly those to whom a person has made the promise, but also the wider society in which the practice of promising is operative—are invested in whether a promise is kept, and is in some sense entitled to reprimand the person who fails to keep that promise. In the general absence of a disposition to react negatively to violated promises, it is difficult to see in what sense other people would be invested in the promise, and thus in what sense the promisor would be obligated to follow through.

Is such a disposition to monitor necessary for the anchoring of a frame principle? I would like to suggest that the answer is yes, but with the caveat that the reality of social institutions is inevitably more complicated than this. In particular, we can recognize that once a society has a system of (something like) Hart's secondary rules, or procedures for generating (something like) Hart's primary rules, then there will often arise a practical necessity to install a special group responsible for monitoring compliance with the primary rules (e.g., law enforcement). (This is not to claim that such

a special monitoring group is strictly necessary for such a system of complex rules, but rather that the embedding of frame principles gives such a group its rationale.) When such a rule (a frame principle) is not directly accepted, this opens the door to the possibility that the group as a whole is not directly invested in such a way that they will actually be disposed to monitor for compliance with the rule. In such circumstances, some disposition to monitor will remain necessary for the anchoring of the principle (since the absence of such a disposition would signify an indifference toward the principle), though it might be delegated to a designated subgroup of monitors.

Thus, collective behavior plays a necessary role in anchoring insofar as people must share a disposition to sanction violations of a deontic frame principle. Again, the frame principle governing promises would be anchored in part by the fact that members of the anchor group are disposed to sanction others if they fail to comply with the frame principle, i.e., if promisors break their promises. As the foregoing discussion suggested, the disposition to sanction defectors from a frame principle appears relevant to the anchoring of that principle. But note that, as with the case of the disposition to monitor for compliance, the disposition to sanction those who fail to comply with a frame principle can be delegated to a subgroup (likely in part to those responsible for monitoring for compliance). However, in that case the complex system of frame principles would still bottom out in the collective acceptance of the system, which here I have described partly in terms of the anchor group's collective disposition to monitor for and sanction defections from that system as a whole.

Dispositions and Anchor Groups

Throughout this discussion, I have suggested in an informal way that the dispositions to believe the content of a frame principle, and to behave as if a frame principle were true, were *collective* dispositions, which can be attributed in some sense to the anchor group as a whole. This raises further questions, of course, about what it means to have a collective disposition and what exactly constitutes an anchor

group. For example, must an anchor group have collective intentionality, a shared intention, or a joint commitment?

By way of addressing these questions about the collective dispositions of anchor groups, I will first suggest that the notion of a collective disposition can be understood as analogous to that of collective intentionality (e.g., a shared belief). Typically collective intentionality is described in terms of a shift from “I-mode” thinking to “we-mode” thinking. For example, John intends to move the piano up the stairs, not insofar as he is an individual, but only insofar as he is a member of a group moving the piano up the stairs. That is, his intention is not *that I move the piano*, but *that we move the piano*. Call such cases of collective intentionality *irreducible*. By contrast, in some cases of collective intentionality, the relevant mental states are *reducible* to individual states plus beliefs about others’ mental states. For example, as Searle (2010) notes, the shared belief *that a quarterback must not intentionally ground a forward pass* is reducible to the individual beliefs *that a quarterback must not intentionally ground a forward pass* and *that others believe that a quarterback must not intentionally ground a forward pass*.

In the first place, then, I believe that the dispositional belief that an anchor group has with respect to a frame principle is a reducible collective belief. In other words, when it comes to anchoring a deontic frame principle, an anchor group’s collective belief should be understood as the individual, non-collective dispositional beliefs of the group members. The “we-intentionality” here is really just the individual “I-intentionality” of each member, together with the belief that others have the same individual belief. However, this is also a case where we face some imprecision or vagueness, insofar as what is required for anchoring is that in general the members of a group believe that the frame principle holds. This imprecision is unavoidable. It would obviously be too strong to claim that every member of the group must have the belief; a single defector from a frame principle does not undermine its anchoring. At the same time, low acceptance certainly counts against the anchoring of

a frame principle. Drawing the line at 50% of group members would be too crude. All that can be said, then, is that the dispositional belief must be prevalent within the group.

Similar considerations hold for the group's disposition to monitor and sanction those who (responsibly) fail to comply with the anchored frame principle. Individual dispositions must be found in general within the group, together with the belief that others are similarly disposed. There must be overlap, though not necessarily uniformity, in how the members are disposed. But apart from a strong presence of dispositions to monitor for compliance and to react accordingly, it again becomes very difficult to see why the deontic frame principle would be anchored.

Are Anchor Groups Plural Subjects?

Some authors (Tuomela, 2013; Gilbert, 2014) have suggested that the groups which anchor social institutions must also think of themselves as constituting a “plural subject” that is the subject of the belief or behavior. But with Searle (2010), I think to the contrary that this is only a requirement for groups whose collective intentionality is irreducible (e.g., that of the symphony orchestra intending to play together). To truly play together, the individuals constituting a symphony orchestra must think of themselves as a symphony orchestra, i.e., as a collection of musicians playing together in concert, with a common aim, the production of a unified piece of music. But, Searle and I would agree, social groups in general do not need to think of themselves as being joined in such a common project in order to constitute a group.

One way of seeing this distinction better is to apply an informal test for whether a collective mental state is reducible. Specifically, this test selects two (sets of) statements and assesses whether the one can be translated into the other without loss to semantic content. First, consider the statements:

R1 We intend *that we row the boat*.

R2 I intend *that I row the boat* & believe *that you intend to row the boat*.

It is not difficult to imagine a situation in which **(R2)** is true, and yet **(R1)** is false: you and I are each trying to row the boat, but want the other person to stop rowing. In this case, something essential is missing, namely, the idea that we are rowing *together*, or that each individual intention is only intended as *part of* the group intention.

Now consider the statements **(FM1-FM2)**, for “Fiat Money”):

FM1 The society believes *that paper bills printed by the U.S. Treasury are money*.

FM2 Each member of that society believes *that paper bills printed by the U.S. Treasury are money* and that the other members of society believe *that paper bills printed by the U.S. Treasury are money*.

In this case, there does not seem to be any scenario in which **(FM2)** would be true, but **(FM1)** false. What this suggests is that there is nothing to **(FM1)** over and above **(FM2)**: the society’s collective belief in the institution of U.S. fiat money is nothing more or less than individual member’s believing this and having beliefs about others’ beliefs. More generally, this example suggests that collective dispositional beliefs about frame principles need not be irreducible in the way Gilbert and Tuomela maintain.

Anchor Groups and Group Variables

In the last chapter, I introduced three *group variables* as a way of spelling out features a group may or may not have, namely, whether that group’s constitution depends in one way or another on (1) collective goals, (2) prior social institutions, and (3) intentional opting in. Now many groups will have these features, e.g., a typical corporation. But if we are approaching these group variables from the

perspective of what ultimately anchors a system of deontic frame principles, then we should reject the view that such an anchor group will necessarily have any of these features. In the first place, a group which has such collective dispositions (to believe, to behave) need not share any collective goals, i.e., any goals which they are aiming to accomplish insofar as they are a group. Rather, the individuals constituting this group might have fundamentally incompatible goals they are trying to achieve. The mere fact that the members accept a common set of frame principles for governing social interactions does not entail that they believe that those principles promote a common goal. To see this point more clearly, consider the following fictional example:

Purple World

Purple World is a society composed of two main subgroups, Reds and Blues. As a group the Reds share the goal of promoting the advantage of Reds in Purple World and the disadvantage of Blues, and vice versa for the Blues as a group. As a whole, however, the residents of Purple World collectively accept a frame principle governing the printing of fiat money, which Reds and Blues alike use in their everyday financial transactions.

In **Purple World**, a group collectively accepts an institution of fiat money, but *ex hypothesi* does not share a common goal. In fact, the goals of the respective subgroups are diametrically opposed to one another.⁷ Therefore, it would be a mistake to hold that an anchor group necessarily share a common goal or intention.

Likewise, the existence of such an anchor group does not necessarily depend on previous social institutions—quite the opposite, as some anchor groups are what ultimately make social institutions possible. Lastly, membership in such an anchor group obviously does not require that a

⁷ It might be objected that the Reds and Blues do, in fact, share a goal, namely: promoting one's own advantage and anybody else's disadvantage. But this would be a mistake. Analogously, one might argue that because Democrats and Republicans both want someone to be president, they both want the same one to be president—an example of the quantifier shift fallacy described by Geach in (1980). Or it would be like the case of a boy telling his father, "I'm the fastest," and his father replying, "You're right; I *am* the fastest!" The father's joke relies on ignoring the referent of 'I' in his son's utterance. In a similar fashion, the objection that Reds and Blues share the same goal relies on ignoring the actual content of their goals and the meaning of "one's own advantage," which for Reds is indexed to specifically Reds.

member explicitly opt into the group. In reality, most people simply find themselves as members of a social group, and grow to accept the frame principles anchored by that group without conscious reflection.

For these reasons, I am led to the conclusion that an anchor group need not be a well-structured group, such as a corporation, or a group with a shared goal, or anything of the sort. Rather, an anchor group can be a relatively loose-fitting aggregate of individuals who share dispositions which become manifest in various beliefs and behaviors relevant to a frame principle. For a group to anchor a frame principle just is for that group to share such dispositions.

Putting the Pieces Together

At long last, I am in a position to sum up this account of what anchors deontic frame principles. Many such frame principles are anchored in facts which are anchors in virtue of further frame principles. To reuse an example from earlier, the frame principle governing what counts as first-degree murder is anchored by facts about legislation, judicial decisions, etc., which anchor that frame principle only because of some further frame principle whose purpose is to set up conditions under which new frame principles can be anchored. Such anchoring facts I have described as intermediate anchors, since they rely on additional, more fundamental anchors. In tracing the anchoring of frame principles from one intermediate anchor to another, one will inevitably find that the whole system of nested frame principles is anchored by a group which collectively accepts that system. My goal in this section has been to shed some light on what is required for such collective acceptance.

Drawing together various pieces from the foregoing discussion of collective beliefs and behaviors, the following picture of deontic anchoring thus emerges. Some rules gain their deontic force from facts about what people with power or authority have done (e.g., made a declaration or issued a command), but these facts are intermediate anchors. Anchor groups need not directly

collectively accept such rules or the authority that undergirds them, but they do serve as their indirect anchor by means of accepting a whole system of frame principles, which give rise to further (i.e., nested) frame principles. This provides a way of explaining more formally how collective acceptance is tied to the emergence of obligations, claims, rights, permissions, powers, immunities, and so forth. A group might anchor a frame principle which directly governs the emergence of an obligation (e.g., a frame principle about promising), or it might instead indirectly anchor a frame principle which sets up the conditions under which an individual in authority can bind his subordinates to a course of action. Similarly, it might—whether directly or indirectly—anchor the frame principle in virtue of which that individual gains authority. Just as there is no one way to ground an obligation, so too there is no one way to anchor the frame principles which create the conditions under which such entities can be grounded. That being said, their anchoring will ultimately depend upon the collective acceptance of a system of such principles, which will express itself in widespread dispositional beliefs about that system and the dispositions to sustain that system through monitoring for compliance. As those dispositions wane, that system of frame principles erodes.

The previous few chapters have broached two questions—one about what grounds deontic phenomena like obligation, and the other about what anchors deontic institutions. From the vantage point of the Grounding-Anchoring Framework, I have attempted to develop an answer to these questions at some length. Undoubtedly there is more that could be said about these questions, e.g., concerning the makeup of specific systems of frame principles. The goal of this chapter was, thus, not to propose a comprehensive answer to these questions, but first to demonstrate that the Grounding-Anchoring Framework is a powerful tool for addressing these questions, and second to defend the thesis that collective acceptance should be cashed out in terms of an anchor group's dispositions to believe, monitor, and respond. This leaves one other major task for this dissertation: an account of where exactly these deontic phenomena fall within an ontological picture of the world. What *are* they

exactly? I begin this task by looking to a proposal initially formulated by Reinach, but defended more recently by Barry Smith (2003a; 2003b; 2008) among others.⁸

⁸ E.g., Crosby (1983), Mulligan (1987a), Almeida, et al. (2012), and Brochhausen, et al. (2013).

Toward Deontic Ontology

In *The Apriori Foundations of the Civil Law*, Adolf Reinach suggests that deontic phenomena belong to a metaphysically peculiar category of entity:

Through the act of promising something new enters the world. A claim arises in the one party and an obligation in the other. What are these curious entities (*Gebilde* [more literally, “structures”])? They are surely not *nothing*. How could one eliminate a nothing by waiving or by retracting or by fulfilling? But they cannot be brought under any of the categories with which we are otherwise familiar. They are nothing physical; that is certain. One might rather be tempted to designate them as something psychical or mental, that is, as the experiences of the one who has the claim or the obligation. But cannot a claim or obligation last for years without any change? Are there any such [corresponding] experiences? And further: are not claims and obligations really there even when the subject does not have or need not have any experiences, as in sleep or in the loss of consciousness. Recently one has begun to recognize again, in addition to the physical and the psychical, the distinct character of ideal objects. But the essential mark of these objects, such as numbers, concepts, propositions, is their timelessness. Claims and obligations, by contrast, arise, last a definite length of time, and then disappear again. Thus they seem to be temporal objects of a special kind of which one has not yet taken notice. (Reinach 1913/1983, pp. 8-9)

This passage contains a number of interesting claims and arguments. First, Reinach claims deontic phenomena are “surely not nothing,” but rather “entities” (Crosby’s translation) or “structures” (*Gebilde* in German) of a distinctive type. They exist in some sense; a claim or an obligation is something rather than nothing. When someone makes a promise, the fabric of the social world is somehow changed.

Second, Reinach argues that claims and obligations are entities that do not fit comfortably within the categories of the physical, mental, or ideal (abstract). Even if obligations are grounded in the physical world, the world of brute facts, they obviously have no physical properties themselves. It is admittedly more tempting, Reinach remarks, to attempt to identify claims and obligations with some set of mental states. Indeed, later in *Apriori Foundations*, Reinach criticizes other accounts of deontic reality on grounds that they all fall into the trap of psychologizing obligations and their associated acts of promising, e.g., by trying to reduce promises to the intention to perform the promised action (ibid., pp. 34-46). But even though an agents' mental states might be relevant to which obligations exist, or even to what obligations are, it would be a mistake (he argues) to identify a mental state (e.g., a belief) with an obligation. Lastly, he argues an obligation is not an ideal or abstract object (e.g., the number four), whose distinguishing mark is their "timelessness" (ibid., p. 9), existing "nowhere and nowhen" (Russell, 1912, pp. 55-6). But however abstract an obligation might seem when compared to physical objects or mental states, an obligation has a temporal mode of being: it comes into existence at one time, and ceases to exist at another. Therefore, Reinach concludes, claims and obligations are *sui generis*, "temporal objects of a special kind of which one has not yet taken notice."

Drawing in part from Reinach's account, Barry Smith (2003a; 2003b; 2008) has argued that claims and obligations are indeed special, "quasi-abstract" entities. More precisely, he describes them as "freestanding Y terms," which is to say they are Y terms in Searle's sense (i.e., status functions) but of a sort which lack any corresponding X term (i.e., grounding in the physical world). In this chapter, I discuss Smith's argument for this Reinachian approach and raise the concern that it unduly reifies certain deontic phenomena. I then lay the groundwork for an alternative approach by introducing a few ontological categories that will help us situate deontic phenomena within an ontological picture of the world.

Freestanding Y Terms

In an earlier chapter, I described Searle's (1995) theory of social ontology, which suggests that all status functions (e.g., money, promising) are externally grounded features of concrete physical objects (e.g., slips of paper) or events (e.g., an act of speaking). He also describes this aspect of his theory, equivalently, as the view that social reality is constructed when we collectively accept that X terms (the brute physical reality) count as Y terms (the overlain social reality). Interestingly, one of Searle's examples of this is the case of electronic banking, where the status function of money is imposed, he claims, on the magnetic properties of computer hard drives—the “blip[s] on a computer disk” (Searle, 1995, p. 56)—rather than concrete physical objects like coins, bills, beaver pelts, or cigarettes.

Barry Smith (2003a; 2003b) has objected that such “blips” no more count as money than a record of employees counts as an employee.¹ The magnetic configuration of a hard drive is, in effect, a digital representation of the content of a bank statement, but such a bank statement is a record of one's money, not the money itself. Instead, Smith argues, there are some cases of status functions which exist, and yet lack any corresponding X term with which they could be identified. For example, money is a status function, which can be assigned to coins and slips of paper, but the money of which there is a record in one's electronic bank account is money, despite the fact that it does not coincide with any concrete, physical phenomenon. Thus, there seem to be some “freestanding Y terms,” i.e., status functions which do not coincide with any physical objects. Similarly, Smith (2008) argues, although Searle's “X counts as Y” formula works for imposing the status function of queen upon a chess piece of the sort used in normal play, it overlooks cases of blindfold chess. In a game of blind

¹ As Smith puts it (2003a, p. 26): “The blips in the bank's computers merely represent money, just as the deeds to your property merely record or register the existence of your property right. The deed is not identical with your property right and nor does it count as your property right. An IOU note, similarly, records the existence of a debt; it does not count as the debt. It is an error to run together records pertaining to the existence of free-standing Y terms with those free-standing Y terms themselves, just as it would be an error to regard as the X terms underlying obligations, responsibilities, duties and other deontic phenomena the current mental acts or neurological states of the parties involved.”

chess, there is no board and no pieces, but rather the two players merely call out moves to each other (“Pawn to E4,” etc.). Thus, a game of blind chess involves the movements of kings, queens, rooks, pawns, etc., but there are no concrete, physical objects to which we can point and say, “This piece is the king,” and thus no movement of pieces either. In other words, there is an assortment of Y terms, but no X term.² Because of examples like these, Smith concludes, we need to amend Searle’s theory of social ontology to include the possibility of status functions that have no physical basis.

The introduction of freestanding Y term, Smith argues, also provides us with a way to describe the ontology of entities like “property rights, debts, claims, obligations,” etc. He argues, for example, that much as in the cases of electronic money and blind chess, claims and obligations are social entities which are not identical to any underlying physical objects or events. For example, the deed to a house might “record or register the existence of the property right” (Smith, 2003b, p. 289), much in the same way an electronic record of funds records the existence of money, but that deed is not the property right itself. Similarly, he argues, suppose that a valuable Dutch painting is found in the cellar of the Lucca family. After an extensive 10-year investigation, it is discovered that, a century prior, the accredited owner of the painting bequeathed the painting to the family. Consequently, the family is

² Smith invokes Reinach when he describes a game of blind chess as an example of freestanding Y terms, which he here calls “quasi-abstract entities”: “And the same applies also in the case of blind chess: only here, the giant chess board is absent. What we have instead are images of a chess board in the minds of the players. But these images, too, cannot constitute the game of chess – for they are present also when board and pieces truly do exist... A second alternative answer to our question might consist in the claim that, when we play blind chess, then *there is no game at all*. It is merely *as if* such an entity exist – as in the cinema it is merely as if the represented events were actually taking place in the theater. The players in a game of blind chess are, according to this account, just pretending to play chess, as a pianist may pretend to play the piano by touching the keys but without actually depressing them. This is a doctrine of fictionalism... [Rather, a] game of blind chess is what we shall call a *quasi-abstract pattern*, something that is, like abstract entities such as numbers or forms, both non-physical and non-psychological, but at the same time, through its association with specific players and a specific occasion, tied to time and history. A quasi-abstract pattern thus has two properties that are normally assumed to be incompatible. On the one hand it has no physical parts, and is not able to stand in physical relations of cause and effect. But on the other hand it is an historical entity, which means that its existence is tied to a certain interval of time and to certain actions of specific players. Already Plato would have regarded such a combination of properties as something impossible. For Plato the *forms* are essentially non-historical, indeed atemporal; the objects participating in these forms are essentially bound to time and change. To do justice to phenomena like the blind chess game, we need to recognize that there are entities of a third sort, entities which are both abstract (non-physical) but yet historical (they are tied to time),” (Smith, 2008, pp. 36-37).

issued official documents certifying their ownership of the painting. Thus, the investigation reveals that, prior to the investigation, the family had rightful ownership of the painting, despite the fact that they lacked the official documentation, and nobody alive remembered the bequeathing. The family had a certain property right (i.e., a Y term) for a ten-year span, which apparently could not be identified with any concrete, physical entity (its corresponding X term). Moreover, even when documentation is issued, it stands as a mere record of the right, not as the right itself, which preexisted the documentation by a century.

Consequently, Smith proposes that rights, obligations, debts, etc., are freestanding Y terms. They are status functions—genuine features of social reality—but ones which are not identified with any physical object or process in the way that applies, for instance, to the status functions associated with dollar bills or driver licenses. Echoing Reinach, Smith maintains that they belong to a quasi-abstract, *sui generis* category of entity, one which has simply been overlooked in the history of philosophy. Such a type of entity is neither physical (it has no physical properties) nor psychological (it is not identifiable with any mental states). Moreover, such an entity is not a traditionally abstract entity like a number or proposition, since claims and obligations do not exist atemporally. Therefore, he concludes that we should include within our ontological account a new category of entity, and that the introduction of such a category explains what certain deontic entities (e.g., claims and obligations) are.

In his *Making the Social World* (Searle 2010), Searle responds to Smith’s argument regarding freestanding Y terms. His response relies on understanding status functions as symbolic “placeholders” for deontic powers:

[Such] freestanding Y terms do not bottom out in concrete objects, but they do bottom out in actual people who have the deontic powers in question. So there is no object or person who is the corporation, but there are the president, the board of directors, the stockholders, and others, and the deontic powers accrue to them. A corporation is just a placeholder for a set of

actual power relationships among actual people. The same holds for electronic money and blindfold chess. The owner of the money and the possessor of the queen have the relevant powers. (Searle, 2010, pp. 21-2)

In this passage, Searle seems to be making a number of different claims about freestanding Y terms. First, some status functions are freestanding in the sense that these status functions exist despite the fact that there is no identifiable X term which we could designate as the phenomenon which counts as the status function. Second, no status functions are freestanding in the sense of being wholly unanchored from the world of brute, physical reality, because status functions “consist of” deontic powers (as Searle puts it), which belong ultimately to concrete, physical objects, most notably people. Third, the sense in which status functions consist of deontic powers is that such freestanding status functions are merely placeholders for individuals’ deontic powers.

Understanding Searle’s reply hinges on the question of how to interpret the claim that freestanding status functions are placeholders for deontic powers. On one reading, Searle means that status function words (e.g., ‘corporation’) are merely roundabout ways of referring to deontic powers. Thus, if I say, “John is a member of a corporation,” I am not really asserting or implicating that there is some entity that is a corporation; I am only saying: John has a certain power with respect to some other people, but for convenience’s sake I will use the term ‘corporation’. But if this is what Searle means, then the first claim above—that freestanding status functions do really exist—seems questionable.³ Moreover, if this reading is correct, then it would raise the question of whether all status functions, and not merely freestanding ones, are merely eliminable placeholders for deontic powers. If, on the other hand, Searle means that freestanding status functions do exist, though they are grounded in the brute, physical world by means of their relation to deontic powers, then Searle’s

³ In fact, this reading might be supported by Searle’s suggestion that status functions are nothing more than a “massive fantasy” (Searle 2010, p. 201).

response to the problem of freestanding Y terms becomes somewhat more perplexing, since we do not have a clear picture of what that relationship is. In any case, I will leave Searle's account of deontic powers aside, since I discussed it already in Chapter 2 (see pp. 20-24).

Reference, Reification, and Deontic Phenomena

Perhaps Reinach and Smith are correct that claims and obligations belong to a *sui generis*, freestanding or quasi-abstract category of entity. But in response to this proposal, one might wonder whether claims and obligations are *entities* in their own right at all. Although we have been supposing that deontic phenomena are genuine features of social reality, it is not clear whether nouns like 'claim' and 'obligation' denote discrete entities (a claim-entity, an obligation-entity). Alternatively, a phrase like "Ollie's obligation" might be an oblique way of describing the various events, mental representations, etc., involved in Ollie's promise, rather than denoting a distinct entity standing in relation to Ollie (namely, his obligation).

In looking at cases of claim and obligation, we should not naively assume that 'obligation' denotes a discrete entity, rather than serving the linguistic function of summing up a complex configuration of people, events, mental representations, and more—e.g., my saying the words, "I promise to *A*," my intention to do *A*, your expectation that I do *A*, your belief that I ought to be blamed if I responsibly fail to *A*, and so forth—but none which can be identified as "the obligation itself," standing apart from these entities. Claims and obligations would thus not be "curious entities," as Reinach thought; they would not be entities at all. Reinach and Smith would merely have been tricked by the surface-level language into thinking that our deontic terms denotes deontic entities.

In short, this is a problem of whether terms like 'claim' and 'obligation' denote entities, and whether an ontological account of deontic phenomena needs to regard claims and obligations as special types of entity, or whether instead we can do justice to the ontology of the deontic phenomena

without them. Now framing the problem in this manner naturally raises the question: what are some specific deontic phenomena we want an ontological account to do justice to? I will focus on four.

First, an ontological account of deontic phenomena needs to do justice to the apparent effects of actions wherein agents commit to performing (or omitting) future actions: promising, signing employment contracts, signing mortgage papers, agreeing to terms and conditions, etc. Consider a typical case of promising. Reinach rightly notes that, in an act of promising, it certainly seems as though the world has changed in a few specific ways. I said, “I promise,” and intuitively my promise has *done something* besides sending vibrations through the air. In making you a promise, I have bound myself to you in a new way and have made myself accountable to you. I have acquired a new status and a new task. And because of my act of promising, there seems to be a new fact, namely, *that I am obligated*. Similar phenomena appear in other acts of commitment, e.g., the assumption of a certain role or responsibility, as in the signing of an employment contract. Here, the particular acts one will be bound to perform are likely not specified fully ahead of time, as in a promise. Rather, one is making a commitment to instantiate certain types of acts in particular circumstances. Thus, an employment contract often does not specify every particular act the employee will be required to perform, but rather outlines the role the employee is meant to realize in acts of a certain type. Thus, a typical contract for an entry-level software engineer will not specify every specific task he will perform as part of his employment (e.g., fixing a particular bug in a particular piece of code), but rather the types of acts he is expected to perform both as a member of the organization generally (e.g., a general code of conduct) and as a person with that role within the organization (e.g., the responsibility to follow best practices within his field). But like a typical case of promising, the signing of an employment contract seems to carry with it certain deontic effects: the signee is now required to show up for work and to carry out his duties.

Second, this ontological account needs to do justice to the phenomenon of the authority or power to modify social reality. Simply put, some people seem to be in a position to create, modify, transfer, or remove obligations, permissions, and prohibitions, either for themselves or for others. This is most often expressed through binding directives communicated by one party to another: a stop sign, a placard instructing employees to wash their hands before returning to work, an order from a person in a position of authority (e.g., a military commander, a courtroom judge, or a police officer), and so on. In these kinds of cases, the directive communicated by a person has binding force in virtue of having some degree of power or authority over others. But another interesting set of cases involves the power one has over oneself to place oneself under obligation through an act of commitment. In promising, I exercise a sovereignty over my own future actions by binding myself to a course of action, by limiting my own future permissions. Yet another set of interesting cases concerns the use of authority or deontic power to modify another person's authority. This kind of exercise of authority does not directly modify an agent's obligations or permissions, but rather modifies that agent's ability to modify agents' obligations or permissions. Thus, if a senior manager promotes an employee to junior manager, then the senior manager exercises his authority, not only to impose new responsibilities on the employee, but to invest that employee with powers of his own. And if later the senior manager exercises his authority by firing that junior manager, then the junior manager loses not only the responsibilities tied to his role, but his authority as well. Yet another dimension to this phenomenon involves the exercise of deontic power in transferring rights and responsibilities, e.g., exercising my powers of ownership over a piece of property in selling it to another person. To sum up, then, our ontological account needs to do justice to the power or authority underlying the actions which alter the fabric of deontic reality.

Third, this ontological account needs to capture the various ways in which deontic phenomena are interrelated, in a way that we grasp intuitively. For example, if John promises to mow Mary's lawn,

then not only is he obligated to mow it, but he lacks the permission to slack off instead, and Mary possesses a claim against John that he mow it. Similarly, if Mary possesses deontic immunity from being coerced into practicing a religion, then this entails that John lacks the deontic authority to coerce her into practicing a religion. Thus, an approach to the ontology of deontic reality should also account for the ways in which these various elements of deontic reality seem to be interwoven.

Fourth, this ontological account needs to accommodate the evident temporality of deontic phenomena that is distinct from the temporality of associated mental experiences on the part of person's involved. As Reinach pointed out, claims and obligations “arise, last a definite length of time, and then disappear again” (1913/1983, p. 9). Thus, an adequate description of the ontology underlying promises, contracts, law, etc., needs to include an explanation of the relation of our claims, obligations, and permissions to times. Ontologically, what does it mean to have an obligation for a specified period of time?

Toward an Alternative Proposal

There are perhaps additional deontic phenomena that would need to be accounted for in order to fill out a complete ontological picture. But these four stand out as of central importance to developing an ontology of deontic reality. With all that in mind, I would like to begin to explain an alternative to Reinach and Smiths' proposals: deontic phenomena can be accounted for ontologically, not in terms of *sui generis*, quasi-abstract categories of entity, but in terms of combinations of shared representations and shared dispositions, together with the ways in which those shared representations and dispositions can be captured and communicated in various artifacts (policies, laws, traffic signs, etc.).

As a way of summarizing the general thought behind this alternative proposal, suppose John is driving and approaches an intersection with a stop sign. What is the ontological situation here? At first glance, it might look as if the only relevant entities involved are John, his car, and the sign. But

what is it about that sign that makes John legally obligated to stop at the intersection? In a sense I will spell out more fully later, the sign is an artifact created to bear and convey certain prescriptive content. The content is in part about approaching motorists stopping their vehicles, but it is also about the consequences of failing to stop (e.g., crashing into another vehicle, striking a pedestrian, or getting a speeding ticket). Another element that is essential to this content is the fact that both the artifact and its content were created within the context of an anchored system of legal frame principles, and that the monitoring and enforcement of compliance with the directive has been delegated to a particular group. This is reflected on the level of the content itself: the meaning conveyed by the features of the artifact (the word, the shape and color of the sign, etc.) includes the fact that the directive for motorists to stop is a legally imposed requirement. Thus, ontologically speaking, what we have is a rather complex situation. There is the artifact itself, its prescriptive content as produced within the context of a particular system of legal frame principles, the various components of that system of frame principles (in this case, legislators, law enforcement personnel, procedures for changing laws and for electing officials, etc.), and the underlying collective acceptance of that system. On all these levels, shared representations and dispositions have a role to play.

This account is admittedly a bit messy, but only, I think, because social reality is inherently complex. For exposition's sake, we can focus on artificially simple examples (e.g., a group of people collectively accepting a frame principle about promising), but ultimately we need to grapple with the fact that deontic phenomena appear only within complex systems of frame principles. It may not be easy to parse every entity involved in the requirement that John stop at the intersection, yet even there I believe we can point to all the essential ontological ingredients: there are shared representations of the world, including representations of how people are disposed to react in the face of certain behaviors, which can also be shared in some way by artifacts like documents, legal codes, standards, policies, procedures, service contracts, confidentiality agreements, utility bills, traffic signs, handicap

parking signs, and so on. This is all one needs, I believe, to account in general for the ontology underlying obligations, claims, responsibilities, permissions, authorizations, prohibitions, and authority within social and institutional contexts.

The rest of this chapter is concerned with describing a few of those principal ontological ingredients in greater detail, in particular, what I mean by *representation*, *content*, *prescriptive content*, *disposition*, and *role*. In doing so, I will be drawing heavily on an existing ontological framework advanced in large part by Barry Smith, especially Basic Formal Ontology (BFO), the Information Artifact Ontology (IAO), and the Mental Functioning Ontology (MFO).⁴ This is obviously no coincidence: what I hope to show is not only how these ontological accounts of representation, content, etc., can help shed light on the ontology underlying deontic phenomena, but also why this approach should be appealing to Smith, insofar as it builds upon his existing ontological commitments. After summarizing and (in some cases) elaborating upon this approach, I will in the following chapter explain how these components can be brought together to provide an account of several basic deontic phenomena.

Representation and Dependence

Implicit throughout my discussion of social ontology has been the notion that groups of people have similar representations of the world, e.g., the shared representation of a slip of paper as a dollar bill, tied in turn to dispositions to behave as though the dollar has monetary value. For Searle (1980; 1983; 1984), representation is bound up with the relation of “aboutness” or intentionality, i.e., “that property of mental states and events by which they are directed at or about, or of object and states of affairs in

⁴ Basic Formal Ontology is described in (Arp, et al., 2015). Basic Formal Ontology is also under review for publication as an ISO standard (see: <https://www.iso.org/standard/74572.html>). The Information Artifact Ontology is partially described in (Ceusters 2012), and the Mental Functioning Ontology in (Hastings, et al., 2012).

the world” (Searle 1983, p. 1).⁵ Although Searle regards our intentional states as being about things in reality, he also says explicitly that his analysis of representation and intentional attitudes is only meant to elucidate what he calls the “logical properties” of our intentional states, not their ontological status (cf. Searle 1983, p. 14-15). He goes as far as to say that “there is nothing ontological about my use of ‘representation’” (ibid., p. 12), and prefers to speak about the “structure” of intentional states as consisting of attitudes and (propositional) content, which are bound up with the conditions under which the intentional state or speech act would fulfill its aim (e.g., the conditions under which a desire would be satisfied, or an order carried out), while remaining noncommittal about the corresponding ontology (Searle 1983, pp. 16-18; 2010, pp. 27-30).

Smith and Ceusters (2015) have developed a different approach to representation and intentionality by situating it within a more comprehensive ontological framework. In particular, they look to some recent approaches to general ontological categories, information artifacts, and mental functioning in order to present an account of ontological dependence, aboutness, cognitive representation, and the ways in which the content of our cognitive representations can be shared among artifacts like documents and databases.

In the first place, then, their approach rests on a distinction between two types of *ontological dependence*. Traditionally an entity A is said to depend ontologically on B at time t if A cannot exist at t unless B exists at t .⁷ To give a few typical examples: the color of John’s hair depends on John’s hair, since John’s hair color cannot exist unless John’s hair also exists; John’s belief states depend on John’s brain, since his belief states cannot exist unless his brain exists. It is important to note that ontological

⁵ In the modern era, this idea can be traced to Brentano, although Brentano, unlike for example Searle, regards intentionality as essential to every mental state (cf. 1874/1995, p. 89).

⁷ This notion of ontological dependence can be traced back to Aristotle’s distinction between substances and their attributes (“accidents”), which depend on substances (*Categories* 2a11-2b6). For more recent discussions of ontological dependence, and especially of its relationship to the notions of grounding and supervenience, see Fine (1994), Hoffman and Rosenkrantz (1994), Lowe (1998, ch. 6), Schneider (2006), Nolan (2011), as well as the essays contained in Correia and Schneider (2012) and Hoeltje, et al. (2013).

dependence is not a relation of causal dependence. For example, the shape of a statue depends ontologically on the statue, since that shape quality could not exist unless the statue itself existed. But the statue is not the *cause* of the shape. Rather, the shape of a statue depends causally on its sculptor, since it would not have that shape if the sculptor had not given it that shape. But this is not ontological dependence, since the statue's shape can continue to exist long after the sculptor has ceased to exist.

Some authors, however, have distinguished between two species of dependence: *specific* and *generic* dependence.⁸ Arp, et al. (2015, pp. 95, 105) describe this distinction as follows:

A specifically dependent continuant is a continuant entity [i.e., an entity that endures through time, but which exists totally at any given time it exists] that depends on one or more specific independent continuants for its existence... Examples of specifically dependent continuants include the color of this tomato, the pain in your left elbow, the mass of this cloud...

[S]pecifically dependent continuants are thus subject to what we might call the axiom of nonmigration: they cannot migrate from one bearer to another. Some dependent continuants seem, however, to be capable of such migration, as, for example, when you copy a pdf [*sic*] file from one computer to another. Clearly the pdf file is *dependent* on some bearer; for the pdf file to exist, there must be some physical storage device on which it has been saved. But equally clearly, the pdf file can be *moved* from one storage device to another...


To do justice to this and many similar phenomena [our ontology] incorporates the category of *generically dependent continuant*, defined as a continuant that is dependent on one or other independent continuants that can serve as its bearer.

Specific dependence is the most familiar type of ontological dependence, since in most cases, if *A* depends on *B*, then *A*'s existence depends specifically on *B*'s existence. Again, if a statue ceases to exist, then its shape quality ceases to exist as well. It is not as if that shape quality (i.e., that *particular* shape quality, not that *type* of shape quality) could persist by leaping from one statue to another. Note that the above definition leaves open the possibility that *A* might depend specifically on both *B* and *C*, such that if either *B* or *C* ceases to exist, then *A* ceases to exist.

⁸ The distinction between specific and generic dependence is defined in Basic Formal Ontology (BFO) (Arp, et al., 2015). For other discussions of specific and generic dependence, cf. B. Smith (1982; 1992), Arp and Smith (2008), Jansen (2008), and Galton (2014).

In other cases, however, although an entity's existence does depend on some other non-dependent entity, it does not seem to be dependent on any particular non-dependent entity. Thus, when Leo Tolstoy wrote *War and Peace*, he wrote a novel which depends for its existence upon some material bearer. There is no *War and Peace* unless there is some physical or electronic copy (unless somebody has committed all 1,400 pages' worth of content to memory). Thus, *War and Peace* depends ontologically upon there being some copy of *War and Peace* in some medium (e.g., patterns of ink in a book, patterns of magnetism on the sector of hard drive composing a PDF file). At the same time, the existence of *War and Peace* does not depend upon any particular copy. If my copy is destroyed, then the content can nevertheless persist in your copy.

This distinction between types of dependent continuants, and the nature of their dependency, can be summed up in the following diagram, which depicts a set of these entities at two times:



In this diagram, all the specifically dependent continuants (SDCs), generically dependent continuants (GDCs), and independent continuants (ICs) exist at t_1 , but at t_2 , ICs 1, 2, 4, and 5 cease to exist (marked in grey). This, in turn, has ramifications for whether their associated SDCs or GDCs continue to exist. Specifically, SDC 1 ceases to exist because its sole bearer ceases to exist; SDC 2 ceases to exist because one of its bearers ceases to exist (though note that IC 3 continues to exist), and it was dependent specifically on the existence of both; GDC 1 ceases to exist because its sole bearer ceases to exist, but GDC 2 survives because at least one bearer remains (namely, IC 6).

Representation and Content

In light of this distinction, the developers of the Information Artifact Ontology (IAO) have defined an *information content entity* (ICE) as a generically dependent entity (more specifically, a *continuant* entity) that depends on a material entity and that stands in a relation of aboutness toward some entity.⁹ Ceusters (2012) and Smith and Ceusters (2015) propose revising this definition, such that an ICE stands in a relation of aboutness to “some portion of reality,” which includes not only entities, but also universals, relations, and configurations of entities (facts, obtaining states of affairs). On this approach, then, the content of a sentence or a PDF file is not taken to be abstract and atemporal proposition,¹⁰ but rather something at once dependent on a material bearer of the content and capable

⁹ The category *information content entity* was developed within the Information Artifact Ontology (IAO), though it is not entirely clear how to determine its authoritative definition. There is a version encoded in the Web Ontology Language (OWL) and published at the OBO Foundry (<http://www.obofoundry.org/ontology/iao.html>) and on GitHub (<https://github.com/information-artifact-ontology/IAO>), where the term “information content entity” is defined more or less informally as: “A generically dependent continuant that is about some thing.” The definition I invoke here is based on Smith and Ceusters (2015, p. 1).

¹⁰ That is, this approach to content differs significantly from the traditional view among analytic philosophers that the semantic content of sentences are *propositions*: abstract entities that are the bearers of truth and falsity, and that serve as the object of our beliefs, desires, and other psychological attitudes. That being said, whether propositions exist, and if so what they are, is of course controversial. However, any further discussion of the debates surrounding propositions would take us far afield from the subject at hand. Philosophers in the analytic tradition typically trace this view back to Frege (1892/1952), who suggests that the “sense” of a sentence is a “thought,” by which he means “not the subjective performance of thinking but its objective content, which is capable of being the common property of several thinkers” (p. 62). It may be possible to trace this notion of proposition as far back as Seneca, who highlights the distinction between the material bodies the senses observe and the content our minds think about following an observation. Such content, he

of migrating between material bearers. Smith and Ceusters (2015) further suggest that ICEs are *concretized* by the qualities of their material bearers, where a *quality* is taken to be a specifically dependent continuant entity that is “fully exhibited or manifested or realized in” some independent continuant entity (Arp, et al., 2015, p. 96). Thus, a piece of content might depend generically upon a particular copy of a book, but it would also be concretized in the patterns of ink inscribed onto (and specifically dependent upon) the pages of that book. In other mediums (e.g., digital), that same generically dependent content would get concretized in other patterns of qualities.

That being said, it is important to note that the aboutness relation between an ICE and a portion of reality is derivative.¹¹ To borrow an example from Smith and Ceusters (2015, p. 2), suppose a nurse measures a patient’s pulse rate, and then records it in the patient’s medical chart.

When the nurse enters the string *72 beats per minute* in the medical chart of the patient, then there is an ICE that is concretized in the ink (or pixel) pattern exhibited on the chart, which inherits its aboutness from the aboutness of what we shall call the nurse’s *direct cognitive representation* of the pulse. The latter is a (binary) relational quality; it links the nurse causally to the target of his observations. It is on this basis that, by entering data, he creates an ICE that is also tied relationally to its target in reality.

As the above example notes, the aboutness of the ICE the nurse creates is in some sense derivative: the pattern of ink (or pixels, or what have you) comes to concretize the information content entity that is about that patient’s pulse rate because the nurse entered that data as a record of his observation (i.e., a cognitive representation). The nurse’s cognitive representation of the pulse rate has

says, “is not body, [but] a certain declarative fact [*enuntiativum*] concerning body—called variously an ‘utterance,’ [*effatum*] a ‘declaration,’ [*enuntiatum*] a ‘statement.’ [*dictum*]” (*Ad Lucilium*, 117, 13; cf. Seneca 1925).

¹¹ Smith and Ceusters (2015, p. 3): “We shall presuppose in what follows that information artifacts do not bear information in and of themselves, but only because cognitive subjects associate representations of certain sorts with the patterns which they manifest. We thus view the aboutness that is manifested by information content entities in accordance with the doctrine of the ‘primacy of the intentional’ [Chisholm], according to which the aboutness of those of our representations formulated in speech or writing (or in their printed or digital counterparts) is to be understood by reference to the cognitive acts with which they are or can in principle be associated.” The now common distinction between “intrinsic vs. derivative” intentionality can be found, among other places, in Searle (1980; 1983; 1984) and Fodor (2009).

intentionality in the primary sense since the mind characteristically has the capacity to direct itself toward reality. By contrast, the medical chart is only about the pulse rate in a derivative way, i.e., because those patterns of ink have been in some way imbued with meaning from the outside by agents capable of primary intentionality. Or we could think of the example of a correctly functioning and attached ECG machine registering the patient's heart rhythm. In this case, the meaningfulness of its recordings would hinge on the possibility of minded agents regarding them as meaningful.

Accordingly, Smith and Ceusters' approach to the ontology of content, aboutness, and information artifacts adds an important qualification to IAO's original characterization of ICEs by explaining more clearly the relationship between ICEs and representations. A representation they define as "a quality which *is about* or is intended to *be about* a portion of reality" (Smith and Ceusters 2015, p. 3). The most significant feature of this definition is that it describes the relation of aboutness as holding not between a portion of reality and an ICE (a generically dependent continuant), but between a portion of reality and a representation (the specifically dependent continuant that can concretize that ICE).¹² That being said, ICEs still have a role to play in this revised approach. Returning to the earlier example, there is still an ICE concretized by the medical record and capable of migrating (or being copied, printed, etc.), and thus of being concretized in the qualities of other material bearers, including a person's cognitive representations. The concretizing qualities of artifacts, then, may be said

¹² Based on Smith and Ceusters (2015), it is not entirely clear whether part of their proposal is that information content entities do not, strictly speaking, stand in an aboutness relation to anything, or whether they stand in a different kind of aboutness relation (e.g., secondhand aboutness?), or whether this is necessary any longer. This concern arises because of the definitions and axioms found toward the end of their paper. On page 4, they define the "is_about" relation such that its domain is "representations." Now a representation, according to their definition on page 3, is a type of quality, which is a type of specifically dependent continuant. If no entity is simultaneously a specifically dependent continuant and a generically dependent continuant, and if information content entities are generically dependent continuants, then no representation could also be an information content entity. This means, in turn, that information content entities would fall outside of the domain of the "is_about" relation, as defined. Therefore, at the end of their article, it is not entirely clear in what sense information content entities would be (directly) about portions of reality. Again, one possibility is that they are only indirectly about reality in the sense that they are related to agents' cognitive representations, or perhaps because of their concretization by (non-cognitive) representations. Alternatively, perhaps information content entities stand in a different sort of aboutness relation, with a domain inclusive of generically dependent continuants.

to be “conformant to” a cognitive representation if they both concretize the same information content entity (cf. Smith and Ceusters 2015, p. 4).

Another notable feature of this definition is that it allows for a representation to fail to be about a portion of reality. For example, if John sincerely believes *that Hamilton was a President of the USA*, then this belief as a whole may fail to be about a portion of reality, even if components of his belief are about portions of reality (e.g., Hamilton). That is to say, although the relation of aboutness is veridical, a given representation is not necessarily veridical, though a representation is only a representation *of* a portion of reality if it is veridical (i.e., if it is *about* that portion of reality).¹³

A final, closely related component of this account is how they accommodate cases of “misinformation” (2015, pp. 2-3) and non-veridical (or non-referring) representations. On their view, if I believe that *Barack Obama is president of Russia*, then taken as a whole my belief does not represent any actual configuration in reality. However, this representation can be broken down into a number of “representational units” (ibid., p. 4)—*Obama, the role of president, Russia*, and so forth—each of which represents a real portion of reality. Thus, although my belief as a whole fails to be about a portion of reality, it contains some genuinely representational units, each of which represents some portion of reality (entity, type, relation, etc.).

Representation, Direction of Fit, and Prescriptive Content

By the above definition, representations must at least be intended to be about a portion of reality. But this raises an important question about the role representations play when we aim not so much to represent the world as it really is (or was), but to change the world in accordance with our desires, intentions, requests, instructions, commands, and so on. That is to say, we might consider cases where

¹³ Cf. the definition of *is_a_representation_of* in (Smith and Ceusters 2015, p. 4).

a representation or content is seemingly not meant to fit the world, but rather where the world is meant to change to fit one's representation of it, or what we earlier called "world-to-mind" and "world-to-word" directions of fit (cf. Searle 1975, pp. 345-6). It is one thing for me to believe or assert that I am drinking coffee, and another for me to desire or intend that I drink coffee, or to command John, "Go drink coffee." If in these latter cases I am not attempting to represent an actually obtaining configuration within reality, but rather trying to bring about a new configuration within reality in accordance with my representation, then one might wonder if it is a mistake to define representations in terms of their (intended) veridicality. Put another way, this is a question of whether Smith and Ceusters' view rules out the possibility of our representational content having what Searle calls a world-to-mind (or world-to-word) direction of fit.¹⁴

This question is important because many deontic phenomena intuitively seem to involve such a world-to-mind or world-to-word direction of fit. As Searle (e.g., 1975, p. 346) notes, this holds true

¹⁴ In the more traditional analytic view, this is all accounted for in terms of diverse *psychological attitudes* toward a *proposition* (or *propositional content*), understood as the object of the mind's intentional *states* (or *modes*). (For a typical example of this traditional approach, see (Quine 1956) or (Velleman 1992).) Thus, we have on the one hand the proposition *that Washington is the first president of the USA*, and on the other an array of possible attitudes toward that proposition: a student's *belief* that he is the first president, an eighteenth-century person's *desire* that he be the first president, and so forth. Thus, the mind is directed intentionally toward propositional content, and the type of intentionality is then specified as belief, desire, etc. But this view does not cohere well with Smith and Ceusters' approach, for a few reasons. Most importantly, as we have seen, Smith and Ceusters characterize intentionality as a relation between cognitive representations (or, derivatively, content entities) and portions of reality.

By contrast, as Searle (1983, pp. 16-18; 2010, pp. 27-30) has pointed out, the traditional analytic approach regards propositional content itself as the object of intentionality: my belief that Washington was the first president turns out not to be about Washington, his presidency, the temporal priority of his presidency relative to other presidents, or the configuration of all those elements, but about an abstract propositional entity. Thus, we would end up with a view on which cognitive representations stand in a relation of aboutness toward content, but which would then (presumably) stand in some unspecified relation toward a portion of reality itself. This may or may not be a defensible view to adopt—for his part, Searle (1983, pp. 16-18) suggests it involves a fundamental confusion over the nature intentionality—but it does not seem to cohere well with the approach I have been developing thus far. Furthermore, this approach would put a strain on the veridicality (or intended veridicality) constraints around the relation of aboutness, as described by Smith and Ceusters. In a certain sense, this traditional analytic account would likewise guarantee that the object of our intentional states exists in reality: if the object of my intentionality is a proposition, then even if I believe a false proposition, the proposition nonetheless exists. However, the intent behind Smith and Ceusters' account of aboutness is for the mind to cognitively represent entities, relations, configurations, etc., as found in reality. As Smith and Ceusters (2015, p. 2) point out, however, our representations are not typically representations of information content entities, but of the reality which content is about (e.g., Washington, his presidency, etc.). Thus, although this traditional approach easily accommodates cases of world-to-mind direction of fit, it only does so because none of our intentional states are not necessarily veridical in the first place. We are still left with the question, then, of how to account for these cases within the approach outlined by Smith and Ceusters.

for speech acts like Directives and Commissives. A superior's command is not about what his subordinate is or was doing, or even necessarily about what he believes he will in fact do, but rather is in some sense about what he should do. In fact, he is commanding an action be performed only if the state of affairs in which it is performed does *not* obtain (yet). Or what if the subordinate never complies with the Directive? Can the content of the superior's command be in any way directed toward a state of affairs in which the subordinate performs the action if that state of affairs never obtains? The veridicality constraints around Smith and Ceusters' account of representation suggest the answer is a resounding: No.

Fortunately, I believe Smith and Ceusters' approach can accommodate cases of world-to-mind and world-to-word direction of fit without too much trouble. The most important thing to note here is a certain asymmetry between cases of mind-to-world (or word-to-world) and world-to-mind (or world-to-word) directions of fit. The former case is, in a certain sense, exclusively representational, in that one is not trying to do anything other than represent the world the way it is at a given moment. The latter, by contrast, involve more than representations. Phenomena like desire and intention do involve representations of reality, but they also involve (among other things) dispositions to behave.¹⁵ If I desire to drink a cup of coffee, then I am not attempting to veridically represent the state of affairs in which I am drinking coffee. Rather, I am representing certain portions of reality (me, coffee, the activity of drinking, the anticipated effects of coffee, etc.), and I also have a disposition to go get some. And besides that, there are also undoubtedly a host of mental or bodily processes involved in the fact

¹⁵ For his part, Searle avoids this problem by differentiating between representation and intentionality. He does not regard cases of world-to-mind or world-to-word direction of fit as cases where the intentional state is *about* a non-obtaining state of affairs. For example, my desire to drink a cup of coffee is not *about* a non-existent (cf. 1983, p. 17). At the same time, my intentional state of desire *represents* its own "conditions of satisfaction," i.e., the state of affairs in which I am drinking coffee. This is because, for Searle, representation is not the same as intentionality, and so does not play the same ontological role as it does within Smith and Ceusters' account. Rather, "representation" is merely linguistic shorthand for a complex of "logical notions" he uses to explicate his theory of speech acts and intentional mental states. (Cf. 1983, p. 12).

that I desire to get some.¹⁶ Thus, the ontology underlying phenomena like having a desire and intention would simply involve going beyond representations.

Put otherwise, desire involves representations, or more specifically what Smith and Ceusters call representational units, but it also involves more than representations. In the case of my desire, the representational units are about portions of reality relevant to the desire (*me, a cup, coffee, the activity of drinking*, etc.). One thing that is important to note here is that representational units need not be about particulars; they might instead be about types (universals), relations, or configurations of entities, relations, and types within reality. Thus, I can represent myself, my particular beige cup, or a particular portion of coffee sitting in my coffee pot. But even if there does not yet exist a *particular* activity of drinking for my representation to be about, that representation might be about that *type* of activity and about the relation of instantiation. We might even describe my desire as aiming at the instantiation of that type of activity under the right conditions. There are two kinds of directedness existing side by side: intentionality not only in the sense of my mind's representations of the world, but also in the sense of the aim of my action.

What about cases of information content entities which seem to have a world-to-word direction of fit (e.g., a document that details how to assemble a bookcase)? The Information Artifact Ontology (IAO)¹⁷ might suggest one way to address this question. IAO includes a subclass of *information content entity* called *directive information entity*, which in turn includes subclasses such as *action specification* and *objective specification*. In each case, the intent here appears to be to describe species of content entities whose representations are in some sense actionable, much like in Searle's cases of

¹⁶ A fuller account of the ontology of desire falls outside the scope of this dissertation. But, taking a cue from the Mental Functioning Ontology, I believe we can point to various dispositions and processes that would be involved in, e.g., my desire for a cup of coffee. These could include dispositions like my preferences, or processes like wanting or planning. Cf. (Hastings, et al., 2012) and the version of the Mental Functioning Ontology published on the OBO Foundry website (<http://www.obofoundry.org/ontology/mf.html>).

¹⁷ Unless otherwise indicated, the terms and definitions are taken from the version of IAO available at the OBO Foundry website (<http://www.obofoundry.org/ontology/iao.html>) and on GitHub (<https://github.com/information-artifact-ontology/IAO>).

world-to-word direction of fit. Thus, *directive information entity* is defined as: “An information content entity whose concretizations indicate to their bearer how to realize them in a process.” (Cf. also Ceusters’ identical definition in (2012, p. 3).)

IAO’s approach is promising insofar as it aims to capture prescriptive, world-to-word content while maintaining the (intended) veridicality of the aboutness relation.¹⁸ That being said, I do not find it feasible to adopt this particular definition of directive or prescriptive information for my purposes. One issue with this definition is its suggestion (as formulated) that the ICE’s concretizations indicate to the bearer of the ICE how to realize them (i.e., the concretizations) in a process. If, as is currently the case in BFO and IAO, concretization only holds between generically dependent continuants and qualities (cf. Smith and Ceusters 2015, p. 4), then the concretizations of that ICE are not the kind of thing that could be realized in processes, as qualities are not realizable entities (cf. Arp, et al., 2015, pp. 95-9).¹⁹ Another potential issue is that if the bearer of this ICE is an artifact, then, according to this definition, either the concretizations of that ICE in the qualities of the artifact indicate to their bearer—i.e., the artifact itself—how to realize something in a process, or the ICE is not a directive ICE at all. If the former, then the concretization of the directive content of an instruction manual would indicate to the manual itself (i.e., its bearer) how to realize a process, which seems mistaken. If the latter, then we need something other than *directive information entity* to shed light on the sense in which an artifact like an instruction manual can have content that prescribes an action or set of actions.

Instead I would like to suggest that, in order to do justice to such cases of prescriptive content, we will need supplement Smith and Ceusters’ account with an account of ICEs that include some notion of the *addressee* of an ICE, its *purpose*, and its *force*. These are elements which belong potentially

¹⁸ Unlike, for example, the approach to directive information entities as found in the Common Core Ontologies (<https://github.com/CommonCoreOntology>). Cf. Jensen, et al. 2018.

¹⁹ Barry Smith tells me (in personal correspondence) that, in the future, BFO will allow processes to serve as concretizations. Even then, however, the definition provided in IAO would need to be revised: the directive information entity could theoretically be concretized in a process, but not concretized as a realizable entity.

to many types of ICEs, but which will also help highlight what is peculiar about prescriptive content in particular. What is pivotal to note here is Smith and Ceusters' claim that ICEs are not abstract entities, but entities with a particular context and history. The record of the patient's pulse, for example, is meaningful only because it was created by the nurse *as* a recording of that pulse. If the nurse had written the same words, "72 beats per minute," but in a radically different context and for a different purpose (e.g., as a character's lines in a novel she's writing), it would not be the same ICE. Part of what makes it that particular ICE is the context and circumstances under which it was created.

Addressee, Purpose, and Force

First, then, I would like to propose that most ICEs have an *addressee* or intended *audience*.²⁰ In some cases, this is a particular individual (e.g., if John sends a text message to Mary, or if I write a reminder to myself) or group (e.g., if a teacher sends an email to her students). In other cases, the addressee or audience of an ICE is not as particularized; it might be totally unrestricted, or else addressed to a particular class of people. For example, a traffic sign is not addressed to John or Mary in particular, but to any and all motorists driving down the road. Or, to return again to Smith and Ceusters' example, the nurse who records the patient's pulse rate creates an ICE meant for any healthcare professional charged with that patient's care, as well as for the patient's information about his own state of health.

Second, most ICEs have a discernible *purpose* or *point* that is reflected in the ICE itself. In specifying what an ICE or its concretization is about, we need to take into consideration what the content is intended to communicate. In many cases, the purpose is to make a claim or assertion about

²⁰ Throughout this section, I say "most ICEs," as opposed to "all ICEs," because of some potential edge cases. For example, suppose John has a cognitive representation of Mary, and then idly writes down on a piece of paper, "Mary". In this case, he has apparently created an ICE, concretized by the pattern of ink on that paper, but he may have no particular addressee or purpose in mind. It is an atomic representation. In most cases, ICEs are created to assert, argue, persuade, prescribe, request, instruct, command, declare, etc., which is to say, with an inherent audience, purpose, and force (e.g., John's request to Mary that she call him), as I describe below.

the way the world is or was. Again, the nurse who records the patient's pulse rate is, in the act of creating the ICE, making an assertion about the patient's pulse rate at that time. By contrast, the patient's doctor may write "Verapamil, 80mg, 3x/day" with the intent, not of describing what medication the patient is taking, but of prescribing what the patient should take. The two ICEs differ not only in what they are about (a pulse rate and a medication, respectively) but also in what their creator is trying to accomplish (faithfully record information about the patient vs. guide the patient's actions).

This kind of distinction is also illustrated in Anscombe's (1963, p. 56) example of the private investigator who follows around a man with shopping list and records all his purchases. At the end of the day, the two men end up with the same items written on the list, but with two different meanings: the one a record of what the man purchased, the other a directive regarding what he was supposed to have purchased. The ultimate reason for the difference in content is what we have already seen Smith and Ceusters point out, namely, that because the aboutness of an ICE is derivative, its aboutness depends upon the particular context in which it was created. The private investigator created an ICE meant to describe a shopper's actual behavior, whereas the author of the shopping list created an ICE meant to guide the action of the shopper by indicating which items needed to be purchased in order to accomplish the objective of grocery shopping.

In creating ICEs of this sort, their creators perform, not a speech act, but a "document act" (Smith 2012b; 2014); they are trying to "do [something] *with* a document" (Smith 2014, p. 20), and like in the case of speech acts, their acts have something akin to an "illocutionary point" (Searle 1975, pp. 345-6). What I am proposing is that, given the vital connection between an ICE and the circumstances surrounding its origin, this point or purpose is an element inherent to the ICE itself. It is something communicated, or meant to be communicated, when that ICE is concretized.

This is a claim that one might find objectionable. What exactly is the relation between this purpose and the ICE? Is an ICE *about* its purpose? The answer, I believe, is that if one specifies what an ICE (or its concretization as a representation) is about, one will also specify its purpose. If John texts Mary, “Call me,” then John has created an information content entity (a text message).²¹ Much like in the case of the nurse recording the patient’s pulse rate in his medical record, John has generated an ICE, a text message, invested with derivative intentionality. At the very least, the representational content involved is about John, Mary, and phone calls. But it is also about the fact that he wants her to call. In addition to enumerating the entities referred to in John’s text message, we can also identify the point or purpose of that content (getting Mary to call him) as a feature of the content itself. In creating this ICE, John has invested certain patterns of pixels with content representing not only himself or Mary, but also his goal in sending the message. This is not something captured in any individual word of the text message any more than Mary or John are explicitly mentioned, and yet it is an element of the content entity itself insofar as it is meant to be conveyed to Mary.

One test to apply in these cases is whether understanding the point or purpose of an ICE is essential for concretizing that ICE. In general, failing to understand what an ICE is meant to be about would count against a person concretizing that ICE as a mental representation. For example, if a nurse records John’s pulse rate as “72 beats per minute,” and John later reads his medical record and misinterprets “72 beats per minute” as referring to the tempo of a song, John’s cognitive representation would fail to represent what the patient’s pulse rate was at the time it was measured. Similarly, if John misinterprets the prescription “Verapamil, 80mg, 3x/day” as a warning that ingesting this substance in that quantity three times in a single day would be lethal, then he has failed himself to represent what the doctor intended him to represent. That is to say, because he did not understand

²¹ For simplicity’s sake, I will treat the text message as a single information content entity. Admittedly, a more complex artifact might involve multiple information content entities—e.g., a text message consisting of two statements, whose respective patterns of pixels concretize two different information content entities—but that is beside the point here.

the words to mean that he ought to take the medication in that quantity three times per day, he did not concretize the ICE at all.

Turning now to the last of the three elements I wish to highlight, I think we can see that most ICEs will also have a *force* that accompanies its point or purpose.²² As Searle has argued (1975, pp. 346-8), the force of a speech act is not the same as its point or purpose, although it might be thought of as a further specification of that purpose. For example, both a request and a command have the same general “illocutionary point” (i.e., getting somebody to do something), but they do not have the same “illocutionary force,” since one is asking that the action be done while the other is demanding that it be done. The force of content is likewise something which is conveyed when content is communicated. Again, we can differentiate between the request, “Please sit down,” and the demand, “Sit down, or else.” Both sentences overlap in content insofar as they are both about the addressee sitting down, but they differ in what they mean and what else is communicated (e.g., the latter might also suggest that the addressee will be in some way punished for not sitting down).

Like in the case of purpose, we can see that force is inherent to an ICE by considering cases where the intended force of an ICE is misinterpreted, and so the ICE fails to be concretized. If a motorist sees the traffic sign, “Speed Limit: 55mph,” and interprets it as a matter of advice or best practice, then he fails to understand the content it bears. The content of that traffic sign taken as a whole includes the fact that this speed limit is a requirement upon anyone controlling a vehicle on that road, that anyone who travels in excess of 55mph is liable to receive a penalty, and so on. To interpret the sign as a merely advisable upper speed limit would be to fail to have a cognitive representation

²² The notion of *force* can be traced back to Frege (1879/1952), but was developed more fully by Austin (1962). Specifically, Austin distinguishes between a locutionary act, its illocutionary force, and its perlocutionary effect. For example, uttering the sentence “Terry, throw the ball” is a locutionary act because it involves making sounds with a certain syntax and reference (e.g., “Terry” refers to Terry). But we can also describe what this utterance *does*, which in this case is to get Terry to do something, and which has the force of a command. And lastly, this utterance has an effect in the world, e.g., of motivating Terry to throw the ball, which is its perlocutionary effect.

that concretizes that ICE at all. Analogously, we could imagine the sign's manufacturer misinterpreting the intent behind the sign, and producing one which read, "55mph or lower is safest". Intuitively, this would not bear the same content as "Speed Limit: 55mph".

To summarize, then, in order to understand more complicated representations and content, and especially cases of world-to-mind and world-to-word direction of fit, we can decompose representations and content into various representational and non-representational components. In the case of mental phenomena like desire, we can differentiate between the cognitive representations of the entities relevant to the desire being satisfied and the various dispositions and bodily processes involved with that desire (thirst, tiredness, beverage preferences, feelings of longing, etc.). Similarly, in the case of prescriptive content (invitations, commands, traffic signs, etc.), we can differentiate between various representational components. A booklet of instructions for putting together a bookcase is about not only the bookcase or its pieces, but about actions one can take to arrange those types of pieces into a reliable, functional piece of furniture of a particular type. A shopping list is not only about types of products (*bread, milk*, etc.), but about the types of products the addressee should buy in order to achieve a goal (e.g., having enough to eat for the week, or satisfying the desires of the list's author). And a seemingly simple artifact like a sign reading "Speed Limit: 55mph" will bear content that is not only about vehicles or a specific range of speeds, but also about the fact that this is a requirement, backed by threat of punishment—which is to say, that those who are charged with monitoring vehicles' speed on that road are generally disposed to carry out that threat, should they observe you.

Before leaving this topic, I need to offer a few additional remarks and qualifications. Throughout this discussion, I have been assuming for simplicity's sake that these are examples of single ICEs, rather than, e.g., multiple ICEs or aggregates of them. But to clarify, this is not in any way essential to my view. The above examples might instead be aggregates of ICEs, or complex ICEs with

component ICEs as parts. That is likewise compatible with my view because all that is required is that, when taken as a whole, these examples include content about addressee, purpose, and force. When concretized, some of the “representational units” (Smith and Ceusters, 2015, p. 4) involved will represent portions of reality all by themselves (e.g., “Mary” in the request “Call me, Mary”). In other cases, however, the content concretized will go beyond observable patterns of ink. When the nurse records “72 beats per minute,” she does not in that concretization name the patient or the time of the measuring, yet that representation is nevertheless about a particular patient at a particular time. Likewise, the content borne by the traffic sign that reads, “Speed Limit: 55mph,” may be rather complex: it is not about speed limits in the abstract, but about the speed limit on a particular road as backed by a threat of penalty. Again, the traffic sign’s content has this specificity because of the circumstances of its creation, i.e., because it is a particular product of a particular system of legal frame principles.

All that is also not to say the concretizations of an ICE will be necessarily unambiguous across any of these dimensions of audience, purpose, or force. It might very well be the case that, for a given concretization of an ICE, it is not immediately clear whether it is a description of events or a prescription to act, or whom it is addressed to or intended for. But in a similar vein, it may not be obvious from its concretization what portion of reality an ICE is supposed to refer to. Out of context, the words “72 beats per minute” are not very meaningful, nor do they indicate unambiguously that this is a description of a pulse rate, nor whose pulse rate they would be, or at what time. The point is that the ICE created by the nurse is concretized in those patterns of ink or pixels regardless of whether a particular observation of those patterns is causally sufficient for a given person to understand or correctly interpret their meaning.

Lastly, although many of my examples have been about ICEs with a prescriptive intent (requests, commands, etc.), it is important to recognize that other types of ICE most often have

addressee, purpose, and force as well. If John texts Mary, “I am at the store,” then his addressee is Mary, his purpose is to describe his location to Mary, and the force of his message is a definite assertion. By contrast, if he texts her, “I think the bread is in the pantry,” then the force of his message is less a definite assertion than a kind of supposition or mere belief.

Roles

In addition to an ontology of representation and content (and especially of prescriptive content), my account will also invoke the ontological category of *role*. As we have seen already, Searle describes the ontology of the social world in terms of the imposition of *status functions* upon physical objects (e.g., a piece of paper) and events (e.g., an utterance). Such “X terms,” he says, come to “count as Y terms” in the context of a social institution. And as I mentioned earlier in this chapter, Searle has also conceded that it is possible to speak of “freestanding Y terms,” that is, of status functions which do not correspond to any particular physical objects or events, but which nonetheless bottom out in the deontic powers of agents. Like Searle, I believe this notion of status function will be important for spelling out the ontology of deontic phenomena. There is, of course, the general point that deontic phenomena in the social world depend in some way on the creation of status functions. But more specifically, there are the cases I described in an earlier chapter, where the powers, rights, and responsibilities of people are tied to the role they play within a social or institutional context.

Basic Formal Ontology (Arp and Smith 2008; Arp, et al., 2015) includes an account of roles along these lines. As Arp and Smith describe it, a role is:

a realizable entity [roughly, a property of an entity that can be either dormant or manifest] whose manifestation brings about some result or end that is not typical of its bearer in virtue of the latter’s physical structure. Rather, the role is played by an instance of the corresponding kind of continuant entity because this entity is in some special natural, social, or institutional set of circumstances... What is crucial for understanding a role... is that it is a realizable entity

that an independent continuant can take on, but that it is not a reflection of the in-built physical structure of that independent continuant. Certain strains of *Escherichia coli* bacteria have the role of pathogen when introduced into the gut of an animal, but they do not have this role when merely floating around in a pool of water. A heart has the function of pumping blood; but in certain circumstances that same heart can play the role of dinner for the lion. Roles are optional, and they often involve social ascription. This is why a person can play the role of being a lawyer or a surrogate to a patient, but it is not necessary for persons that they be lawyers or surrogates. (Arp and Smith 2008, p. 3)

This characterization of role is broad, encompassing not only roles within specific social or institutional circumstances (e.g., the role of quarterback), but also roles within particular physical circumstances (e.g., the role of pathogen). However, the former are the only that need concern us here. Here we can see a strong similarity with Searle's notion of a status function: such social or institutional roles are features of entities which go beyond their brute physical makeup, and which are instead socially ascribed.

Other authors, drawing on this notion of role, have suggested that the ontology of deontic phenomena can be accounted for in terms of "deontic roles" (Almeida, et al., 2012; Brochhausen, et al., 2013; Chu and Allen 2016). To have an obligation, they maintain, is to have the role of obligor in a social or institutional context, which is realized in the fulfillment of the obligation. Moreover, such roles are understood as concretizing certain quasi-abstract entities called "socio-legal generically dependent continuants," characterized much in the same vein as Smith's freestanding Y terms. But I do not think this approach is quite right. Given the description of roles above, it does make some sense to regard a promisor as taking on the role of obligor. However, this view also obviously relies on Reinach and Smith's characterization of claims and obligations as entities belonging to a *sui generis*, quasi-abstract ontological category. In other words, the only way this view can apparently account for the deontic character of these roles is to appeal to an irreducibly deontic, *sui generis* category of entity, which I am arguing is not the best ontological account we have on offer.

That being said, I do believe social or institutional roles of this sort are an important part of the full ontological account of deontic phenomena. To give some indication of why that would be, consider again Searle's observation that so-called "deontic powers" seem to be tied inextricably to status functions. For example, to have the role of President of the United States is to have certain powers, rights, and responsibilities in a political context, even if the specific powers, rights, and responsibilities associated with that role can evolve (or devolve) over time. To be a bit more precise, we can think of such deontic powers, rights, and responsibilities as specifications of how that social or institutional role can or should be realized. What it means to be the quarterback of a football team includes the specification of what he is entitled to do (e.g., call an audible), meant to do (e.g., pass the ball), or prohibited from doing (e.g., passing so as to intentionally ground the ball) within the context of the game.

From an ontological perspective, I believe this kind of specification is also a matter of content, and at least in part of prescriptive content. Like Searle's Declarations, the content specifying a role has (or at least most often has) both directions of fit at once: it both describes what the role's bearer is empowered (or has authority) to do and prescribes what he may, must, or must not do. For example, a specification of the role of President of the United States would involve describing his powers within that institution, i.e., the ways in which he can modify features of the social or institutional world, including in many cases modifying what other people are empowered, permitted, or required to do. But a full specification (emphasis on *full*) would also detail his responsibilities within that role, and also his deontic limitations, including both what falls outside the scope of his powers and what he is prohibited from doing.

As I will argue in the next chapter, many kinds of deontic phenomena are a matter of having a social or institutional role whose range of realizations has been specified deontically, i.e., in terms of what the person in that role is entitled or empowered to do, or responsible for doing. In some cases,

these will be ways of acting that are allowed or required within the context of that role. In other cases, these will be new ways the person is empowered or authorized to act because of his role (e.g., a superior's authority to issue binding orders to a subordinate).

The Ontological Status of Deontic Phenomena

Adolf Reinach (1913/1983) and Barry Smith (2003a; 2003b; 2008), among others, each highlight the question of the ontological status of social-deontic phenomena like claims and obligations, and propose (albeit in slightly different ways) that this question should be answered by admitting a new category of entity into our ontology: a “quasi-abstract entity” or “freestanding Y term”. I have, in turn, raised the question of whether it is possible, from an ontological point of view, to make sense of deontic reality without introducing such quasi-abstract entities, while simultaneously doing justice to the deontic phenomena, and I have suggested that it would be preferable not to appeal to a *sui generis* category of entity if there are better ontological resources at our disposal.

In this final chapter, I attempt to draw together the various pieces of my account from the foregoing chapters into a coherent whole: within an anchored system of frame principles, people gain certain specified powers, rights, and responsibilities, including the ways in which they can be exercised in the creation, modification, transference, or elimination of regulations over agents’ actions (requirements, obligations, permissions, prohibitions, entitlements) or of their deontic powers. In many cases, these are tied to having a specific role. In all these cases, the scope of these powers, responsibilities, and so on, are ultimately cashed out in terms of shared representations and dispositions, and especially as captured and memorialized within prescriptive content, which allows for things like codification (e.g., within a law or rulebook). Such an account, I believe, gives us the

basic ontology of deontic phenomena within the social world without reifying phenomena like obligations, or placing them in a *sui generis* category.

First, I will revisit Searle's account of deontic powers and explain elements of that account can be brought into my own. Then I turn to explaining how prescriptive acts, prescriptive content, and roles play a central role in the ontology of deontic phenomena. I will also explain how commonly accepted relationships between deontic elements (e.g., the correlation of claim and obligation) fit within this ontological account. I close by considering how this account would handle examples of supposed freestanding Y terms.

Representation, Dispositions, and Frame Principles

Despite the fact that I criticized Searle's (1995; 2010) account of deontic powers in Chapter 2, I believe that he comes close to pinpointing where deontic phenomena like obligations, permissions, and authority fit within an ontology. In his view, deontic phenomena are to be explained in terms of the powers agents possess to compel other agents to do something (cf. Searle, 2010, p. 151). As we have seen already, his account of deontic powers, and especially of so-called "negative deontic powers," runs into problems, especially in its characterization of obligation as a type of negative power possessed by the obligor. Nevertheless, the central place he affords to agential power, and to guiding, regulating, or constraining action, is on the right track. Ultimately, phenomena like permissions and obligations exist in social-ontological contexts as a way of regulating behavior: binding agents to some courses of action and loosening the constraints around others. I believe, then, that Searle's main point, that deontic phenomena is in many respects a matter of the exercise of power within a given society, can be rehabilitated within a fuller ontological account.

The way to do this is to understand the role of frame principles, as backed by threat of sanction (either by the overall anchor group for that system of frame principles, or by some delegated sub-

group of monitors), in enabling or constraining behavior. All this I attempted to spell out in Chapters 3, 4, and 5. One question I did not address adequately, however, was the ontological status of frame principles. What is a frame principle? Is it a type of entity? And if not, what exactly am I referring to on the side of reality when I say there is an “anchored frame principle”? One half of the answer to this question has been suggested already: part of what it means to say there is an anchored frame principle is for there to be a collective *disposition* to monitor for compliance and to sanction in case of non-compliance. Thus, the possible scenarios covered by the frame principle correspond to the conditions under which people are disposed to monitor and sanction. For example, consider a frame principle along the lines of:

(CF) For all x , if x is counterfeit money, then it is illegal to try to use x to make a purchase.

The possible scenarios covered under this frame principle (i.e., the *frame* in Epstein’s parlance) are those in which a person attempts to use counterfeit money. But if this frame principle is anchored, then this is also the same set of possible scenarios in which people would be disposed to monitor for compliance (i.e., if there is a reasonable suspicion, then checking to see if the bill is a fake) and to sanction in case of non-compliance (i.e., if the bill is believed or shown to be fake, then having a negative response, such as refusing the bills or phoning the police). In this respect, using the language of “frame principles” is shorthand for describing how the anchor group is collectively disposed. (Note, again, that this leaves open the possibility that that particular frame principle is anchored by the society at large, or else by a much smaller number of delegated monitors, e.g., if the law is on the books, but not widely known.)

The other half of what it means ontologically for there to be an anchored frame principle is what I am now in a position to explain more clearly: it is for the anchor group to collectively *represent*

the content of that frame principle as holding true. In Chapter 5, I gestured toward this when I suggested that frame principles were anchored in part by collective belief (in a dispositional, not occurrent sense). But given the account of content and representation I outlined in Chapter 6, we can put a bit of a finer point on this claim. Consider again the frame principle (**CF**), and for simplicity's sake suppose that this is collectively accepted directly by the society at large. To be collectively accepted, they must regard the usage of counterfeit money as illegal, i.e., as an act which renders one liable to legal sanctioning, in the sense that law enforcement personnel, among other people, would be disposed to respond negatively. Again, to echo Searle (2010, p. 8), that is not to say that everyone necessarily *approves* of the law, but merely that it is *recognized* as a legal reality.

What this means, ontologically, is that these people share an information content entity (ICE), which, in addition to being concretized in a code of law, is concretized in their respective cognitive representations. This ICE is, in part, about a type of action (purchasing) involving a type of artifact (counterfeit currency); it is addressed to anyone who engages in commerce within that society; it has the purpose of dissuading the usage of that type of artifact in that type of action, and suggests that failing to comply with this directive will result in some form of punishment, as carried out by those duly invested with the authority to make legal judgments and enforce laws (all of which is to say, it has the force of a legally binding requirement). Thus, to say that frame principle (**CF**) is anchored is to say that there is some content shared among multiple concretized representations. But note, again, that this collective representation is also related to a disposition: the anchor group is not always engaging in mental processes that concretize that content and represent it as holding true, but they are disposed to represent it as true.¹

¹ Although belief and disposition to behave often go hand in hand, they can also come apart. If I believe that my cup of coffee is poisoned, then (assuming I do not want to die) I will act in a way consonant with my belief. Yet it is also obviously possible to act contrary to one's beliefs; John's belief that regular exercise will do him good may not translate into actually engaging in regular exercise.

Now the fact that some belief is widely held does not mean that the content of the belief is a frame principle. For example, the ubiquitous belief that Washington was the first president of the USA is not the collective acceptance of a frame principle. This is, rather, a belief that is the consequence of the collective acceptance of a frame principle outlining conditions under which one becomes president of the USA, and then a belief that Washington himself satisfied these conditions.

What sets a frame principle apart from a widely held belief like this is twofold. First, frame principles are usually general in their content because they are meant to apply across many different possible scenarios. Thus, we have legal frame principles governing how all people in a society, or in some other designated group, are allowed or required to behave, rather than a legal frame principle directed at each individual. That is not to say a frame principle could not mention an individual (e.g., a frame principle prohibiting depictions of the person Muhammed), but that it is usually not feasible to limit the frame principle exclusively to individuals.² Second, a frame principle must have both word-to-world (or mind-to-world) and world-to-word (or world-to-mind) directions of fit at once. Recall that, in Searle's (2010, pp. 12-15) view, social reality is constructed out of acts that have "the same logical form as Declarations" (p. 13), which is to say, by acts which have both directions of fit at once. For example, the Declaration "You're fired" is satisfied only if the world is changed in accordance with the words and if the words correspond to the new reality. These need not be explicit Declarative speech acts (e.g., "You're fired"), but they are at least cases where we "linguistically treat or describe, or refer to, or talk about, or even think about an object in a way that creates a reality by representing the reality as created" (p. 13). This is, in short, how the acceptance of frame principles works too: by accepting the frame principle, we make it true.

² This is, however, a logical possibility. We could imagine a fictitious scenario where everyone in a society gangs up on a particular individual and collectively accepts principles targeting him exclusively.

Because of this double direction of fit, the content of a frame principle will never be purely descriptive, but also partly prescriptive in the sense explained in the previous chapter. An anchor group will be disposed to regard a frame principle as holding true (e.g., that a certain vehicle speed is the speed limit), as well as to monitor behavior that accords with the anchored frame principle (e.g., to punish those who go above the speed limit). But the content of a frame principle that one holds true has a prescriptive dimension: a portion of reality is meant to change in accordance with the content, “Speed Limit: 55mph”. As Searle would put it, such content has “the same logical form as Declarations.” And this makes sense. In the case of the content of the speed limit sign, the desired outcome underlying the creation of this content is the bodily safety of all motorists on the road. Like in other cases of prescriptive content, the desire to achieve some goal is translated into prescriptive content meant to guide the actions of its addressees.

Power, Authority, and Roles

The first step toward the creation of particular deontic phenomena like obligations and permissions is the collective acceptance of frame principles through shared dispositions. As I have explained, these dispositions include regarding the content of frame principles as true and being sensitive toward compliance with them. But as I have also argued already, a typical society will have a *system of frame principles*, i.e., a set of frame principles which includes certain nested frame principles. Thus, even if the people in a society need not directly accept every frame principle operative within their society, they must at least collectively accept frame principles that amount to the procedures and mechanisms for creating and anchoring new frame principles (e.g., collectively accepting the voting process for creating officials who can create new laws).

One crucial component in the construction of a system of frame principles is the investment of people with the *authority* to further modify social reality. In particular, we can observe that certain

agents are invested with authority to modify what others are permitted or obliged to do, and that such exercises of power are often the source of agents' particular permissions and obligations: commanding officers issuing orders to subordinates, managers delegating tasks to employees, teachers assigning homework to students, parents laying down rules for their children, governing bodies establishing the speed limit on a particular road, and so on. In each of these cases, the power to modify permissions and obligations is bound up with the specific role the person (or group of people) has. That power itself is not itself a role, yet we can identify in each of these roles a particular *scope of its realization*, which includes the specific ways that role can be realized in the creation, modification, or elimination of permissions and obligations. Every role is by definition a realizable entity, which is grounded externally in the attribution of that role by a social group (Arp, et al., 2015, pp. 98-101), i.e., through the familiar Searlean mechanism of collective acceptance. But, as a realizable entity, a role is the role that it is only in virtue of being realizable in certain types of processes. Accordingly, the representation of a role in general requires specifying how exactly that role can be realized. For example, the role of physician is the role that it is because it is realized in certain characteristic process types (diagnosing disease, consulting with patients, etc.)

The point here is that, in many cases, a specification of the ways in which a role can be realized includes a specification of how that role can be realized in actions that modify others' permissions or obligations. For instance, the ability to issue binding commands to subordinates in certain contexts is part of the full specification of the possible realizations of the commanding officer's role. Specifying the possible realizations of a role is a matter of representation, whether codified and documented in some official manner, or merely in the heads (so to speak) of the people in that social context. For example, the scope of a parent's role—including its responsibilities, permissions, and authority—may be specified by an ICE shared by her society or social group, rather than as codified in a law. In any

case, the content or representation in question is at least partly prescriptive in nature: it is about what the role's bearer can do, but also about how others are meant to respond to it.

Conversely, it might be the case that an individual attempts to overstep the boundaries of his role and impose a requirement or remove a permission, although he lacks the power to modify them. For example, if John's freedom to practice religion includes immunity from working on the sabbath, then Jane might try to issue John a directive requiring him to work on the sabbath, but because of John's immunity her attempt is unsuccessful.

We can also note the ways in which agents can exercise power to modify their own permissions and obligations. One familiar example is promising itself: if John promises to mow Mary's lawn, then he has exercised his power over his own future actions to generate an obligation. One reason this is a genuine exercise of deontic power is that he cannot eliminate the obligation by simply recanting it; he cannot simply choose not to mow without finding himself liable to sanction. A promise is not merely an expression of intention. John has, in a sense, issued a binding directive to himself.³ In promising, he performs an act that creates an information content entity that prescribes that John mow Mary's lawn, and that he is liable to sanction for failing to do so; this ICE is shared by John and Mary, and potentially by anyone else who learns of John's promise.

³ As Searle (1975) argues Commissive have the same direction of fit as Directives. They differ from Directives in that they are not oriented toward guiding the actions of others, but in binding oneself to a course of action. When it comes to my account of deontic phenomena, Commissives break down into essentially two components. First, there is the binding character of a Commissive (e.g., a promise), which I have already explained in terms of an anchored system of frame principles. Second, there is the fact that Commissives similarly guide action, albeit of the speech act's speaker rather than its addressee. But in this latter regard, the content of a Commissive is not radically different in kind from that of a Directive. Both a promise and a command prescribe an action, and both have (or can have) binding deontic force.

Similarly, in the case of Declarations, we have already seen much of their role in an ontological account of deontic phenomena already, especially their general role as enabling the further construction of social institutions. But at this point, I wish to emphasize that Declarations have a dual direction of fit. On the one hand, like an Assertion, they have a world-to-world direction of fit. If a senior manager says, "You are hired," then he is asserting a new institutional fact. On the other hand, like Directives and Commissives, a Declaration also has a world-to-word direction of fit. If that same manager tells another employee, "You're fired," but that (ex-)employee simply ignores the Declaration and continued to behave as though he were not fired, then his actions would fail to match the hiring manger's words; he would be criticized, cajoled, and eventually thrown out by security. (This example is assuming, of course, that the manager in question is exercising his legitimate power to hire and fire.) Here, too, we see how the same dynamic, where content with a world-to-word direction of fit is backed by deontic force, fits into an ontological account of deontic phenomena.

In a similar vein, suppose that Mary writes this up in a contract for John to sign. Before signing, the content of the document is about John's mowing of Mary's lawn. Since John has the power to bind himself to future actions, and since signing one's name at a certain place on a document can count as binding oneself to the terms laid out in that contract, when John signs the document he places himself under an obligation. Before he signs, the document is about John, mowing, and Mary's lawn; it is prescriptive in nature. But after he signs, it gains the force of a requirement, since now part of its content includes John's having made himself liable to sanction for failing to mow her lawn.

Note, however, that in this case John's power to bind himself to a future action may not be tied to any discernible social or institutional role. Even in this example, though, we should recognize that his ability to bind himself to future actions still seems to hinge in some way on collective acceptance, e.g., the acceptance of a principle such that mentally competent adults (unlike children or mentally incompetent people) are capable of binding themselves to perform future actions. In yet other cases, permissions and obligations are rooted in the joint exercise of multiple powers at once. If John signs an employment contract, then he gains certain permissions (e.g., to access his workplace) and obligations (e.g., to perform the duties of his role) in virtue of both his own power to place himself under obligations and in virtue of his employer's power to authorize him and to impose new requirements upon him as an employee.

Next, we can note the ways in which agents can modify others' powers or authority. The clearest cases are, perhaps, when a role itself is gained or lost, as when one manager hires or fires another, lower-level manager. But another type of scenario is when one exercises authority to invest other people with new powers or authority, regardless of whether their role has fundamentally changed. For example, if an amendment to the US Constitution were passed stripping the president of his power to command military forces. Whether or not this involves the elimination of a distinct role (e.g., if there exists a separate role of Commander in Chief), it would be a case in which power is

exercised in the modification of others' powers. Lastly, we can also note cases where agents exercise their power to modify their own powers. Thus, if a pope willingly resigns his office, then he exercises the power of his office in order to divest himself of all papal powers.

To summarize, an important building block in a system of frame principles that gives rise to particular permissions and obligations is the exercise of authority. The nature and scope of this authority may be collectively accepted in a direct way by the group anchoring the system of frame principles as a whole, or it may be accepted indirectly through more proximately accepted procedures designed to invest people with authority. Once that authority exists, however, it becomes possible to exercise that authority by modifying permissions and obligations (others' or one's own), or by modifying powers themselves (others' or one's own). Ontologically speaking, the scope of a given power is determined by the possible scenarios represented in a corresponding frame principle, i.e., the legitimate exercise of the powers corresponds to the collectively accepted specification of the ways in which those powers can be exercised. These are often, but not necessarily, tied to particular roles, such that in gaining or losing a role, one can automatically gain or lose a corresponding set of powers or authority.

Action Regulation

Several examples of deontic phenomena in the last section were concerned not only with what agents are empowered or enabled to do (hire, fire, make promises, issue binding commands, and so on), but also about what agents were *required* or *permitted* to do. The reason is that power and authority, gained through collectively accepted frame principles, is ultimately directed toward the goal of regulating the actions of people in a social context: prohibiting certain actions, requiring others, permitting yet others. In gaining power or authority, one becomes enabled to create prescriptive content which is *binding* upon its addressee, backed by threat of sanction. For example, the sign "Employees must wash their

hands before returning to work” communicates content that prescribes an action to the employees, and the prescribed action, which the sign presents as having the force of a requirement, is in fact a requirement upon the employees if it is backed by the competent authority. But what makes it a requirement? After all, if a random customer were to hang a sign in the restroom, “Employees must dance the tango upon exiting the restroom,” the employees would certainly not be required to dance, despite the fact that the sign presents the action as an imperative. The answer is that prescriptive content takes on binding force when it is backed by a collectively accepted frame principle (or more likely, a system of frame principles), where that collective acceptance is cashed out in terms of a disposition to sanction those who fail to comply. For example, what gives binding force to Jane’s directives over the employees she supervises is that the directive is issued within the context of a collectively accepted (i.e., anchored) system of frame principles. If that system were to fall apart, then her directives would instantly lose their binding force.

Similar considerations hold for other cases of action regulation, especially permission and prohibition. For an action to be prohibited is simply for the commission of that action to be subject to sanction. This prohibition can be communicated in content explicitly (e.g., “No speeding in work zone” explicitly prohibits an action) or implicitly (e.g., the traffic sign “Stop” explicitly calls out the obligation to stop, and implicitly prohibits non-stopping).⁴ Likewise, for an action to be permitted is for it to be presented as among one’s valid (non-sanctionable) options, whether explicitly (e.g., a green traffic light) or implicitly (e.g., if the action is not among the total set of prohibitions).

Keep in mind that such directives can vary greatly in the specificity of their addressee. Some directives prescribe a particular agent to perform an action; others are more generic, prescribing that

⁴ As the examples suggest, by ‘explicit’ I mean that we can identify an ICE with prescriptive content that presents the action as prohibited, i.e., if one were to commit the action, then that person would be liable to sanction. By ‘implicit,’ I mean that the prohibition is logically implied by another prescriptive ICE. Thus, “Stop” presents an action as a requirement, but we can infer a prohibition on non-stopping from it.

any agent in such-and-such circumstances must perform an action. Again, the content of the utterance, “Go away, John,” is directed to a particular addressee, since it specifies John as the individual who is to perform the action, whereas the content of a stop sign is not directed at any specified addressee, but rather to anyone who approaches the sign in a vehicle. Consequently, the group which concretizes these different kinds of directive content in their dispositions will vary as well. For example, the content of the stop sign is likely concretized in the dispositions of law enforcement officers to monitor whether motorists stop at the intersection. Note, however, that the deontic force of this piece of prescriptive content would stem not only from the law enforcement’s disposition to monitor and sanction, but also from the collectively accepted system of frame principles which make those law enforcement roles possible.

Note again, too, that there is no one way to ground a deontic phenomenon like an obligation. Any and all grounding conditions will be set up by anchored frame principles, which leaves open the possibility that there are numerous ways for an agent to become obliged. Another possibility is that the action prescribed by the imperative content is obligatory if it was established as such by a frame principle, e.g., the frame principle establishing that if an action is promised, then it is obligatory. This frame principle, in turn, might be anchored directly by collective acceptance, or else by a Declaration, the conditions for which are established by another frame principle. For example, if the frame principle that all employees must arrive to work on time is directly anchored by collective acceptance, then the group monitoring for tardiness will be disposed to sanction those who arrive late. Or, if the frame principle is anchored by the fact that the boss declared that employees must not be late, then that fact of the Declaration would need to have been established by a further frame principle, e.g., one which establishes that if somebody is a boss, and if that person declares in her capacity as boss a new work-related rule, then that rule is anchored. All this was laid out in Chapter 5.

Thus, my proposal is that the question of what agents are required or permitted to do just is the question of which directives prescribing those actions, as backed by a collectively accepted system of frame principles, exist. When an agent is the subject of such a directive, his action becomes regulated: ways of acting become opened or closed for him. There is a directive, issued with binding force, that an agent *must* perform an action. The agent is obligated. There is another directive, issued with binding force, that an agent *may* perform an action. This directive has binding force because violation of the permitted course of action is likewise met with sanction. This just is the agent's permission. There are no extra entities required. There are only the entities having to do with regulating that agent's actions, and there is the deontic force supplied by collective acceptance.

This account also explains why it is true in an important sense that many obligations, etc., do “bottom out” in deontic powers, as Searle suggests. Many directives come into existence because of the exercise of deontic powers, but those powers are founded in turn upon a collectively accepted system of frame principles. I should note again, however, that this is only one paradigmatic case of what grounds an obligation. For example, an obligation to pay for damages might be grounded instead upon a frame principle dictating that whoever causes damage to another's property is obligated to pay for those damages. In such a scenario, the obligation is not grounded in the exercise of a deontic power (e.g., issuing a command), but simply in the fact that the obligor did something which caused damage.

Deontic Relations?

Before drawing this discussion to a close, I will offer a few remarks on deontic *relations*. First, in Reinach's account of the ontology of deontic phenomena, he highlights the relationship between claim and obligation: if John has an obligation to Mary, then Mary has a claim upon him. This is one indicator, Reinach suggests (1913/1983, pp. 12-3), that we have *a priori* knowledge of these “peculiar

entities.” When the obligation is “over and against another” (p. 12), then we understand there to be a corresponding claim as a matter of necessity, much in the same way one might understand a mathematical relationship *a priori*. Barry Smith takes this line of thought a bit further and suggests that such claims and obligations exhibit a “mutual dependence” upon one another, such that the one entity cannot exist without the other (cf. Smith 2003b, p. 289).

If claims and obligations are discrete, quasi-abstract, or freestanding entities, then I agree it makes sense to think of them as standing in such a relation. But given that my aim is to provide an alternative account, which omits such quasi-abstract entities from the ontological picture, one might wonder how my account accommodates this apparent relation between claims and obligations.

The approach I have developed throughout Chapters 6 and 7 can indeed accommodate this relationship between claims and obligations, albeit not by positing a relation of mutual dependence between a claim-entity and an obligation-entity (since my view posits no such entities). Assume for the moment that all obligations have corresponding claims, i.e., that for any situation in which an agent is obligated or required to perform an action, there is some person or group of persons that has a claim against him that he does.⁵ Then, on the approach I have developed here, what we have is a directive prescribing that the obligor perform an action; this action is presented as a requirement upon him; it is also presented as being backed by threat of sanction. Moreover, part of the content of that requirement is that the claimant has the right to seek restitution for a broken promise.

⁵ As Reinach points out (1913/1983, p. 12), this need not be the case. Reinach differentiates between “relative obligations,” which are obligations to another person (e.g., John’s obligation *to Mary* to mow her lawn), and “absolute obligations,” which present an action as required, but which is not tied to any particular person’s claim upon the obligor. Although Reinach admits that the latter type of obligation is up for debate, he also claims “it is beyond any doubt that they are *a priori* possible” (pp. 12-3). That is to say, even if all obligations involve some degree of wider accountability, monitoring, and liability to sanction, it seems possible to specify someone’s obligations without identify any particular claimant or group of claimants especially invested in that obligation. One possible stark example of what he calls “absolute obligation”: John promises to bury his elderly father beside his mother. Now suppose that there is no existence after death, so that when John’s father dies, he exists no more. Then, if John is indeed obligated to bury his father beside his mother as he promised, then he will have an obligation to perform an action, but nobody *to whom* he is obligated, since *ex hypothesi* his father exists no more.

So consider again John's promise to mow Mary's lawn. The relevant accepted frame principle governing promising might include the content that: "whomever the promise was made to has the right to seek restitution for the broken promise." On the level of the particular directive addressed to John (that he must mow her lawn), it is understood that, if John fails to mow Mary's lawn, then Mary has the right to seek restitution from him (an apology, a compensatory action, etc.). This right to seek restitution would be cashed out, in turn, in terms of Mary's deontic powers, e.g., to impose an additional obligation on John to make it up to her. Thus, John's obligation corresponds to Mary's claim in the sense that, when the requirement upon John is fully specified, that specification includes some notion that Mary is especially entitled to him performing the action, or to seek restitution should he fail to do so.

The second set of deontic relations I wish to consider were presented back in Chapter 2. As I mentioned there, Hohfeld draws a distinction between first-order deontic elements (permissions, claims) and second-order deontic elements (powers, immunities). As a reminder, what differentiates these two sets of deontic phenomena is that first-order elements exist to regulate actions and that the second-order elements exist to modify the first-order elements. Second-order Hohfeldian deontic elements are not directly responsible for regulating behavior in the same way that their first-order counterparts are. Having a power or immunity might even imply the existence of such a regulation, but they are clearly not identical. For example, it might be the case that if A is immune from φ ing, then B is prohibited from trying to coerce A into φ ing. But the purpose of a power or immunity is not to regulate an agent's behavior, but rather to empower agents to generate, modify, transfer, or eliminate regulations or to be protected against other agents' power. My own account suggests much the same. As we can see from the above discussion, my account of the ontology of deontic phenomena maps fairly well onto Hohfeld's two orders of deontic elements: frame principles give rise to action regulations and powers, the latter of which can further modify action regulations and powers.

Freestanding Y Cases Reconsidered

My approach, then, suggests that we do not need to hold that obligations or permissions are entities that belong to a special category of quasi-abstract entities or freestanding Y terms. There is an ontologically more conservative account which finds a place for the same deontic phenomena within a general schema of ontological categories. To this point, I will now explain how this alternative approach handles two cases of allegedly freestanding Y terms as presented by Barry Smith (2003a; 2003b; 2008), namely, cases of debts and property rights. My goal is not to argue that the freestanding Y approach in general is false. For example, I will not address Smith's comparatively non-deontic cases of freestanding Y terms, namely, electronic money and blindfold chess. Rather, I am merely arguing that positing the existence of a Reinachian, *sui generis* category of entity is not necessary to account for these deontic phenomena.

First consider Smith's (2003b, pp. 288-9) case of debts. Smith rightly points out that a record of a debt—a piece of paper, or a file on a hard drive—is not the debt itself. Smith draws the conclusion that such debts are freestanding entities. But given the account of entities like obligation developed in this chapter, I am in a position to explain what that debt is without recourse to freestanding Y terms: the debt can be explained in terms of an information content entity that prescribes an act of repayment. This directive entity can be concretized in the qualities of that document, but might be concretized only “in the heads” of the debtor and creditor. There is no need to posit a separate entity, “the debt itself,” so long as there is a directive to pay back the amount of the debt (plus the specified interest), and so long as that directive is backed in the appropriate way by threat of sanction (e.g., the agreement was a legally binding contract). (And I am presuming, of course, that the emergence of this directive met the grounding conditions established by a relevant anchored frame principle.)

Next, recall Smith's (2003b, pp. 289-90) case of the Lucca family and its valuable Dutch painting. In this example, a family owns and rightfully possesses the painting for a century prior to any official documentation certifying that ownership. Smith rightly points out that this certificate is not an entity which counts as the property right, but merely an official record of it. Since the right over this property is not identical to the document which makes record of the right, Smith draws the conclusion that this right is a freestanding Y term.

But there is another way to explain this case without recourse to freestanding Y terms at all. As it turns out, this property right is quite a complex bundle of deontic phenomena. At the very least, it is a permission to possess, modify, or destroy the painting, a prohibition against unauthorized persons taking or modifying it, and the power to loan it out. Thus, representing the fact that Lucca family owns the painting will require more than identifying some one entity referred to by 'the property right'. For simplicity's sake, however, I will focus on that part of the property right that involves the *permission* to possess the painting. On my account, the permission to use the painting would likewise be explained in terms of an information content entity with prescriptive content. When they first discover the painting, they believe that they own it, though this is not yet certified. Their representation of their own permission to use it is at least *intended* to represent reality. What the later investigation into the history of the painting establishes is that the facts surrounding the Lucca family's possession of the painting meet the relevant grounding conditions, and therefore their representation is accurate. The official documentation, then, would simply be a way of testifying to the fact that their believed permission to use it is in fact grounded in the appropriate way and carries binding force. If, on the contrary, the investigation had established that Grandpa Lucca had stolen the painting 100 years ago, then it would have turned out that their belief about the situation was false, and carried no deontic weight.

This response to the Lucca case suggests that if there had been no record or evidence of rightful ownership, then there would not have been any facts about who owned the painting. And I think that is entirely correct. Suppose, for example, that the investigation into the Lucca painting found no record of an accredited owner that lived a hundred years prior. Suppose that the investigation turned out no evidence whatsoever. The painting has simply been found, and nobody has any idea about how the Lucca family came into its possession. In that case, we can easily imagine the relevant frame principle which would give the present Lucca family ownership, roughly: if a person has possession of an item, and there are no competing claims, then that person has rightful ownership of that item. Prior to discovering it, I do not think it makes any sense to suggest that there was a permission to use it. The permission only comes into the picture when they discover the painting. The fact that there is no competing claim to the painting means that the grounding conditions are met, and thus that the directive takes on binding force. Note that I am not claiming that their knowledge of the lack of any competing claims bestows deontic force, but that the actual lack of competing claims does. That the investigation suggests the lack of competing claims merely gives the family good reason to believe the fact which does, in fact, ground their rightful ownership.

Again, the argument of this section is not meant to discount the possibility that obligations, claims, rights, and so on, are quasi-abstract entities or freestanding Y terms. Rather, it is merely meant to provide an alternative way to address Smith's main examples of allegedly freestanding Y terms, but without recourse to freestanding Y terms. According to this alternative, the way to spell out the ontology of our rights, obligations, etc., is to look to what action regulations, power specifications, and frame principles exist.

Conclusion

At the outset of this dissertation, I noted two ontological puzzles that arise when we consider ordinary cases of deontic phenomena, namely, (1) in virtue of what do people gain obligations, permissions, powers, etc., within social contexts, and (2) what is the ontology underlying these phenomena? I have aimed to develop an account that sheds light on both questions by drawing on existing frameworks within social ontology: deontic phenomena arise within systems of frame principles, anchored by collective dispositions, and they can be accounted for ontologically by looking to collective dispositions and representational content. Such phenomena can be grounded in a variety of ways, which depends in turn on which frame principles are anchored within a social context.

This account has a few benefits, three of which I will sum up here. First, this approach to accounting for the grounding and ontology of deontic phenomena is more ontologically conservative than an alternative approach that would characterize deontic entities as Reinachian quasi-abstract or freestanding entities (Smith, 2003a; 2003b; Almeida, et al., 2012; Brochhausen, et al., 2013). Second, this account provides an explanation of the connection between group attitudes, social context, speech acts, and the emergence of obligations, permissions, etc., through frame principles, and thus provides a clearer picture of what it means to say that a social obligation, liberty, norm, etc., is operative relative to a particular social context. Third, this account provides a way to accommodate the anchoring of diverse perspectives on social obligations, liberties, norms, etc., insofar as it is possible to anchor multiple competing frame principle within a subgroup within a society.

This dissertation falls short of providing a formal-ontological account of deontic phenomena, e.g., a formal representation of all the relevant deontic types and relations. But my hope is that, in aiming to do justice to certain general features of the deontic dimension of the social world, this account can eventually prove useful in formally representing them. From a practical point of view,

this would be the next logical step, though I do not pretend that this account has supplied everything necessary to develop such a formal-ontological representation.

Bibliography

- Almeida, M.B., Slaughter, L., and Brochhausen, M. (2012). Toward an Ontology of Document Acts: Introducing a Document Act Template for Healthcare. In P. Herrero, H. Panetto, R. Meersman, T. Dillon (eds.), *On the Move to Meaningful Internet Systems: OTM 2012 Workshops. Lecture Notes in Computer Science*, 7567. Berlin: Springer.
- Anscombe, G.E.M. (1958). On Brute Facts. *Analysis*, 18(3), pp. 69-72.
- Anscombe, G.E.M. (1963). *Intention*. Cambridge, MA: Harvard University Press.
- Aristotle. (1984). *Complete Works of Aristotle*, 2 vols. J. Barnes (ed.). Princeton: Princeton University Press.
- Armstrong, D.M. (1978). *A Theory of Universals: Universals & Scientific Realism Volume II*. Cambridge: Cambridge University Press.
- Armstrong, D.M. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Arp, R., Smith, B., and Spear, A.D. (2015). *Building Ontologies with Basic Formal Ontology*. Cambridge, MA: MIT Press.
- Audi, P. (2012). A Clarification and Defense of the Notion of Grounding. In (Correia & Schneider, 2012, pp. 101-121).
- Austin, J. (1832/1995). *The Province of Jurisprudence Determined*. W. Rumble (ed.). Cambridge: Cambridge University Press.
- Austin, J.L. (1953). How to Talk: Some Simple Ways. *Proceedings of the Aristotelian Society*, 53, pp. 227–246.
- Austin, J.L. (1962). *How to Do Things with Words*. J.O. Urmson (ed.). Oxford: Clarendon Press.
- Beebe, H., Effingham, N., and Goff, P. (2011). *Metaphysics: The Key Concepts*. New York: Routledge.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, UK: Cambridge University Press.

- Bicchieri, C. and Muldoon, R. (2014). Social Norms. In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), URL = [<https://plato.stanford.edu/archives/spr2014/entries/social-norms/>](https://plato.stanford.edu/archives/spr2014/entries/social-norms/).
- Bickhard, M.H. (2003). Some Notes on Internal and External Relations and Representation. *Consciousness and Emotion* 4(1), pp. 101-110.
- Binmore, K. (2008). Do Conventions Need to be Common Knowledge? *Topoi*, 27, pp. 17-27.
- Bird, A. (1998). Dispositions and Antidotes. *The Philosophical Quarterly*, 48, pp. 227-34.
- Bispinck-Funke, C. (2017). On the Question of How Social Rules and Social Norms Exist. *Phenomenology and Mind*, 13, pp. 110-18.
- Bittner, T., Donnelly, M., and Smith, B. (2004). Endurants and Perdurants in Directly Depicting Ontologies. *AI Communications*, 13(4), pp. 247-258.
- Brentano, F. (1874/1995). *Psychology from an Empirical Standpoint*. London: Routledge.
- Brochhausen, M., Almeida, M.B., and Slaughter, L. (2013). Towards a Formal Representation of Document Acts and the Resulting Legal Entities. *Johanssonian Investigations: Essays in Honour of Ingvar Johansson on His Seventieth Birthday*. Walter de Gruyter.
- Brock, J.E. (1981). An Introduction to Peirce's Theory of Speech Acts. *Transactions of the Charles S. Peirce Society*, 17(4), pp. 319-326.
- Breuker, J.A., Valente, A., and Winkels, R.G.F. (1997). Legal Ontologies: A Functional View. In P. R.S. Visser and R.G.F. Winkels (eds.), *Proceedings of the First International Workshop on Legal Ontologies (LEGONT '97)*, pp. 23-36. Melbourne, Australia, 4 July 1997.
- Carnap, R. (1937/2002). *The Logical Syntax of Language*. Amethe Smeaton (trans.). Chicago: Open Court Press.
- Chalmers, D., Manley, D., and Wasserman, R. (eds.). (2009). *Metametaphysics*. Oxford: Oxford University Press.
- Chu, Y. and Allen, R.B. (2016). Formal Representation of Socio-Legal Roles and Functions for the Description of History. *Proceedings of the 20th International Conference on Theory and Practice of Digital Libraries, TPD L 2016*, pp. 379-385. Hannover, Germany, September 5–9, 2016.
- Correia, F. (2008). Ontological Dependence. *Philosophy Compass*, 3(5), pp. 1013-1032.
- Correia, F. and Schneider, B. (eds.). (2012). *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.

- Crosby, J.F. (1983). Reinach's Discover of the Social Acts. *Aletheia*, 3, pp. 143-194.
- Daly, C. (2012). Skepticism about Grounding. In (Correia & Schneider, 2012, pp. 81–100).
- Donohue, B. (2018). Toward a BFO-Based Deontic Ontology. Forthcoming in *Proceedings of the 8th International Conference on Biomedical Ontology*. Newcastle-upon-Tyne, United Kingdom.
- Dunn, J.M. (1990). Relevant Predication 2: Intrinsic Properties and Internal Relations. *Philosophical Studies*, 60, pp. 177-206.
- Durkheim, E. (1893/1984). *The Division of Labor in Society*. New York: Free Press.
- Durkheim, E. (1895/1964). *Rules of Sociological Methods*. New York: Free Press.
- Epstein, B. (2015). *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford: Oxford University Press.
- Fine, K. (1994). Ontological Dependence. *Proceedings of the Aristotelian Society*, 95, pp. 269-290.
- Fine, K. (2001). The Question of Realism. *Philosophers' Imprint*, 1, pp. 1-30.
- Fine, K. (2010). Some Puzzles of Ground. *Notre Dame Journal of Formal Logic*, 51, pp. 97–118.
- Finlay, S. and Snedegar, J. (2014). One Ought Too Many. *Philosophy and Phenomenological Research*, 89(1), pp. 102-124.
- Fodor, J. (2009). Where Is My Mind? Review of Andy Clark, *Supersizing the Mind*, *London Review of Books*, 31(3).
- Frege, G. (1879/1952). Begriffsschrift. Translated by P.T. Geach. In P.T. Geach and M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*, pp. 1-20. New York: Philosophical Library.
- Frege, G. (1892/1952). On Sense and Reference. Translated by M. Black. In P. Geach and M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*, pp. 56-78. New York: Philosophical Library.
- Galton, A. (2014). On Generically Dependent Entities. *Applied Ontology*, 9(2), pp. 129-53.
- Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L. (2002). Sweetening Ontologies with DOLCE. In *International Conference on Knowledge Engineering and Knowledge Management*, pp. 166-181. Berlin: Springer.

- Gangemi, A., Sagria, M.-T., and Tiscornia, D. (2005). A Constructive Framework for Legal Ontologies. In V.R. Benjamins, P. Casanovas, J. Breuker, A. Gangemi (eds.), *Law and the Semantic Web*, pp. 97-124. *Lecture Notes in Computer Science*, 3369. Berlin: Springer.
- Garcia, J. L. A. (1987). Constitutive Rules. *Philosophia*, 17, pp. 251–270.
- Giddens, A. (1984). *The Constitution of Society: Outline of the Theory of Structuration*. London: Hutchison.
- Gilbert, M. (1989). *On Social Facts*. London: Routledge.
- Gilbert, M. (1996). *Living Together: Rationality, Sociality, and Obligation*. Lanham, MD: Rowman & Littlefield.
- Gilbert, M. (2000). *Sociality and Responsibility: New Essays in Plural Subject Theory*. Lanham, MD: Rowman & Littlefield.
- Gilbert, M. (2014). *Joint Commitment: How We Make the Social World*. New York: Oxford University Press.
- Gomez-Perez, A., Fernandez-Lopez, M., and Corcho, O. (2004). *Ontological Engineering: with examples from the areas of knowledge management, e-commerce and the semantic web*. London: Springer.
- Guarino, N. and Giaretta, P. (1995). Ontologies and Knowledge Bases: Towards a Terminological Clarification. In N. Mars (ed.), *Toward Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing (KBKS '95)*, pp. 25-32. University of Twente, The Netherlands. Amsterdam: IOS Press.
- Habib, A. (2018). Promises. In E.N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Spring 2018), URL = <<https://plato.stanford.edu/archives/spr2018/entries/promises/>>.
- Harman, G. (1996). Moral Relativism. In G. Harman and J.J. Thomson (eds.), *Moral Relativism and Moral Objectivity*. Cambridge: Blackwell.
- Hart, H.L.A. (1955). Are There Any Natural Rights? *Philosophical Review*, 64, pp. 175-191.
- Hart, H.L.A. (1961). *The Concept of Law*. Oxford: Clarendon.
- Hastings, J., Ceusters, W., Jensen, M., Mulligan, K., and Smith, B. (2012). Representing mental functioning: Ontologies for mental health and disease. *Proceedings from: Towards an Ontology of Mental Functioning – Third International Conference on Biomedical Ontology*, July 22, 2012.
- Hayes, E.C. (1911). The Classification of Social Phenomena. *American Journal of Sociology*, 17(1), pp. 90-118.

- Heidegger, M. (1962). *Being and Time*. J. Macquarrie and E. Robinson (trans.). Oxford: Basil Blackwell.
- Heil, J. (2009). Relations. In R. Le Poidevin, et al. (eds.), *Routledge Companion to Metaphysics*, pp. 310–21. London: Routledge.
- Heil, J. (2012). *The World as We Find It*. Oxford: Clarendon Press.
- Hendricksen, Christopher. (2014). An Analysis of the Grounding Relation. Doctoral dissertation at the University of Illinois at Urbana-Champaign, Department of Philosophy.
- Herrestad, H. and Krogh, C. (1995). Deontic Logic Relativised To Bearers and Counterparties. In J. Bing and O. Torvund (eds.), *25 Years Anniversary Anthology NRCCL*. Oslo: TANO Publishers.
- Hindriks, F. (2009). Constitutive Rules, Language, and Ontology. *Erkenntnis*, 71(2), pp. 253-275.
- Hodgson, G.M. (2006). What Are Institutions? *Journal of Economic Issues*, 40(1), pp. 1-25.
- Hoekstra, R., Breuker, J., Di Bello, M., and Boer, A. (2007). The LKIF Core Ontology of Basic Legal Concepts. In P. Casanovas, M.A. Biasiotti, E. Francesconi, and M.T Sagri (eds.), *Proceedings of the Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT 2007)*, June 2007.
- Hoeltje, M., Schneider, B., and Steinberg, A. (eds.). (2013). *Varieties of Dependence: Ontological Dependence, Grounding, Supervenience, Response-Dependence*. Munich: Philosophia Verlag.
- Hoffmann, K. (1987). Reinach and Searle on Promising — A Comparison. In Mulligan (1987b), pp. 91-106.
- Hoffman, J. and Rosenkrantz, G.S. (1994). *Substance among Other Categories*. Cambridge: Cambridge University Press.
- Hofweber, Thomas. (2009). Ambitious, Yet Modest, Metaphysics. In (Chalmers, et al., 2009, pp. 260-289).
- Hohfeld, W.N. 1919. *Fundamental Legal Conceptions*. W. Cook (ed.). New Haven: Yale University Press.
- Hume, D. (1740/1975). *A Treatise of Human Nature*. 2nd edition. L. A. Selby-Bigge (ed.), P. H. Nidditch (rev.). Oxford: Clarendon Press.
- Hume, D. (1777/1975). *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. L.A. Selby-Brigge (ed.), revised 3rd edn., P. H. Nidditch (ed.). Oxford: Clarendon Press.

- Husserl, E. (1900/1984). *Logische Untersuchungen*. The Hague: Martinus Nijhoff.
- Ibn Sina (Avicenna). (1973). *The Metaphysica of Avicenna (Ibn Sina)*. P. Morewedge (trans.). London: Routledge & Kegan Paul.
- Jansen, L. (2008). Categories: The Top-Level Ontology. In (Munn and Smith, 2008), pp. 173-96.
- Jensen, M., Cox, A.P., Donohue, B., and Rudnicki, R. (2018). Problems with Prescriptions: Distangling Data about Actual versus Prescribed Entities. Forthcoming in *Proceedings of SPIE Defense and Security: Ground/Air Multisensor Interoperability, Integration, and Networking for ISR IX*.
- Kant, I. (1785/2012). *Groundwork of the Metaphysics of Morals*. M. Gregor and J. Timmerman (trans.). Cambridge: Cambridge University Press.
- Leibniz, G.W.F. (1989). *Philosophical Essays*. R. Ariew and D. Garber (trans., eds.). Indianapolis: Hackett Publishing Company.
- Lewis, D. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, D. (1997). Finkish Dispositions. *The Philosophical Quarterly*, 47, pp. 143-58.
- Lipps, T. (1899). *Die ethischen Grundfragen: zehn Vorträge*. Hamburg/Leipzig: Verlag von Leopold.
- Lipps, T. (1909). *Leitfaden der Psychologie*. Leipzig: Engelmann.
- Lowe, E.J. (1998). *The Possibility of Metaphysics*. Oxford: Oxford University Press.
- Lowe, E.J. (2013). A Neo-Aristotelian Substance Ontology: Neither Relational nor Constituent. In A. Tahko (ed.), *Contemporary Aristotelian Metaphysics*, pp. 229-48. Cambridge: Cambridge University Press.
- MacBride, F. (2011). Relations and Truth-Making. *Proceedings of the Aristotelian Society*, 61, pp. 159–76.
- Martin, C.B. (1994). Dispositions and Conditionals. *The Philosophical Quarterly*, 44, pp. 1-8.
- McLaughlin, B.P. (1995). Varieties of Supervenience. In E. Savellos, and U. Yalcin (eds.), *Supervenience: New Essays*, pp. 16-59. Cambridge: Cambridge University Press.
- Merrill, G.H. (2010a). Ontological Realism: Methodology or Misdirection? *Applied Ontology*, 5(2), pp. 79-108.

- Merrill, G.H. (2010b). Realism and Reference Ontologies: Considerations, Reflections, and Problems. *Applied Ontology*, 5 (3-4), pp. 189-221.
- Miller, S. (2014). Social Institutions. E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition), URL = <<https://plato.stanford.edu/archives/win2014/entries/social-institutions/>>.
- Molnar, G. (2003). *Powers: A Study in Metaphysics*. Oxford: Oxford University Press.
- Mulligan, K. (1987a). Promisings and other Social Acts: Their Constituents and Structure. In Mulligan (1987b), pp. 29-90.
- Mulligan, K. (ed.). (1987b). *Speech Act and Sachverhalt: Reinach and the Foundations of Realist Phenomenology*. Dordrecht: Martinus Nijhoff.
- Mumford, S. (1998). *Dispositions*. Oxford: Oxford University Press.
- Munn, K. and Smith, B. (eds.). (2008). *Applied Ontology: An Introduction*. Frankfurt: Ontos Verlag.
- Nolan, D. (2011). Categories and Ontological Dependence. *The Monist*, 94(2), pp. 277-301.
- North, D.C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge, UK: Cambridge University Press.
- Owens, D. (2006). A Simple Theory of Promising. *Philosophical Review*, 115(1), pp. 51-77.
- Owens, D. (2008). Promising without Intending. *Journal of Philosophy*, 105(12), pp. 737-755.
- Peirce, C.S. (1904/1976). New Elements. In Carolyn Eisele (ed.), *The New Elements of Mathematics by Charles S. Peirce, Volume 4, Mathematical Philosophy*. The Hague: Walter de Gruyter.
- Plantinga, A. (1974). *The Nature of Necessity*. Oxford: Oxford University Press.
- Poincare, H. (1898). On the Foundations of Geometry. *The Monist*, 9, pp. 1-43.
- Popper, K. (1945). *The Open Society and Its Enemies, Vol. 1*. Princeton: Princeton University Press.
- Prichard, H.A. (1968). *Moral Obligation and Duty and Interest*. New York: Oxford.
- Quine, W.V.O. (1948). On What There Is. *The Review of Metaphysics*, 2(1), pp. 21-38.
- Quine, W.V.O. (1956). Quantifiers and Propositional Attitudes. *The Journal of Philosophy*, 53(5), pp. 177-87.
- Randsell, J. (1971). Constitutive Rules and Speech-Act Analysis. *The Journal of Philosophy*, 68, pp. 385-399.

- Raven, M. (2013). Is Ground a Strict Partial Order? *American Philosophical Quarterly*, 50, pp. 191-199.
- Rawls, J. (1955). Two Concepts of Rules. *Philosophical Review*, 64(1), pp. 3-32.
- Rawls, J. (1972). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Raz, J. (1975). *Practical Reasons and Norms*. Oxford: Clarendon.
- Raz, J. (1977). Promises and Obligations. In P. M.S. Hacker and J. Raz (eds.), *Law, Morality and Society: Essays in Honor of H.L.A. Hart*, pp. 210-228. Oxford: Clarendon.
- Reid, T. (1785/2002). *Essays on the Intellectual Powers of Man*. Derek Brookes (ed.). University Park: Pennsylvania State University Press.
- Reid, T. (1788/2010). *Essays on the Active Powers of Man*. Knud Haakonssenn and James A. Harris (eds.). University Park: Pennsylvania State University Press.
- Reinach, A. (1913/1983). The Apriori Foundations of the Civil Law. J.F. Crosby (trans.). *Aletheia*, 3, pp. 1-142.
- Rescorla, M. (2017). Convention. E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), URL = [<https://plato.stanford.edu/archives/sum2017/entries/convention/>](https://plato.stanford.edu/archives/sum2017/entries/convention/).
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In R. Hale and A. Hoffman (eds.), *Modality: Metaphysics, Logic, and Epistemology*, pp. 109-136. Oxford: Oxford University Press.
- Russell, B. (1903). *The Principles of Mathematics*. London: Norton.
- Russell, B. (1912). *The Problem of Universals*. London: William & Norgate.
- Ryle, G. (1949). *The Concept of Mind*. London: Penguin Books.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schaffer, J. (2009). "On What Grounds What". In (Chalmers, et al., 2009, pp. 347-383).
- Schelling, T.C. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Schneider, B. (2006). A Certain Kind of Trinity: Dependence, Substance, Explanation. *Philosophical Studies*, 129, pp. 393-419.
- Schumann, K. and Smith, B. (1990). Elements of Speech Act Theory in the Work of Thomas Reid. *History of Philosophy Quarterly*, 7(1), pp. 47-66.

- Schuppe, W. (1881). *Grundzuge der Ethik und Rechtsphilosophie*. Warsaw, Poland: Wilhelm Koebner.
- Searle, J.R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Searle, J.R. (1975). A Taxonomy of Illocutionary Acts. *Minnesota Studies in the Philosophy of Science*, 9, pp. 344-369.
- Searle, J.R. (1980). Mind, Brains and Programs. *Behavioral and Brain Sciences*, 3(3), pp. 417-57.
- Searle, J.R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Searle, J.R. (1984). Intentionality and Its Place in Nature. *Synthese*, 61(1), pp. 3-16.
- Searle, J.R. (1995). *The Construction of Social Reality*. New York: The Free Press.
- Searle, J.R. (1998). *Mind, Language, and Society*. New York: Basic Books.
- Searle, J.R. (2005). What is an Institution? *Journal of Institutional Economics*, 1, pp. 1-22.
- Searle, J.R. (2010). *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.
- Seneca. (1925). *Ad Lucilium: Epistulae Morales*, vol. 3. Translated by R.M. Gummere. New York: G.P. Putnam's Sons.
- Shapiro, S.J. (2014). Massively Shared Agency. In M. Vargas and G. Yaffe (eds.), *Rational and Social Agency: The Philosophy of Michael Bratman*. Oxford: Oxford University Press.
- Shiffrin, S. (2012). Are Contracts Promises? In A. Marmor (ed.), *The Routledge Companion to the Philosophy of Law*. London: Routledge.
- Sider, T. (1996). Intrinsic Properties. *Philosophical Studies*, 83(1), pp. 1-27.
- Sider, T. (2012). *Writing the Book of the World*. Oxford: Oxford University Press.
- Smith, B. (ed.). (1982). *Parts and Moments. Studies in Logic and Formal Ontology*. Munich: Philosophia.
- Smith, B. (1984). Ten Conditions on a Theory of Speech Acts. *Theoretical Linguistics*, 11, pp. 311-30.
- Smith, B. (1990). Towards a History of Speech Act Theory. In A. Burkhardt (ed.), *Speech Acts, Meanings and Intentions: Critical Approaches to the Philosophy of John R. Searle*, pp. 29-61. Berlin/New York: Walter de Gruyter.

- Smith, B. (1992). An Essay on Material Necessity. *Canadian Journal of Philosophy*, 22, pp. 301-322.
- Smith, B. (1998). Basic Concepts in Formal Ontology. In N. Guarino (ed.), *Formal Ontology in Information Systems*, pp. 19-28. Amsterdam: IOS Press.
- Smith, B. (2003a). John Searle: From Speech Acts to Social Reality. In B. Smith (ed.), *John Searle*, pp. 1-33. Cambridge: Cambridge University Press.
- Smith, B. (2003b). The Ontology of Social Reality. *American Journal of Economics and Sociology*, 62(2), pp. 285-299.
- Smith, B. (2004). Beyond Concepts: Ontology as Reality Representation. In A. Varzi and L. Vieu (eds.), *Proceedings of FOIS 2004: International Conference on Formal Ontology and Information Systems*. Turin, 4-6 November 2004.
- Smith, B. (2008). Searle and de Soto: The New Ontology of the Social World. In B. Smith, D. Mark, and I. Ehrlich (eds.), *The Mystery of Capital and the Construction of Social Reality*, pp. 35-51. Chicago and La Salle: Open Court.
- Smith, B. (2012a). Classifying Processes: An Essay in Applied Ontology. *Ratio*, 25(4), pp. 463-488.
- Smith, B. (2012b). How to Do Things with Documents. *Rivista di Estetica*, 50, pp. 179-198.
- Smith, B. (2013). Diagrams, Documents, and the Meshing of Plans. In A. Benedek and K. Nyiri (eds.), *Visual Learning, vol. 3: How to Do Things with Pictures: Skill, Practice, Performance*, pp. 165-179.
- Smith, B. (2014). Document Acts. In A.K. Ziv and H.B. Schmid (eds.), *Institutions, Emotions, and Group Agents: Contributions to Social Ontology*, pp. 19-31. Dordrecht: Springer.
- Smith, B. and Ceusters, W. (2010). Ontological Realism: A Methodology for the Coordinated Evolution of Scientific Ontologies. *Applied Ontology*, 5(3-4), pp. 139-188.
- Smith, B. and Ceusters, W. (2015). Aboutness: Toward Foundations for the Information Artifact Ontology. *Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO)*. Lisbon, Portugal (CEUR 1515).
- Smith, B. and Żelaniec, W. (2012). Laws of Essence or Constitutive Rules? Reinach vs. Searle on the Ontology of Social Entities. In F. De Vecchi (ed.), *Eidetica del Dritto e Ontologia Sociale: Il Realismo di Adolf Reinach*, pp. 83-108. Milan: Mimesis.
- Smith, M. (1987). The Humean Theory of Motivation. *Mind*, 96(381), pp. 36-61.
- Spear, A.D., Ceusters, W., and Smith, B. (2016). Functions in Basic Formal Ontology. *Applied Ontology*, 11, pp. 103-128.

- Stamper, R.K. (1991). The Role of Semantics in Legal Expert Systems and Legal Reasoning. *Ratio Juris*, 4(2), pp. 219-244.
- Stamper, R.K. (1996). Signs, Information, Norms and Systems. In B. Holmqvist and P. B. Andersen (eds.), *Signs of Work*. Berlin: De Bruyter.
- Steiner, H. (1994). *An Essay on Rights*. Oxford: Blackwell.
- Sugden, R. (1998). The Role of Inductive Reasoning in the Evolution of Conventions. *Law and Philosophy*, 17: 377-410.
- Sugden, R. (1986/2004). *The Economics of Rights, Co-operation, and Welfare*. 2nd edn. New York: Palgrave Macmillan.
- Thomson, J.J. (1990). *The Realm of Rights*. Cambridge, MA: Harvard University Press.
- Tuomela, R. (1995). *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford: Stanford University Press.
- Tuomela, R. (2002). *The Philosophy of Social Practices: Collective Acceptance View*. Cambridge, UK: Cambridge University Press.
- Tuomela, R. (2005). We-Intentions Revisited. *Philosophical Studies*, 125: 327-69.
- Tuomela, R. (2007). *The Philosophy of Sociality: The Shared Point of View*. Oxford: Oxford University Press.
- Tuomela, R. (2013). *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.
- Turner, J. (1997). *The Institutional Order*. New York: Longman.
- Valente, A. (1995). *Legal Knowledge Engineering: A Modelling Approach*. Amsterdam: IOS Press.
- Vanderschraaf, P. (1995). *Learning and Coordination: Inductive Deliberation, Equilibrium, and Convention*. London: Routledge.
- Varzi, A. (2010). On the Boundary between Material and Formal Ontology. In B. Smith, R. Mizoguchi, and S. Nakagawa (eds.), *Interdisciplinary Ontology*, Vol. 3. Tokyo: Keio University.
- Velleman, J.D. (1992). The Guise of the Good. *Nous*, 26(1), pp. 3-26.
- Visser, P. R.S. and Bench-Capon, T.J.M. (1998). A Comparison of Four Ontologies for the Design of Legal Knowledge Systems. *Artificial Intelligence and Law*, 6, pp. 27-57.

Weirich, P. (1989). Conventions and Social Institutions. *The Southern Journal of Philosophy*, 27(4): 599-618.

Williams, B. (1981). *Ought* and Moral Obligation. In *Moral Luck*, pp. 114-123. Oxford: Oxford University Press.

Wilson, J.M. (2014). No Work for a Theory of Grounding. *Inquiry: An Interdisciplinary Journal of Philosophy*, 57(5-6), pp. 535-579.

Wittgenstein, L. (1953/2009). *Philosophical Investigations (German-English)*, 4th edn. G.E.M. Anscombe, P. M.S. Hacker, and Joachim Schulte (trans.). Malden, MA: Blackwell.