Collective Abstraction*

Jon Erling Litland

Abstract

This paper develops a novel theory of abstraction—what we call collective abstraction. The theory solves a notorious problem for non-eliminative structuralism. The non-eliminative structuralist holds that in addition to various isomorphic systems there is a *pure* structure that can be abstracted from each of these systems; but existing accounts of abstraction fail for non-rigid systems like the complex numbers. The problem with the existing accounts is that they attempt to define a *unique* abstraction operation. The theory of collective abstraction instead simultaneously defines a *collection* of distinct abstraction operations, each of which maps a system to its corresponding pure structure. The theory is precisely formulated in an essentialist language. This allows us to throw new light on the question to what extent structuralists are committed to symmetric dependence. Finally, we apply the theory of collective abstraction to solve a problem about converse relations.

Keywords: Structuralism, Non-rigid structures, Abstraction, Generation, Essence

*An early version of the material in this paper was presented at a workshop at the University of Oslo in 2017. Thanks to the participants. Thanks especially to Øystein Linnebo for many discussions about collective abstraction and to Ray Buchanan for discussions about generativism more generally. I am extremely grateful to Tim Button for comments that showed me that a previous version of the theory contained a serious flaw. Finally, I would like to thank the anonymous referees and the editors for persistent and constructive criticisms that greatly improved the paper both substantially and presentationally.

1 Introduction

Abstract objects pose a metaphysical problem: what *are* they? They also pose a (meta)semantic problem: how can we *refer* to them? Generativism provides an answer to both questions. Abstract objects are, by their nature, values of certain generative operations, different types of abstract objects being the values of different operations: sets being the values of the set-formation operation, sequences being the values of the sequence-formation operation, and so on.[1] We can *refer* to abstract objects generated by an operation Σ by coming to understand an operator "$\mathbf{\Sigma}$" standing for that operation.[2] Expressions of the form "$\mathbf{\Sigma}(b)$" (where "b" is a name for an object to which Σ is applied) then serve as "canonical" names for objects that are values of the operation Σ.[3] Unfortunately, existing generativist accounts have a serious limitation: they are unable to generate indiscernible objects like the square roots i and $-i$ of -1. To explain the problem let us introduce the main theme of this paper: structural abstraction.

There are many systems of objects that are structured like the natural numbers, but which of them are *the* natural numbers? This is the Identification Problem of Benacerraf 1965. The generativist offers an answer: there is a generative operation of (structural) *abstraction*—call it A. This operation takes a system S and an object a in S and generates an object $A(S, a)$. The operation

[1]For an early statement of generativism, see Fine 1991. For an application to composition in particular, see Fine 2010b. As a way of characterizing abstract objects, generativism is an instance of what Lewis (1986, 84-86) calls the "Way of Abstraction"; for more on the Way of Abstraction and generativism, see Rosen 2020.

[2]We will use "operator" to mean a part of language and "operation" to mean a worldly item for which an operator stands.

[3]Prominent defenders of this view include Dummett (1973, 1991), Wright (1983), and Hale and Wright (2001).

is defined by an abstraction principle of the following form:

(SA) A generates the same object from (S,a) and (T,b) iff a plays the same role in system S as b plays in system T.

If we take the third element of one system with the natural number structure and the third element of another, A will generate the same object from both—we may call it the (ordinal) number 2. More generally, what the natural numbers *are*—so the generativist says—are the objects generated by this operation A from any system that has the natural number structure.

As shown by Linnebo (2008) and Linnebo and Pettigrew (2014) this works well for rigid systems (that is, systems without non-trivial symmetries); but, as is well known, it fails catastrophically for non-rigid systems (that is, systems with non-trivial symmetries)—a canonical example being the complex numbers. The generativist would want to say that the square roots of -1—that is, i and $-i$—are what can be abstracted from any objects $\mathbf{i}, -\mathbf{i}$ that play the roles of roots of -1 (in some system $\mathbb{D}$). But it turns out that any role played by $\mathbf{i}$ is also played by $-\mathbf{i}$ (and *vice versa*).[4] The structural abstraction operation A therefore has to generate the *same* object both from $\mathbf{i}$ and from $-\mathbf{i}$; absurdly, we cannot generate *two* roots of -1.

This problem of how to account for indiscernible objects like i and $-i$ has bedeviled non-eliminative structuralists for decades.[5] The main contribution of this paper is the development of a novel account of structural abstraction—what we call *collective abstraction*. Briefly, we give up the idea that there is a *unique* abstraction operation A; rather, there are several distinct operations

[4]For details and other examples see § 2.2 and § 4.

[5]See, e.g., Brandom 1996, Burgess 1999, Keränen 2001, Linnebo 2008, Linnebo and Pettigrew 2014, Normand 2018, Wigglesworth 2021, Button 2006, Parsons 2004, Ladyman 2005; Leitgeb and Ladyman 2008.

that give rise to the same abstracts. These operations are not individually defined by abstraction principles like (SA); rather, there is a single principle that defines these operations collectively (or simultaneously). We rigorously develop this theory of collective abstraction using a higher-order theory of essence. While this involves some heavy-duty metaphysics, the pay-offs are worth it. For the reader's benefit here is an overview of the main results.

In § 2 we develop a generativist account of abstraction on rigid systems. We begin in § 2.1 by giving a generativist account of the familiar Fregean case of abstracting directions from lines; in §§ 2.2 to 2.4 we then generalize these ideas to structural abstraction on rigid systems. The core mathematical results are not novel (see e.g., Linnebo and Pettigrew 2014); what is novel is the use of a higher-order logic of essence (§§ 2.2 to 2.3). This allow us to state the generativist's view with unprecedented precision.

In § 3, we show that invoking a higher-order logic of essence affords a precise development of generativist non-eliminative structuralism. Non-eliminative structuralists hold that in addition to various systems of objects there are *pure* structures containing pure positions canonically representing those systems. The generativist solves two problems for the non-eliminative structuralist. First, while non-eliminative structuralists have made many suggestive remarks about the nature of pure structures and pure positions we still lack a systematic metaphysics for pure structures. We develop such a metaphysics in § 3.1. In § 3.2 we then use the essentialist framework to investigate on what the pure structures and positions depend. One consequence is worth highlighting. It has often been claimed that a pure position cannot be made sense of on its own but only together with its fellow pure positions. This is false: each pure position has an individual essence that distinguishes it from all other pure positions. The non-eliminative structuralist's second problem is that, even equipped with a metaphysics of pure structures, she still faces the Identification Problem: amongst all the systems that are structured like the natural numbers there is a unique "pure" structure, but what makes this *the* natural number structure? The essentialist framework provides the solution here too: the pure structures are the unique systems whose natures are exhausted by representing the systems isomorphic to them (§ 3.3).

The above generativist account only works for rigid systems; the heart and most significant contribution of the paper is in § 4 where we develop the account of collective abstraction. In §§ 4.1 to 4.2 we generalize the idea of a Fregean abstraction principle by laying down a principle that simultaneously defines a collection of abstraction operations. In § 4.3 we then examine the nature of the pure structures and positions generated by collective abstraction. Strikingly, we can no longer make sense of pure positions individually: they no longer have individual essences that distinguish them from each other. (We might say that they are "entangled".) In § 4.4 we discuss the relationship between individual essences and haecceitistic properties like *being identical to p*: while the pure positions do not have individual essences there is, for every pure position p, the property of being identical to p. In § 4.5 we return to the Identification Problem.

In § 5 we turn to the metasemantic problem: how can we refer to pure positions? In particular, how can we refer to indiscernible pure positions like i and $-i$? We answer this question by developing an account of reference relative to an abstraction operation, absolute reference being reference with respect to all abstraction operations. It then turns out that we can (absolutely) refer to the roots of -1—that is, to $i, -i$—though not to any one of them. In this way we partially vindicate a Fregean conception of objects as the possible referents of terms.

The techniques of collective abstraction are applicable elsewhere in metaphysics. In § 6 we illustrate this by developing a novel version of positionalism about relations, solving a notorious problem about converse relations. In § 7 we conclude. A technical appendix that establishes the consistency of the theory of collective abstraction by modifying the model theory for higher-order contingentism is available at <INSERT LINK TO APPENDIX HERE>.

2 Generativism: the Rigid Case

There is an intuitive distinction between generative and non-generative operations.[6] The operation *set of*, e.g., is generative in the sense that we can explain what the set $\{a, b\}$ *is* by noting that it is what is formed by applying the *set of* operation to the plurality a, b. In contrast, the *tallest of* operation is non-generative. If Bob is the tallest of the *xx*, his being the tallest of the *xx* is not part of his identity—it is not part of what he *is*.

To explain this distinction between generative and non-generative operations we use the notion of essence. An operation G is *generative* iff for all x, if $x = Ga$, then it is essential to x that there is some y such that $x = Gy$. A possibly different explanation uses the notion of real definition.[7] An operation G is generative if for any x such that $x = G(y)$, for some y, there is some z such that x can be *defined*—in the sense of *real* definition—as the value of G on input z. For this latter approach to work it is important that an object can have several distinct definitions. Suppose, to take a familiar example, we think that the operation taking a line to its direction is generative. We might then hold that a given direction d can be defined as the direction of line l_1; but we could equally well have defined d as the direction of any line l_2 parallel to l_1. The direction d has a manifold of possible definitions—what Fine (1994b, 66–69) calls an *essential manifold*. Since we do not yet posses a satisfactory logic of real definition, we will talk in terms of essence when we give precise

[6] The terminology goes back to N. Goodman 1958. (He uses "relations that generate ".)

[7] Some authors have tried to use essence (and ground) to (really!) define real definition—see e.g., Correia 2017, Rosen 2015. If their approaches work the two approaches come to the same.

formulations of the generativist's theses; however, we will often use the idiom of real definition for illustrative purposes.

To define a generative operation we have to specify three conditions:[8]

(i) application-conditions—telling us when the generative operation can be applied;

(ii) identity-conditions—telling us under what conditions the values of the operation are identical; and

(iii) existence-conditions—telling us under what conditions the values of the operation exist.

Rather than explain how these conditions are to be understood in the abstract we informally elucidate them using the familiar example of direction abstraction.[9]

2.1 The case of directions

The Fregean wishes to define an operation d that, given a line l, generates its direction $d(l)$. The application condition for d states that for any line l its direction $d(l)$ is defined. Many lines have the same direction, so if $x = d(l)$ it is not essential to x that it be the direction of l; however, the application condition should ensure that if $x = d(l)$ then it is essential to x that it can be abstracted from *some* line. Finally, while it is not essential to x that it be the direction of l it is no accident that x can be generated from l so the application condition should ensure that it is essential to x together with l that $d(l) = x$.

In general, the identity conditions for a generative operation G are given by specifying an equivalence relation $\approx_G$ on its inputs; the outputs $G(a)$ and $G(b)$ are to be identical if and only if

[8]These are some of the principles discussed in Fine 2010b, 570.

[9]The next section is inspired Yablo and Rosen 2020, though they focus mainly on the case of Fregean abstraction of numbers. See also Fine 2002, 29-31.

$a \approx_G b$. The identity-condition for directions, famously, is *parallelism*. We thus have the following identity criterion (using $l_1 \parallel l_2$ to mean that the lines l_1 and l_2 are parallel):

(DI) $\quad d(l_1) = d(l_2) \leftrightarrow l_1 \parallel l_2$

The existence condition states that a direction x exists if and only if a line l such that $d(l) = x$ exists; in fact, it is natural to take the existence condition to say that the existence of x is *grounded* in the existence of any l such that $x = d(l)$.[10]

The generativist now holds that the direction-of operation d is implicitly defined by its satisfying the above application, identity, and existence conditions. For this to work we have to take the nature of the operation d to be *exhausted* by its satisfying these conditions. To see this, let us focus on the identity condition (DI). An operation e's satisfying (DI) requires only that e is an operation mapping lines to objects such that parallel lines are mapped to the same object and non-parallel lines are mapped to distinct objects. But if there is *one* such operation there are many—for instance, there is the operation that is exactly like the genuine direction-of operation except that it maps the lines parallel to some particular line l to the emperor Diocletian. How, then, could (DI) define a unique operation? Is not the best we can hope for that (DI) defines the property of being a direction-of operation, that is, the property $\lambda O.\forall l_1 \forall l_2 (O(l_1) = O(l_2) \leftrightarrow l_1 \parallel l_2)$?[11] Requiring that d also satisfies the application and existence conditions further constrains matters, but it seems unlikely that there is a unique operation satisfying all of these conditions.

[10]Let us understand ground as constitutive explanation: for any proposition ϕ, if it is the case that ϕ, we can ask in what its being the case that ϕ consists and the answer to that question will be the grounds for ϕ. For more on ground and constitutive explanation, see Fine 2012a and Dasgupta 2017.

[11]The view that this is all that abstraction principles can do is defended in Antonelli 2010 and is further developed in Boccuni and Woods 2020.

While there are many operations that satisfy the above application, identity, and existence conditions, the generativist contends that there is a unique operation d the essence of which is exhausted by satisfying them. It is this that allows the generativist to solve the Caesar Problem. This, recall, is the problem of explaining why a given direction—$d(l)$ say—is not identical to a Roman emperor. The key is that—according to the Generativist—the *values* of d are themselves fully defined by being the values of d. So if $x = d(l)$ happened to be Diocletian, this would have to be a consequence of the nature d together with the nature of l. But this is impossible: the nature of d is exhausted by the application, identity, and existence conditions, and these conditions make no mention of Diocletian—or any other Roman emperor.[12]

We now turn to making these ideas more precise and generalizing them to structural abstraction operations.

2.2 Type-theoretic background

The generativist's view turns on the relationship between systems of objects, certain operations, and the objects that are the values of those operations. This is perspicuously expressed in the language of higher-order logic.[13] Specifically, we will work in a relational type theory. We have the basic type e

[12] The generativist's solution to the Caesar problem is inspired by Yablo and Rosen 2020. But one important difference from their view is that the generativist takes it as part of the definition of the generative operations that their values are essentially values of them. This allows the generativist to get around some objections that Wright (2020, 319-320) raised against Rosen and Yablo's view.

[13] Throughout the paper we take for granted that higher-order logic is an appropriate framework for doing metaphysics, though for the applications in the present paper we do not have to take a stand on whether higher-order logic is irreducible to first-order logic. The reader who prefers to think in first-order terms can take quantification into predicate position to be first-order quantification over properties; quantification over operations to be quantification over certain relations between

of objects. We have two ways of forming complex types. If $\tau_0, \ldots, \tau_{n-1}$ are types, then $\langle \tau_0, \ldots, \tau_{n-1} \rangle$ is a type. If τ is a type, then $[\tau]$ is a type. Intuitively, $\langle \tau_0, \ldots, \tau_{n-1} \rangle$ is the type of relations between items of types $\tau_0, \ldots, \tau_{n-1}$ and $[\tau]$ is the type of pluralities of items of type τ. Note that $\langle\rangle$ is a type—the type of 0-ary relations, or as we will call it here: the type of propositions.

We assume that our typed language contains, for each type τ, infinitely many variables $x_0^\tau, x_1^\tau, \ldots$ of that type (the superscripts indicating the type of the variable). We introduce constants of various types as we need them. We assume that we have the usual logical operations and that the quantifiers can bind variables of any type. We will understand the quantifiers possibilistically as ranging over all items that are possible (relative) to a world; to express that an item of type τ (actually) exists we use an existence predicate for items of type τ.[14]

Whenever t is an expression of type $\langle \tau_0, \ldots, \tau_{n-1} \rangle$ and $s_0, \ldots, s_{n-1}$ are expressions of types $\tau_0, \ldots, \tau_{n-1}$ respectively, $t(s_0 \ldots s_{n-1})$ is an expression of type $\langle\rangle$. When t is an expression of type $[\tau]$ and s is an expression of type τ, $t(s)$ is an expression of type $\langle\rangle$. Expressions of type $\langle\rangle$ are referred to as sentences. In addition, we allow non-vacuous λ-abstraction with respect to any variables. Thus, when ϕ is a sentence containing the variable x^τ, $\lambda x^\tau.\phi$ is a monadic predicate ("being such that $\phi(\text{it}^\tau)$"); and when ϕ is a sentence containing the variables $x_0^{\tau_0}, \ldots, x_{n-1}^{\tau_{n-1}}$ then $\lambda x_0^{\tau_0} \ldots \lambda x_{n-1}^{\tau_{n-1}}.\phi$ is an n-ary predicate.

To help readability we adopt a range of conventions. We normally use uppercase letters $R, S, T, \ldots$ as variables of relational type, leaving their exact type to be determined by context.

pluralities and relations, and so on. For defenses of the intelligibility and irreducibility of higher-order logic see e.g., Williamson 2003, 2013, Dorr 2016, J. Goodman 2017, and Prior 1971.

[14]We opt for a possibilistic understanding of the quantifiers for two reasons. First, this is in keeping with how the quantifiers are understood in the logic of essence (Fine 1995b, 2000b). Second, while we could understand the quantifiers actualistically this would complicate the statement of the generativist's principles.

While we officially have a distinct existence predicate E^{τ} for each type τ, we use E with typical ambiguity for any of these predicates. For instance, we write $E(x) \wedge E(R)$ instead of $E^{\langle e \rangle} x \wedge E^{\langle\langle e,e \rangle\rangle}(R)$. Since we work in a relational type theory we officially think of functions as functional relations[15]—though if F is a functional relation, we typically write $Fx = y$ instead of Fxy. Moreover, to stay close to mathematical practice we often use $f, g, \ldots$ as variables for functions. Finally, when $a, b, c, \ldots$ are all expressions of a given type τ we write $[a, b, c, \ldots]$ to mean the plurality of type $[\tau]$ that contains exactly (the denotations of) $a, b, c, \ldots$. If I is a plurality and there is a function that associates with each i in I an item a_i we write $[a_i \colon i \in I]$ to mean the plurality of all such a_i. (We here assume that all the a_i are of the same type.) If $t^{[\tau]}$, and $s^{[\tau]}$ are two pluralities we write $t \subseteq s$ to mean that $\forall x^{\tau}(t(x) \to s(x))$.[16]

We will think of a *system*[17] in higher order terms as consisting of a set-sized[18] plurality of objects $D^{[e]}$, and a binary relation $R^{\langle e,e \rangle}$ defined just on D. Formally, a system is a relation $S^{\langle [e], \langle e,e \rangle \rangle}$ such that $S(D, R)$ for exactly one set-sized plurality D and exactly one relation R; moreover, R only holds between objects in D and if Rab, then it is necessary that Rab.[19] If S is a system we often

[15]A relation F is functional iff for all x, y, z if Fxy and Fxz, then $y = z$.

[16]We will occasionally also use $\subseteq$ to mean the sub*set* relation, but context will disambiguate.

[17]This adapts the terminology of Shapiro 1997; see also Linnebo and Pettigrew 2014.

[18]We—following Linnebo 2008, 75—restrict our attention to set-sized pluralities to avoid running into the Burali-Forti paradox. Merely imposing the restriction that systems be set-sized is, of course, not ultimately satisfactory. However, the problems raised by non-rigid systems arise already for set-sized systems, so it is reasonable to restrict our attention to that easier case. Though we cannot go into this here, the theory of collective abstraction can be generalized to various ways of dealing with non-set-sized systems. For instance, it generalizes to the "dynamic" approach to abstraction defended by Linnebo (2009, 2018, 51-76). Alternatively, we could adopt a theory of classes like the one in Fine 2005 or Linnebo 2006 and let a system consist of a class together with a relation on that class.

[19]The restriction to a single binary relation is just for convenience. We can allow arbitrarily many

write $S = \langle D_S, R_S \rangle$ to indicate that D_S, R_S are the unique D, R such that $S(D,R)$. We call D_S the *domain* of S; and we call the objects that are amongst the D_S the *elements* of S. We sometimes abuse notation and write Sa to mean that $D_S a$. (Recall that Da means that a is one of the objects in D.)

Let $S_0 = \langle D_0, R_0 \rangle$ and $S_1 = \langle D_1, R_1 \rangle$ be two systems. We say that f is an isomorphism between S_0 and S_1 iff $f\colon D_0 \to D_1$ is a bijection such that for all a, b in D_0 we have

$$R_0(a,b) \Leftrightarrow R_1(f(a), f(b))$$

If f is an isomorphism between S_0 and S_1 we write $f\colon S_0 \cong S_1$; we write $S_0 \cong S_1$ to mean that there is some f such that $f\colon S_0 \cong S_1$. As usual, if $f\colon S \cong S$—that is, if f is an isomorphism from S to S itself—we say that f is an *automorphism*. For each $S = \langle D, R \rangle$ the identity function id_S defined by $\mathrm{id}_S(a) = a$ for all $a \in D$ is an automorphism on S. We say that S is *rigid* if for all automorphisms $g\colon S \cong S$, $g = \mathrm{id}_S$; S is *non-rigid* otherwise.[20]

Let $\mathcal{S} = [S_i\colon i \in I]$ be a plurality of systems. A *candidate abstraction operation* (or simply *candidate*) on $\mathcal{S}$ is a ternary relation A such that

(i) whenever $A(S, a, p)$ then $S = \langle D_S, R_S \rangle$ is in $\mathcal{S}$ and a is in D_S;

(ii) whenever $S = \langle D_S, R_S \rangle$ is in $\mathcal{S}$ and a is in D_S, there is a unique p such that $A(S, a, p)$;

relations of arbitrary (even infinite) arity; and we can also allow functions and designated objects. This will not affect the points made.

[20] A slight complication: officially, functions are relations, and since we work in a fine-grained setting we cannot assume that a relation that is (necessarily) coextensive with id_S is identical to id_S. Strictly speaking, what we should require for rigidity is just that if g is any automorphism on S, then $g(a) = a$, for all a. We will suppress this subtlety in what follows.

(iii) for all S_i, S_j in $\mathcal{S}$, if $S_i \cong S_j$ there is a unique isomorphism $f_{ij}\colon S_i \cong S_j$ such that for all a in D_{S_i} we have $A(S_i, a, p)$ iff $A(S_j, f_{ij}(a), p)$.

For readability we will use functional notation and write $A(S, a) = p$ instead of $A(S, a, p)$.

Given any plurality $\mathcal{S}$ of (small) systems there are several distinct candidate abstraction operations on $\mathcal{S}$: for each isomorphism class C of systems simply pick a system $|S|$ of that class, and for each system S in C pick an isomorphism $f\colon S \cong |S|$. Then let $A(S, a) = f_S(a)$. Distinct choices of $|S|$ and f_S yield distinct operations. It is this fact that lies at the core of Benacerraf's Identification Problem. If $\mathcal{S}$ is a plurality of systems having the natural number structure, a candidate abstraction operation for $\mathcal{S}$ offers a candidate answer to the question: which are *the* natural numbers? A crucial first step in solving the Identification Problem is to determine which candidate is *the* abstraction operation.

2.3 Essentialist background

We now turn to developing a precise way of expressing the generativist's essentialist claims; but before we dive into the details, a methodological remark. A lot of effort has been spent on developing the formal foundations for essentialist metaphysics;[21] but nobody has developed an essentialist metaphysics of a particular type of object in formal detail. Below we will do exactly that. This serves two purposes. First, it shows how the generativist's metaphysics can be precisely expressed. Second, it illuminates the essentialist framework itself: the generativist's metaphysics provide natural and well motivated illustrations of many delicate essentialist distinctions.

We adopt a Finean framework for expressing essence. When P is a predicate $\Box_P$ is the sentential

[21] An inexhaustive list: Correia 2000, 2006, 2012; Correia and Skiles 2019; Ditter 2020, n.d.; Fine 1994a,b, 1995b, 2015; Teitel 2019.

operator: "it is true in virtue of the items that are P that ...". Since we are working in a type-theoretic framework we also allow higher-order predicates.[22] For instance, if R is a binary relation between objects, then $\lambda x^{\langle e,e\rangle}.x = R$ is a predicate and so $\Box_{\lambda x^{\langle e,e\rangle}.x=R}$ is the sentential operator: "true in virtue of R". Since we officially treat functions as functional relations a candidate abstraction operation A is simply a certain higher-order relation, and so we can ask what is essential to A. We also allow $\Box$ to be indexed by a finite *list* of predicates; $\Box_{P_0,P_1,\ldots,P_n}$ is the operator, "true in virtue of the items that are P_0 together with the items that are $P_1,\ldots$, together with the items that are P_n". As usual, if P is a predicate of the form $\lambda u.(u = x \vee u = y \vee \ldots)$ we write $\Box_{x,y,\ldots}$ instead of $\Box_{\lambda u.(u=x \vee u=y \vee \ldots)}$. We extend this to the case where the subscript is a list, writing $\Box_{a,P,R}$ for $\Box_{\lambda x^e.x=a,\lambda x^{\langle e\rangle}x=P,\lambda x^{\langle e,e\rangle}.x=R}$.

So much for notation: what notion of essence do we have in mind? Throughout, we will work with a generalization to type-theory of Fine's consequential logic of essence (Fine 1995b, 2000b). Assume that for each type τ we have divided the expressions of that type into logical and non-logical expressions. Essence is closed under logical consequence in the sense that if $\Box_P\phi$ and ψ is a consequence of ϕ and any non-logical expressions in ψ occur also in ϕ or P then $\Box_P\psi$. Call this F-consequence. To see how this works, consider that Socrates is essentially a man (formally: $\Box_s Ms$). He is thus also essentially a man or such that something is self-identical—$\Box_s(Ms \vee \exists x x = x)$. But even though his being a man or a city-dweller is a logical consequence of his being a man, he is not essentially a man or a city-dweller. (Formally: $\neg\Box_s(Ms \vee Cs)$.) The reason is that the non-logical predicate "being a city-dweller" does not occur in Ms.[23]

[22]The importance of asking about the nature of higher-order entities is stressed in Correia 2006; Correia and Skiles 2019. The formalization we adopt here follows the pioneering Ditter n.d.

[23]It is also possible to formulate the generativist's theory using a notion of constitutive essence—for instance, by using the framework of Ditter n.d. However, working with constitutive essence is notationally cumbersome, and the important metaphysical committments of generativism can be adequately formulated using just consequential essence.

Let $\bigwedge$ be the predicate $\lambda p^{\langle\rangle}.p \neq p$. Then $\Box_{\bigwedge}$ is the operator, "true regardless of the nature of any propositions" or "true in virtue of just logic" (cf. Fine 1995b, 246). If ϕ is a logical truth formulated in purely logical vocabulary, then $\Box_{\bigwedge}\phi$; but if ϕ is a logical truth containing non-logical vocabulary this is not so. What we, however, should expect is this: if ϕ is a logical truth, then $\Box_{\phi}\phi$. Each logical truth ϕ is true in virtue of the nature of the proposition ϕ. More generally, for each type τ let $\mathrm{Log}^{\langle\tau\rangle}$ be the predicate $\lambda x^{\tau}.\Box_{\bigwedge}Ex$. Log holds of an item of type τ if it is true in virtue of just logic that the item exists. Intuitively, Log holds of exactly the purely logical items of type τ.

2.4 Defining the abstraction operation

What do the application, identity, and existence conditions for the structural abstraction operation look like?

The application condition for the abstraction operation A tells us, first, that whenever $S = \langle D_S, R_S\rangle$ is a system and an object a is in D_S then $A(S,a)$ is defined. Secondly, the application condition tells us that it is essential to the values of A that they are values of A. That is, if $p = A(S,a)$ then it is essential to p that it can be abstracted from *some* object in some system—that is, $\Box_p \exists T \exists b A(T,b) = p$. Thirdly, if $p = A(S,a)$ this is no accident; it rather lies in the nature of p together with S,a that p can be abstracted from S,a—that is, $\Box_{p,S,a} A(S,a) = p$. (Note, however, that it is *not* essential to p that it can be abstracted from S and a in particular.) Putting all this together—using $\forall S$ as a quantifier restricted to (rigid, small) systems—the application condition is:

(AC) a. $\forall S \forall a (S a \rightarrow \exists p A(S,a) = p)$

b. $\forall S \forall a \forall p (A(S,a) = p \rightarrow \Box_p \exists T \exists b A(T,b) = p)$

c. $\forall S \forall a \forall p (A(S,a) = p \rightarrow \Box_{p,a,S} A(S,a) = p)$

The equivalence relation associated with the structural abstraction operation is the relation $\approx_A$ that holds between (S,a) and (T,b) iff there is an isomorphism $f\colon S \cong T$ such that $fa = b$.[24] The identity-condition is then:[25]

(IC) $\forall S \forall a \forall T \forall b (A(S,a) = A(T,b) \leftrightarrow (S,a) \approx_A (T,b))$

The existence condition states that a pure position p exists if and only if there is some system S and an element a of S such that $A(S,a) = p$ and each of S,a and A exists.

(EC) $\forall S \forall a \forall p (p = A(S,a) \rightarrow \Box_p (Ep \leftrightarrow \exists T \exists b (Tb \wedge A(T,b) = p \wedge E(T) \wedge E(b) \wedge E(A))))$

We will understand this as making the grounding claim that the existence of p is *grounded* in the existence of S,a, and the existence of A itself. Since the points we make in this paper do not turn on the precise formulation of the ground-theoretic existence condition we relegate a precise statement to this footnote.[26]

The generativist now holds that the abstraction operation A is implicitly defined by its satisfying

[24]In higher-order terms: $\approx_A = \lambda S \lambda a \lambda T \lambda b. \exists f (f\colon S \cong T \wedge fa = b)$.

[25](IC) is the precise version of the informal principle (SA) above. Apart from inessential notational differences this is the principle that Linnebo and Pettigrew (2014, 274–278) call "Frege Abstraction"; see also Linnebo 2008, 76.

[26] Use $<$ as a sentential operator for strict full ground, and $\leq$ as a sentential operator for distributive weak full ground. (For explanation of these notions see, e.g., Fine 2012a,b.) We take $<$ to have type $\langle [\langle\rangle], \langle\rangle \rangle$—that is, it is a relation between pluralities of propositions and propositions; and we take $\leq$ to have type $\langle [\langle\rangle], [\langle\rangle] \rangle$—a relation between pluralities of propositions. The existence condition is then:

(GEC) $\forall S \forall a \forall p (p = A(S,a) \rightarrow \forall \Gamma (\Gamma < Ep \leftrightarrow \exists T \exists b (Tb \wedge A(T,b) = p \wedge \Gamma \leq [E(T), E(b), E(A)])))$

What (GEC) ensures is that any grounds for the existence of a pure position p has to ground the existence of an object and a system from which p can be abstracted.

the above application, identity, and existence conditions. That is, the abstraction operation A is defined by:

(EA) $\quad$ (AC) $\wedge$ (IC) $\wedge$ (EC)

Within the essentialist framework we can express this precisely as follows. First, we need so say that it is essential to A that (EA) holds: that is, we have $\Box_A$(EA). By itself this is not enough: this leaves it open that there is more to the essence of A than just satisfying (EA). But just like in the case of the *direction-of* operation (§ 2.1 above) we want the nature of A to be exhausted by its satisfying (EA). To ensure this we require an "elimination" principle, saying that if ϕ is also true in virtue of the nature of A, then ϕ is a logical consequence of (AC), (IC), and (EC). We express this schematically as follows. If ϕ is a sentence, let $((\text{EA}) \to \phi)^A$ be the result of replacing each non-logical expression in (EA) $\to \phi$ with a distinct variable and taking its universal closure. The claim that the nature of A is exhausted by its satisfying (EA) is then expressed by the following schema:

$$\Box_A \phi \to \Box_{\bigwedge}((\text{EA}) \ \to \phi)^A$$

3 Generativist Structuralism

Non-eliminative structuralists hold that given any isomorphism class of (small) systems $\mathcal{S}$ there is a canonical representative $|S|$ of that isomorphism class—what we programmatically call the *pure structure* corresponding to $\mathcal{S}$. But non-eliminative structuralists have not developed a systematic metaphysics of such pure structures. The generativist offers an account of pure structures that—we will argue—gives non-eliminative structuralists what they want.

3.1 Pure structures generativist style

We begin by defining the *pure positions*. For any system $S = \langle D_S, R_S \rangle$, let $|D|_S$ be the plurality such that $|D|_S\, y$ iff $y = A(S, b)$, for some b in D_S. It is easy to show that if $S_0 \cong S_1$ then $|D_0|_{S_0} = |D_1|_{S_1}$.[27] $|D|_S$ is the plurality of pure positions generated by A from the system S. We next have to define a "pure" relation $|R|_S$ on $|D|_S$ such that the system $\langle |D|_S, |R|_S \rangle$ is isomorphic to S. But—as stressed by Hellman (2007, 546)—there is an embarrassing multitude of relations between which to choose. For let $|D|_S$ be the pure positions abstracted from D_S for some system $S = \langle D_s, R_S \rangle$. Suppose $|S| = \langle |D|_S, |R| \rangle$ is a candidate for being the pure structure and let π be a permutation of the pure positions $|D|_S$. Let $|R|^\pi$ be the relation that holds between the pure positions p, q iff $|R|$ holds between $\pi(p), \pi(q)$ and let $|S|^\pi = \langle |D|, |R|^\pi \rangle$. It is routine to see that $|S|$ is isomorphic to $|S|^\pi$. Moreover, $|S|$ and $|S|^\pi$ have the same pure positions: why should we choose $|R|$ over $|R|^\pi$ as our pure relation?

A natural first response is that not all R are such that $|S| = \langle |D|_S, R \rangle$ satisfy[28]

[27] This will follow from the more general Theorem 4.2 or see Linnebo and Pettigrew 2014.

[28] The apt term "Auto-Abstraction" is taken from Zanetti 2020 where it is credited to Linnebo. (Auto-Abstraction) is the generativist's way of expressing the idea that in the pure structure the positions occupy the roles they represent; in the terminology of Linnebo and Pettigrew 2014, 282 "places-as-offices" coincide with "places-as objects".

The perceptive reader will have noted that (Auto-Abstraction) together with (GEC) yields counterexamples to the orthodox view that partial ground is irreflexive. In fact, the problem does not just arise for structural abstraction: Donaldson (2017) noted that the same phenomenon arises for Hume's Principle understood as a grounding claim. However, these counterexamples to irreflexivity are of a comparatively benign kind. Since the existence of a pure position is also grounded in the existence of objects that are not pure positions, the self-grounding is not of that most problematic kind where a fact is its own sole ground. The literature on the "puzzles of ground" following Fine 2010a contains numerous proposals that can be extended to deal with the problems caused by (Auto-Abstraction). (See e.g., Woods 2018, Krämer 2013, Correia 2014, Litland 2015, 2020, (Auto-Abstraction) For every p in $|D|_S$, we have $A(|S|, p) = p$.

However, while we want the pure relation to satisfy (Auto-Abstraction) it is not uniquely characterized by satisfying it. Natural essentialist assumptions—to which the generativist is committed—ensure that there are many distinct relations Q such that $\langle |D|_S, Q\rangle$ is isomorphic to S and satisfies (Auto-Abstraction). For let Q be one relation such that $\langle |D|_S, Q\rangle$ is isomorphic to S. Let Q_d be the relation that holds between the pure positions p_0, p_1 iff Q holds between them and the emperor Diocletian is self-identical. Since different propositions are true in virtue of the nature of Q than in virtue of the nature of Q_d, Q and Q_d are distinct relations.

There is, however, a natural candidate for the relation $|R|_S$. To define this relation we need the assumption that for every (small) system S there is purely logical property L_S (in the sense of § 2.3 above) such that L_S applies to all and only the systems isomorphic to S.[29] We define the pure relation $|R|_S$ as follows

(Pure Relation)

$$|R|_S = \lambda xy.\exists T \exists u \exists v(L_S(T) \wedge x = A(T,u) \wedge y = A(T,v) \wedge R_T uv)$$

In words, for the relation $|R|_S$ to hold between two pure positions p_0, p_1 is for there to be a system $T = \langle D_T, R_T\rangle$ such that T is of the isomorphism type of S and such that $p_0 = A(T,a)$ and $p_1 = A(T,b)$ and R_T holds between a and b.

For each system $S = \langle D,R\rangle$ the system $|S| = \langle |D|_S, |R|_S\rangle$ is isomorphic to S (see Theorem 4.2).

deRosset 2021, Lovett 2020. This is not the place to adjudicate between these competing views.

[29] The assumption that L_S is logical is justified by the Tarski-Sher account of the logical operations, see e.g., McGee 1996. The property L_S is definable in $L_{\infty,\infty}$—for details see Button and Walsh 2018, 399-408.

We claim that $|S|$ is an excellent candidate for being the privileged representative—the pure structure—of the isomorphism class of S. Before we can make that case we need to get clearer on the nature of $|S|$ and how $|S|$ differs from other proposed candidates for being the pure structure. We begin with three preliminary points.

First, for the generativist pure positions are objects in just the same way as mere elements are; moreover, the pure positions have their properties in exactly the same way as mere elements have theirs. This sets the generativist apart from the view of Nodelman and Zalta (2014) who hold that pure positions do not instantiate the structural properties had by the elements—they merely *encode* them.

Second, the generativist has an advantage in that she does not just offer a theory of *sui generis* objects made for non-eliminative structuralism; for the generativist, pure positions are just one type of generated object amongst many. This sets the view apart from the theory of *ante rem* structures in Shapiro 1997 and the theory of pure graphs in Leitgeb 2020, 2021.

Third, and relatedly, like any other generated objects, the existence of the pure positions is *grounded* in the existence of the systems from which they can be abstracted. This, too, sets generativism apart from Shapiro's *ante rem* structuralism: for the *ante rem* structuralist the pure positions exist independently of the systems from which they can be abstracted.

However, the most striking differences between generativism and other accounts of the pure structures have to do with ontological dependence and purity.

3.2 Ontological dependence and purity in the rigid setting

Many structuralists have promulgated striking dependence claims, claiming that the pure positions depend on each other or that they depend on the pure structures. Here is Shapiro seeming to endorse both claims.

> The number 2 is no more and no less than the second position in the natural number structure; and 6 is the sixth position. Neither of them has any independence from the structure in which they are positions, and as positions in this structure, neither number is independent of the other. (Shapiro 2000, 258)

We use the essentialist framework of § 2.3 to evaluate these claims.[30]

Say that x *rigidly* depends on y if $y \neq x$ and y is a non-logical constituent of a proposition true in virtue of the nature of x; formally: $x \neq y \wedge \neg\text{Log}(y) \wedge \exists P \Box_x Py$.[31] Say that x *generically* depends on the Fs if F is non-logical and it is true in virtue of the nature of x that there is an F; formally, $\neg\text{Log}(F) \wedge \Box_x \exists y F y$.[32]

What, in these senses, do the pure positions depend on? Let p be a pure position and let $\mathcal{S}$ be the property of being a system from which p can be abstracted. (That is, $\mathcal{S} = \lambda T.\exists b A(T, b) = p$.) Since $\Box_p \exists T \mathcal{S}(T)$ it follows that p generically depends on the systems from which p can be abstracted. Consider now the property $\lambda z.\exists S\, \exists x(A(S, x) = p \wedge p \neq z \wedge \exists y A(S, y) = z)$ of being a pure position distinct from p that can be abstracted from any system from which p can be abstracted; call this property P.[33] We then have that $\Box_p \exists x P x$, and so p depends generically on the pure positions that

[30] For some other discussions of these claims see Linnebo 2008; Thompson 2016; Wigglesworth 2018.

[31] This follows Fine's suggestion that x depends on y if "y is a constituent of a proposition that is true in virtue of the identity of x or, alternatively, if y is a constituent of an essential property of x" (Fine 1995a, 275).

[32] For more on definitions along these lines see Fine 1994b, 1995a,b and Correia 2005. The attempt at defining dependence in terms of essence is not without its critics (see e.g., Wilson 2020; Koslicki 2012) but even if one thinks that what is defined by the above does not deserve the honorific "dependence" or if one thinks that other notions also deserve that honorific, the above two definitions clearly pick out metaphysically interesting notions.

[33] We are setting aside the case of the pure structures whose domain consist of a single pure

can be abstracted from the same systems as p. This also shows that the pure positions weakly depend on each other in the sense of Linnebo 2008, 78. According to Linnebo two entities x, y are said to *weakly depend* on each other if any way of generating x must make use of entities that suffice to generate y. In this sense, two pure positions p_0 and p_1 in $|D|_S$ do depend on each other: if $A(S,a) = p_0$, there is bound to be some object b in S such that $A(S,b) = p_1$.

It should also be clear that a pure position p rigidly depends on no objects: while it is essential to a pure position p that it be the result of abstracting from *some* system and *some* object in that system, there is no system and no object in that system such that it is essential to p that it be abstracted from *that* system and *that* object. The pure structure, on the other hand, depends on the pure positions. For if p is in $|D|_S$, then it is essential $|D|_S$ that p is in it, and it is essential to $|S|$ that its domain is $|D|_S$.

This does not mean that there is nothing on which the pure positions rigidly depend: they all depend on the abstraction operation itself.[34] In contrast, the abstraction operation A does not rigidly depend on the pure positions. The essence of A is exhausted by satisfying (EA) and since no non-logical constants (of any type) except A occur in (EA), no proposition containing a pure position will be true in virtue of the nature of A. The abstraction operation, as the saying goes, "knows nothing" about *particular* pure positions. Note, in particular, that the equivalence relation $\approx_A$ does not rigidly depend on any pure positions[35] Observe also that while it is essential to A, S, a taken together that there is a unique x such that $x = A(S,a)$—formally, $\Box_{A,S,a}\exists!xA(S,a) = x$—this is a quantificational claim, and is not *about* the particular object which is that unique x.

position.

[34]Though see footnote 51 below for some worries about this.

[35]To foreshadow: the situation will be more complicated in the non-rigid setting—see § 4.3 below.

Under what condition does the abstraction operation itself exist? The following view is natural. Abstraction operations are the kinds of items the nature of which it is to exist. By successfully laying down the application, identity, and existence conditions for the abstraction operation we have thereby ensured that $\Box_A E(A)$. The "successfully" does work here: as the literature on the "bad company" objection shows, what it takes for a putative definition of an abstraction operation to be successful is not at all a trivial matter.[36] The generativist only claims that *if* the definition is succesful, then it is essential to the defined operation that it exists.[37]

Let us now consider the extent to which distinct pure positions have distinct essences. Let p_0, p_1 be two distinct pure positions in $|S|$. It is essential to p_0, p_1 that they are distinct.[38] This is a collective essence: it is essential to *them* that *they* be distinct. An interesting question is whether this collective essence of p_0, p_1 can be "factorized" into a logical consequence of the essence of p_0 together with the essence of p_1. Let us remove any suspense: factorization holds in the rigid but fails in the non-rigid setting.

For any rigid system $S = \langle D, R\rangle$ and any object a in D, there is a logical predicate L^a_S such that L^a_S holds between an object b and a system $T = \langle D_T, R_T\rangle$ iff $D_T b$ and there is an isomorphism

[36]For a survey of various proposed constraints on abstraction principles, see Linnebo 2011 and Cook and Linnebo 2018. (Note, though, that their concern is with abstraction principles in the form of Hume's Principle.)

[37]This formulation just tells us that it lies in the nature of A to exist. That does not yet tell us what *grounds* the existence of A. Here there are several natural views. We could take inspiration from Dasgupta 2014 and hold that the existence of A is grounded in $\Box_A E(A)$, and hold that $\Box_A E(A)$ itself is an "autonomous" fact, a fact not apt for being grounded. Or we could hold that $E(A)$ is zero-grounded in the sense of Fine 2012a, 47-48. Or maybe we should just say that $E(A)$ is ungrounded. Which way to go here is obviously important, but since the problem of what grounds facts that are essentially true arises for everyone the generativist has no special problem here.

[38]$\forall x \forall y(x \neq y \rightarrow \Box_{x,y} x \neq y)$ is a theorem of the logic of essence (Fine 1995b, 255-256).

$f\colon S \cong T$ with $f(a) = b$. Let p_0, p_1 be two distinct pure positions in $|D|_S$ and consider the logical predicates $L^{p_0}_{|S|}$ and $L^{p_1}_{|S|}$. It is essential to p_0 that it is the result of abstracting from any object and system that has $L^{p_0}_{|S|}$. Formally: $\Box_{p_0}\forall S\,\forall x((L^{p_0}_{|S|}(S,x) \rightarrow A(S,x) = p_0) \wedge \exists S\,\exists x L^{p_0}_{|S|}(S,x))$. It is essential to p_1 that it is the result of abstracting from any object and system that has $L^{p_1}_{|S|}$. Formally: $\Box_{p_1}\forall S\,\forall x((L^{p_1}_{|S|}(S,x) \rightarrow A(S,x) = p_1) \wedge \exists S\,\exists x L^{p_1}_{|S|}(S,x))$. It is also a logical truth[39] that there is no S and no b such that $L^{p_0}_{|S|}(S,b)$ and $L^{p_1}_{|S|}(S,b)$. Thus, $\Box_{\wedge}\forall S\,\forall x\forall y(L^{p_0}_{|S|}(S,x) \wedge L^{p_1}_{|S|}(S,y) \rightarrow x \neq y)$. Thus $p_0 \neq p_1$ is a logical consequence of what is essential to p_0 together with what is essential to p_1.

Summing up, the generativist's view about dependence is that the pure positions depend on each other weakly but not rigidly. The pure positions depend weakly on the pure structure (and it on them); however, the pure positions do not depend rigidly on the structure, rather it depends rigidly on them. Moreover, distinct pure positions have distinct essences. As for the claim that the pure positions have "no independence from the structure", the generativist emphatically rejects this. Once the pure positions have been generated they take on a life of their own and are free to enter into other systems, alone or together. This is an advantage for the generativist: it allows her to take at face value constructions that embed pure structures in larger systems.

In addition to the dependence claims mentioned above, many structuralists have been tempted by a Principle of Purity stating that the pure positions have only *structural* properties.[40] Such a principle of Purity is related to the position that the pure positions are "incomplete" in not having an internal composition or nature.[41] It is well-known that naïvely stated Purity is false—a given

[39] In an appropriate infinitary logic.

[40] By "property" we here mean a relation between a system and the elements of its domain. A property is *structural* iff for all systems $S = \langle D_S, R_S\rangle$ and $T = \langle D_T, R_T\rangle$, all isomorphisms $f\colon S \cong T$, and all a in D_S we have $P(S,a)$ iff $P(T, f(b))$.

[41] For a good discussion of the problems with such an incompleteness claim see Linnebo 2008, 62-66.

pure position might be the Pope's favorite abstract but this is not a structural property—or even contradictory—the property of having only structural properties is not a structural property.

An initial fallback position for the structuralist is to hold that the only essential properties of the pure positions are structural.[42] Indeed, Linnebo and Pettigrew (2014, 276-278) have shown how to develop this view rigorously. (They use "fundamental" instead of "essential", but this makes no difference for current purposes.) But this is too restrictive: the pure positions are essentially abstract, but being abstract is not a structural property.

An advantage of generativist structuralism is that we do not have to rely on intuitions to determine which properties the pure positions have essentially: an account of their essential properties falls directly out of general generativist principles. For according to the generativist the essential properties of a generated object will be determined by the application, identity, and existence conditions for the generative operation the value of which it is. So if $p = A(S, a)$ the essence of p is exhausted by p's being what can be abstracted from any (T, b) having the property L^p_S. The generativist thus defends:

(Purity$_G$) The only properties a pure position has essentially are properties definable (using the vocabulary of the logic of essence) from the abstraction operation A.

(Purity$_G$) comes as close as possible to capturing the idea that the only properties the pure positions have essentially are the structural properties.

[42] For a critical assessment of other approaches see Linnebo and Pettigrew 2014, 270–272.

3.3 The Identification Problem

$|S|$ has many of the features non-eliminative structuralists have typically taken the pure structures to have; indeed, since (Purity$_G$) comes as close as possible to the idea that the only essential properties of the pure positions are structural properties, the generativist arguably gives the non-eliminative structuralists as much they could hope for. However, someone who presses the Identification Problem may well respond by saying that all we have done is describe a new type of platonic object, and that merely doing this does nothing to solve the Identification Problem. Take, e.g., the class of ω-sequences. We have shown how this class contains a certain system $|N| = \langle |D|_N, |R|_N \rangle$—the pure ω-structure—that arises by abstraction from all the other systems in the class. But what privileges this structure over all other ω-sequences, what makes it *the* natural number structure? Solving the Caesar Problem—showing that this pure structure is distinct from the mere systems—does not, by itself, help. After all, we know why the von Neumann ordinals are distinct from any other ω-sequence: this hardly makes the von Neumann ordinals the privileged ω-sequence. And observing that the pure structures are, well, *pure* does not help either; this is just to restate that they are special platonic objects.

Qua generativist one could concede. The generativist's main contention is that abstract objects are the values of generative operations; if the pure ω-sequence $|N|$ is not singled out as *the* natural numbers, this does not refute generativism as such. And generativism has much to recommend it even if it does not solve the Identification Problem: it is a principled and general view of the nature of abstract objects; it (once collective abstraction has been developed) allows for the generation of indiscernibilia; it provides a solution to the problem of converse relations; and so on. However, the generativist should not concede. The feature that privileges the pure structures is that they essentially *represent* the systems isomorphic to them.

This representational claim has to be understood in the right way. By making an arbitrary choice of a candidate operation B, any system $S = \langle D_S, R_S \rangle$ can be made to represent the systems isomorphic to it. For let B be a candidate abstraction operation such that for all $T \cong S$ we have $D_S = \{p \colon \exists a(a \in D_T \land B(T,a) = p\}$. *Relative to B*, S represents the systems isomorphic to it in the following sense. An element $p \in D_S$ represents—relative to B—those (T,a) such that $B(T,a) = p$; and the holding of the relation R_S between the elements p_0, p_1 of S represents—relative to B—the holding of R_T beween those a_0, a_1 such that $B(T,a_0) = p_0$ and $B(T,a_1) = p_1$.

But such representation is relative to an arbitrary choice of an operation B. What would it take for a system to non-arbitrarily represent? The natural thought is that a system S non-arbitrarily represents the systems isomorphic to it when its nature is that it so represents. For a system S to represent non-arbitrarily we thus require the following:

(i) there is an operation B such that S essentially represents, relative to B, the systems isomorphic to it;

(ii) the essence of S is exhausted by its so representing relative to B.

A pure structure $|S|$ is determined by its pure positions together with its pure relation; to establish that $|S|$ meets conditions (i) and (ii) we need to investigate the natures of the pure positions $|D|_S$ and the pure relation $|R|_S$.

Starting with the pure positions, suppose $p = A(S,a)$. We have seen (§ 3.2) that there is a logical property $L^p_{|S|}$ such it is essential to p that for any (T,b) such that $L^p_{|S|}(T,b)$ we have $p = A(T,b)$. Thus it is essential to p that it represents, relative to A, the systems from which it can be abstracted. But the essence of p is also exhausted by this: there is nothing more to the essence of p than what follows from its being what can be abstracted from any (T,b) such that $L^p_{|S|}(T,b)$.

As for the pure relations, the nature of $|R|_S$ is exhausted by the fact that its holding between two pure positions p, q just is for p, q to represent some (T, a), (T, b) such that $R_T(a, b)$. (In fact, given standard assumptions in the logic of ground, $|R|_S$'s holding between p, q is grounded in R_T's holding between a, b.) In a slogan: the holding of the representative relation between the representatives represents the holding of the represented relations between the represented.

This shows that the essence of the pure structure $|S|$ is exhausted by its representing, relative to A, the systems isomorphic to it. But why think that $|S|$ is the only system that, with respect to some operation B, has these features? It is true that once one looks at alternative representative systems one finds that they have richer natures. The von Neumann ordinals, for instance, are more than mere representatives of the entries of ω-sequences—the ordinals are *sets*, have *members*, and so on; moreover, for the membership relation $\in$ to hold between two ordinals is not just for those two ordinals to represent prior and posterior elements in ω-sequences. One would, however, like to have a general argument that $|S|$ is the only system with features (i) and (ii). Unfortunately, it is not clear how to develop such an argument; fortunately, there are further conditions on non-arbitrary representation that only the pure structures meet.

Recall that the generativist's general view is that abstract objects are the values of generative operations, and that generated objects exist because something from which they can be generated exists. In particular, the pure structure $|S|$ exists because some mere system from which it can be abstracted exists. We thus have the following condition on non-arbitrary representation:[43]

(iii) It is essential to the system S that it exists because some system that it represents relative

[43]For the generativist this is strictly speaking not a *further* condition. Since A is defined by its satisfying (EA) it is essential to A that any p which is a value of A exists iff (and because) some (S, a) such that $A(S, a) = p$ exists. It then follows—by a "Chaining" principle (see e.g., Fine 1995b, 247–249)—that a pure position exists because something from which it can be abstracted exists.

to B exists.

The final condition on non-arbitrary representation concerns the operation B itself.

(iv) There is no more to the nature of the operation B than is required for conditions (i), (ii), and (iii) to hold.

We can now show that the pure structure is the unique system that non-arbitrarily represents the systems isomorphic to it. According to the generativist the operation A is uniquely defined by (EA); and by inspection of (EA) one sees that what (EA) requires of the abstraction operation A is exactly what is needed for the abstraction operation and the pure structures defined from it to meet conditions (i), (ii), and (iii). Meeting condition (iv) thus amounts to be defined by (EA). It then follows from the generativist's general metaphysics that $|S|$ is the unique system that non-arbitrarily represents the systems isomorphic to it. This, then, is the generativist's solution to the Identification Problem: the pure structure is privileged by being the unique system the essence of which is exhausted by its representing the systems isomorphic to it.

Finally, we should point out that the essentialist framework is—no pun intended—essential for this solution to the Identification Problem. If we could only draw modal distinctions there would be no way of distinguishing between those systems that merely necessarily can be taken to represent the systems isomorphic to them, and those the essence of which is exhausted by their so representing.[44]

[44] It should be stressed that this is not the first attempt at using the representational role of the pure structures to solve the Identification Problem. Fine (1998, 629-931) suggests that we may use the theory of variable (or arbitrary) objects to give an account of pure structures. Since variable objects may reasonably be taken essentially to represent their values this also yields an account where the pure positions essentially represent their values. While a full comparison between Fine's

4 Collective Abstraction

However nice the generativist's account is in the rigid case, it cannot work in the non-rigid case. Consider the complex numbers $\mathbb{C}$. Since there are automorphisms $f\colon \mathbb{C} \cong \mathbb{C}$ with $f(i) = -i$, (IC) yields the absurd result that $A(\mathbb{C}, i) = A(\mathbb{C}, -i)$: we cannot generate *two* roots for -1.

4.1 Collections of abstraction operations

The problem with the generativist's account of the abstraction operation in the non-rigid case is the "the": with non-rigid systems there are many distinct abstraction operations that take us from systems to their pure structures. To see this it is useful to work through a simple case: the graph with two vertices and no edges. In addition to the various "mere" graphs with exactly two vertices and no edges the generativist wants there to be a pure graph $|G| = \langle |V|, |R| \rangle$ with two pure vertices p, q and no edges.

Let us consider how the pure graph $|G|$ may arise by abstraction from two distinct mere graphs G_0 and G_1. G_0 has vertices a_0, b_0 and G_1 has vertices a_1, b_1. p, q can be abstracted from G_0 and G_1 in four ways.[45] There is the way A_0 that generates p from both (G_0, a_0) and (G_1, a_1);[46] there is the

theory of variable objects and the theory of abstraction is beyond the scope of this paper, let us just note that a Finean account of pure structures will—as Fine (1998, 631) is fully aware—result in fairly rich natures for the pure positions: for instance, a Finean pure position will essentially have the non-structural property of being a dependent variable object.

[45]We are for now ignoring the fact that the pure positions also can be generated from themselves; taking that into account gives us 8 ways of abstracting the pure positions; we return to the problem this poses in § 4.3.

[46]It follows from this that G_0 has to generate q from (G_0, b_0) and (G_1, b_1). These details will be suppressed in what follows.

way A_1 that generates q from (G_0, a_0) and (G_1, a_1); there is the way A_2 that generates p from (G_0, a_0) and (G_1, b_1); finally, there is the way A_3 that generates q from (G_0, a_0) and (G_1, b_1). These ways of generating p, q from G_0, G_1 are depicted graphically in figure 1. The (dotted) double lines indicate which elements are to be mapped to the same pure position; the (dotted) single lines indicate *which* pure position they are to be mapped to.

<Figure 1 here>

What the generativist must do is give up the idea that for p to be a pure position is for p to be a value of *the* abstraction operation A. Indeed, given that the pure positions p, q are indiscernible, there *cannot* be a unique abstraction operation A: if there were, p and q would be discernible since one (but not the other) would be identical to $A(G_0, a_0)$. Rather, the generativist should hold that for p to be a pure position is for there to be *some* abstraction operation B such that p is a value of B. The generativist's task is thus to define what it is to be an abstraction operation. As we will put it, the task is to define the property $\mathcal{A}$ of being an abstraction operation.

In defining $\mathcal{A}$ the generativist faces a challenge other platonist positions do not face. Suppose, e.g., we were *ante rem* structuralists in the mould of Shapiro 1997. Then there would just exist a pure graph $|G|$ with pure positions p, q.[47] We could define $\mathcal{A}$ by first defining the abstraction operations in the obvious way: A_0 is the abstraction operation that maps G_0, a_0 and G_1, a_1 to p, A_1 is the abstraction operation that maps G_0, a_0 and G_1, a_1 to q; and so on. The property $\mathcal{A}$ of being an abstraction operation would then just be the property that applies to each operation so defined. But

[47]Recently, Leitgeb (2020, 2021) has shown how to develop a theory of pure graphs that allows indiscernible vertices and edges.

the generativist cannot say this. For the generativist the pure positions are (really) *defined* by being values of some abstraction operation. But if we define $\mathscr{A}$ in terms of the abstraction operations, and the abstraction operations in terms of the pure positions, we have a definitional circle.

4.2 Profiles and bundles

The key to defining the property $\mathscr{A}$ is to lay down analogues of the principles (AC), (IC), and (EC) discussed in § 2.4 above; that is, we must find collective application, identity, and existence conditions. In the main text we state how this can be done, illustrating features of the account using the above trivial two-element graph. A fully precise statement of the theory as well as proofs of relevant results are relegated to the supplementary appendix. <INSERT LINK TO APPENDIX HERE>

It turns out that the collective application condition—in particular finding an analogue of (ACa)—poses the greatest challenge. We begin by observing that if A is an abstraction operation defined on some systems[48] $[S_i : i \in I]$ then A induces an equivalence relation $\approx_A$ on pairs (S_i, a_i) of systems and elements of those systems as follows: $(S_i, a_i) \approx_A (S_j, a_j)$ iff $A(S_i, a_i) = A(S_j, a_j)$. We call $\approx_A$ the *abstraction profile* of A. To illustrate, consider the abstraction operations A_0, A_1, A_2, and A_3 introduced above. The operations A_0, A_1 have the same profile, depicted by the dashed double line in figure 1; this profile differs from the profile of A_2 and A_3 which is depicted by the double line in figure 1.

The application condition for $\mathscr{A}$ must ensure that for every abstraction profile there is an abstraction operation that essentially has that profile; and that for every abstraction operation

[48]For reasons that will become clear in § 4.3 we allow abstraction operations that are defined on some (but not all systems).

there is a profile that the operation essentially has; and finally—since we allow operations that are not defined on every system—that every abstraction operation can be extended to an abstraction operation on any extension of its profile.

The first task is to define what it is to be an abstraction profile. In the rigid case the abstraction profile $\approx_A$ associated with an operation A has a purely logical definition: $\approx_A$ is the relation $\lambda S \lambda T \lambda a \lambda b. \exists f(f\colon S \cong T \wedge f(a) = b)$. In the non-rigid case abstraction profiles do not (in general) have pure definitions. Suppose, e.g., that in the graphs G_0, G_1 the vertices a_0, b_0 are two (qualitatively identical) socks and a_1, b_1 are two other (qualitatively identical) socks. Then there will no purely logical definition of the various abstraction profiles between G_0 and G_1.

But we can give a purely logical definition of what it is to be an abstraction profile. An *isomorphism thread* on a plurality of systems $\mathcal{S} = [S_i\colon i \in I]$ is a plurality $\mathcal{I} = [f_{ij}\colon i, j \in I]$ such that

(i) for all $i, j \in I$, if $S_i \cong S_j$ there is exactly one $f_{ij}\colon S_i \cong S_j$ such that f_{ij} is in $\mathcal{I}$;

(ii) for all $i, j \in J$, if $S_i \cong S_j$ we have $f_{ji} \circ f_{ij} = \mathrm{id}_i$;

(iii) for all $i, j, k \in J$, $f_{ik} = f_{jk} \circ f_{ij}$.

If $\mathcal{I}$ is an isomorphism thread we write $\mathcal{I}(f)$ to mean that f is an isomorphism in $\mathcal{I}$. If $\mathcal{I}$ is an isomorphism thread on $[S_i\colon i \in I]$ define the relation $\approx_{\mathcal{I}}$ by setting $(S_i, a_i) \approx_{\mathcal{I}} (S_j, a_j)$ iff the unique $f_{ij}\colon S_i \cong S_j$ in $\mathcal{I}$ has $f_{ij}(a_i) = a_j$. An abstraction profile is then a relation $\approx$ such that $\approx$ is coextensive with $\approx_{\mathcal{I}}$ for some isomorphism thread $\mathcal{I}$.

Let us write Profile($\approx$) for the claim that $\approx$ is an abstraction profile; and Prof($B, \approx$) for the claim that the operation B has profile $\approx$. If B is an operation we write dom(B) for the plurality of systems on which B is defined. If B, C are two operations we say that C *extends* B if for all S, a such that

$B(S,a)$ is defined $C(S,a)$ is also defined and $B(S,a) = C(S,a)$; if C extends B we write $B \subseteq C$. We then require the following:

(Profile) a. $\forall \approx(\mathrm{Profile}(\approx) \rightarrow \exists B(\mathcal{A}(B) \wedge \Box_B \mathrm{Prof}(B,\approx)))$

b. $\forall B(\mathcal{A}(B) \rightarrow \exists \approx(\mathrm{Profile}(\approx) \wedge \Box_B \mathrm{Prof}(B,\approx)))$

c. $\forall B \forall \approx \forall \approx'(\mathcal{A}(B) \wedge \mathrm{Prof}(B,\approx) \wedge \approx \subseteq \approx' \rightarrow$

$\exists C(\mathcal{A}(C) \wedge \mathrm{Prof}(C,\approx') \wedge B \subseteq C))$

As we already can see from the case of the graphs with two vertices, fixing an abstraction profile does not determine a unique abstraction operation: while the abstraction operations A_0, A_1 agree that a_0, a_1 should be mapped to the same pure position they disagree about *which* of p, q they should be mapped to. The application condition has to describe *all* the abstraction operations with a given profile. The key to this is observing that distinct abstraction operations can be naturally transformed into each other; we illustrate the idea by showing how A_0 can be transformed into A_1.

There are two automorphisms on G_0—the identity automorphism id_0 and the automorphism π_0 that maps a_0 to b_0 (and *vice versa*); similarly, G_1 has the identity automorphism id_1 and the automorphism π_1 that maps a_1 to b_1 (and *vice versa*). If we define $(\pi_0,\pi_1)A_0$ by $(\pi_0,\pi_1)A_0(G_i,c) = A_0(G_i,\pi_i(c))$ then the following simple calculations show that $(\pi_0,\pi_1)A_0$ is A_1.

$$(\pi_0,\pi_1)A_0(G_0,a_0) = A_0(G_0,\pi_0(a_0)) = A_0(G_0,b_0) = A_1(G_0,a_0)$$

$$(\pi_0,\pi_1)A_0(G_0,b_0) = A_0(G_0,\pi_0(b_0)) = A_0(G_0,a_0) = A_1(G_0,b_0)$$

Generalizing, we make the following definition.

Definition 4.1. Let $\mathcal{S}$ be a collection of systems.

(i) An *automorphism bundle* on $\mathcal{S}$ is a plurality μ of automorphisms on systems in $\mathcal{S}$ such that for each S in $\mathcal{S}$ there is a unique automorphism $\mu_S : S \cong S$ in μ.

(ii) If O is an operation defined on $\mathcal{S}$ and μ is an automorphism bundle on $\mathcal{S}$, let μO be the operation defined by $\mu O(S,a) = O(S,\mu_S(a))$.

The application condition has to ensure that the abstraction operations are closed under the application of automorphism bundles; and that for any two abstraction operations with the same domain there is a bundle such that it is essential to the two operations that they can be transformed into each other via the bundle. This gives us:

(Closure) a. $\forall B \forall \mu(\mathcal{A}(B) \to \exists C(\mathcal{A}(C) \land \Box_{B,C,\mu} C = \mu B))$

b. $\forall B \forall C(\mathcal{A}(B) \land \mathcal{A}(C) \land \mathrm{dom}(B) = \mathrm{dom}(C) \to \exists \mu \Box_{C,B} C = \mu(B))$

The conjunction of (Profile) and (Closure) is the collective analogue to (ACa). In the rigid case (ACb) and (ACc) express that the application operation A is generative; in the collective case it is the *property* of being an abstraction operation that is generative. It is essential to the values of any abstraction operation that they are the values of *some* abstraction operation on some input; and if p is the value of the abstraction operation B on input S,a then it is essential to S,a,p taken together that for some abstraction operation C we have $C(S,a) = p$.

We thus end up with the following Collective Application Principle:

(CAC) a. (Profile) $\land$ (Closure)

b. $\Box_{\mathcal{A}} \forall B \forall S \forall x \forall y(\mathcal{A}(B) \land B(S,x) = y \to \Box_y \exists C \exists T \exists z(\mathcal{A}(C) \land C(T,z) = y))$

c. $\Box_{\mathcal{A}} \forall B \forall S \forall x \forall y(\mathcal{A}(B) \land B(S,x) = y \to \Box_{S,x,y} \exists C(\mathcal{A}(C) \land C(S,x) = y))$

In the rigid case there is a unique abstraction operation and its profile has a purely logical definition. Since neither condition is met in the non-rigid case the collective identity criterion is more involved. Observe first that if B is an abstraction operation it has some profile $\approx_B$. This gives us an *intra*operational identity criterion: $B(S_i, a)$ is identical to $B(S_j, b)$ iff $S_i \cong S_j$ and $(S_i, a) \approx_B (S_j, b)$. To obtain an *inter*operational identity criterion observe that if B, C are two abstraction operations it follows by (Profilec) that we can take them to be defined on the same domain. By (Closureb) there is an automorphism bundle μ such that $C = \mu B$. Thus $B(S_i, a) = C(S_j, b)$ iff $B(S_i, a) = B(S_j, \mu_j(b))$, where $\mu_j \colon S_j \cong S_j$ is the unique automorphism on S_j contained in μ. Putting all this together we arrive at the Collective Identity Criterion—(CIC) for short:

(CIC) $\forall B \forall C \forall S_0 \forall S_1 \forall a_0 \forall a_1 (\mathcal{A}(B) \wedge \mathcal{A}(C) \rightarrow$

$(B(S_0, a_0) = C(S_1, a_1) \leftrightarrow \exists \approx \exists \mu (\text{Prof}(B, \approx) \wedge C = \mu B \wedge (S_0, a_0) \approx (S_1, \mu_{S_1}(a_1)))))$

The collective analogue of the existence condition (EC) is the following:

(CEC) $\forall B \forall S \forall a \forall p (\mathcal{A}(B) \wedge p = B(S, a) \rightarrow$

$\Box_p (Ep \leftrightarrow \exists T \exists b \exists C (\mathcal{A}(C) \wedge C(T, b) = p \wedge E(T) \wedge E(b) \wedge E(C))))$

In words: a pure position p exists when a system T, an object b, and an abstraction operation C such that $C(T, b) = p$ all exist. As in the rigid case it is very natural to take this existence condition to be a claim about the grounds for the existence of the pure positions.[49]

The generativist claims that the essence of the property of being an abstraction operation is

[49] With the same conventions as in footnote 26 here is the precise formulation:

(CECG) $\forall B \forall S \forall a \forall p (\mathcal{A}(B) \wedge p = B(S, a) \rightarrow \forall \Gamma (\Gamma < Ep \leftrightarrow \exists T \exists b \exists C (Tb \wedge \mathcal{A}(C) \wedge C(T, b) = p \wedge$ $\Gamma \leq [E(T), E(b), E(C)])))$

exhausted by (CIC), (CAC), and (CEC). We thus lay down both

$$\Box_{\mathcal{A}}((\text{CAC}) \wedge (\text{CEC}) \wedge (\text{CIC}))$$

and the scheme:

$$\Box_{\mathcal{A}}\phi \to \Box_{\bigwedge}((\text{CAC}) \wedge (\text{CEC}) \wedge (\text{CIC}) \to \phi)^{\mathcal{A}}$$

Just as in the rigid case, it is important that (CAC), (CIC), and (CEC) are taken to exhaust the nature of $\mathcal{A}$: while there may be many properties of operations that satisfy these principles, the generativist's claim is that there is but one the essence of which is exhausted by them.[50]

Having thus defined the property of being an abstraction operation the generativist is in a position to define the pure positions and the pure structure. For any system $S = \langle D_S, R_S \rangle$ let the plurality of pure positions $|D|_S$ be defined as follows. $|D|_S\, x$ iff $\exists B \exists T \exists a (T \cong S \wedge \mathcal{A}(B) \wedge B(S, a) = x)$. Even in the non-rigid setting there is still, for every system S, a logical predicate L_S, such that L_S applies to all and only the systems isomorphic to S. We can thus define the pure relation $|R|_S$ in the same way as in the rigid case:

(CPR) $\quad |R|_S = \lambda xy. \exists B \exists T \exists u \exists v (L_S(T) \wedge \mathcal{A}(B) \wedge x = B(T, u) \wedge y = B(T, v) \wedge R_T uv)$

We define the pure structure as $|S| = \langle |D|_S, |R|_S \rangle$. Let B be an abstraction operation and S a system on which B is defined. Define B_S by setting $B_S(a) = B(S, a)$, for each $a \in D_S$. We can establish the following.

Theorem 4.2. *(i) If $S \cong T$ then $|D|_S = |D|_T$.*

[50] If one is doubtful about this one could adopt the analogue for the non-rigid case of the view suggested in footnote 11.

(ii) B_S *is one-to-one and onto.*

(iii) For all abstraction operations B and all systems S on which B is defined $B_S : S \cong |S|$.

Proof: The proof is relegated to the supplementary appendix. [INSERT LINK] □

4.3 Dependence in the non-rigid setting

In the rigid case we saw that each pure position can be made sense of individually. The situation is more complicated in the non-rigid case. To begin, the facts about what the pure positions depend on are almost as in the rigid case; the only difference is that the pure positions only generically depend on the abstraction operations.[51] The situation with the abstraction operations themselves is more interesting.

It is natural to take the property $\mathcal{A}$ of being an abstraction operation to essentially exist; the more challenging question is what to say about the existence of the abstraction operations themselves. It follows from (Profile) that the abstraction operations rigidly depend on their abstraction profiles; this suggests that we can take an abstraction operation to exist when and because its abstraction profile exists. What, then, grounds the existence of an abstraction profile?

The answer is that there is no uniform account. Abstraction profiles are simply certain relations between system-object pairs. Since there is no uniform account of the grounds for the existence of

[51] For the sake of uniformity maybe we should say this in the rigid case as well? We might be misled into thinking otherwise by the fact that in the rigid case there is a unique abstraction operaton. But we cannot conclude from the fact that it is essential to p that there is a unique abstraction operation that has p as its value that it is essential to p that it is the value of *that* operation. (Compare: it is essential to {Socrates} that there is a unique set that has just {Socrates} as a member. (Why? {Socrates} is essentially a set, and it is essential to *set* that everything has a unique singleton.) But it does not follow from this that it is essential to {Socrates} that it is a member of {{Socrates}}.)

relations we should not expect there to be a uniform account of the grounds for the existence of abstraction profiles.

To drive this point home consider first the plurality of systems $\mathcal{G}$ where for each $G = \langle D, R \rangle$ in $\mathcal{G}$ the domain D consists of two possible socks and R is the empty relation. In this case, we cannot define the abstraction profiles without mentioning the possible socks and so the abstraction profiles rigidly depend on the possible socks. It then seems natural to take the existence of the abstraction profiles to be grounded in the (possible) existence of the possible socks. Contrast this with the plurality of systems $\mathcal{G}'$, where for each $G' = \langle D', R' \rangle$ in $\mathcal{G}'$, the domain D consists of a pair of one possible human and one possible rock and R is the empty relation. Consider the abstraction profile $\approx$ that matches up humans with humans (and rocks with rocks). This abstraction profile only rigidly depends on the properties of being human and being a rock. It is natural to take the existence of this abstraction profile to be grounded in the existence of the properties of being human and being a rock.

There is a class of abstraction profiles that raise special problems. A *full* profile is a profile on a collection of systems $\mathcal{S}$ that contains the pure structure of every system it contains. To see the problem posed by these profiles consider the collection of systems $G_0, G_1, |G|$ where G_0, G_1 are our two element mere graphs and $|G|$ is the pure two element graph $|G| = \langle [p, q], |R| \rangle$. The profile $\approx_F$ determined by $(G_0, a_0) \approx_F (G_1, a_1) \approx_F (|G|, p)$ is full. What grounds the existence of $\approx_F$?[52]

Given that p and q are indiscernible it seems that the existence of the profile $\approx_F$ will have to be (partially) grounded in the existence of $|G|$, and in particular, in the existence of p (and q). Let A be an abstraction operation with profile $\approx_F$. Since the existence of A is grounded in the existence of

[52]Similar problems arise for dependence but since the considerations are similar we refrain from discussing this case.

$\approx_F$ it follows by the transitivity of ground that the existence of A is partly grounded in the existence of p. On pain of contradicting the asymmetry of ground the existence of p *cannot* then be grounded in the existence of A. Since $\approx_F$ was an arbitrary full profile and A was an arbitrary operation with profile $\approx_F$ it follows that it is impossible for the existence of the pure positions to be grounded in the existence of full abstraction operations.

It is for this reason that we allowed partial profiles and partial abstraction operations in (CAC). The patterns of ground can become quite intricate. Recall that A_0 is the abstraction operation defined by $A(G_0, a_0) = A_0(G_1, a_1) = p$; let $\approx_0$ be its profile. Let $\approx_F$ be the full profile introduced above and let A_F be the abstraction operation extending A_0 by $A_F(|G|, p) = p$. Then we have the following facts about (partial) ground and non-ground:

$$E(\approx_0) \prec E(A_0) \prec Ep \prec E(\approx_F) \prec E(A_F) \nprec E(p)$$

Turning finally to collective essences, we can use the pure graph $|G|$ to show that some collective essences are not factorizable. Let p, q be the pure positions of $|G|$. We then have $\Box_{p,q} p \neq q$. This collective essence cannot be factorized as a logical consequence of the essences of p, q respectively. For the essence of p is exhausted by being the value of some abstraction operation applied to a system that forms a trivial two element graph. But that is exactly the same as the essence of q. The distinctness of p and q therefore cannot be a logical consequence of their respective essences: their distinctness is "irreducibly collective".

In the rigid case we could make sense of each pure position on its own—the pure position p is the unique object that can be abstracted from any (T, b) such that $L_p(T, b)$. We can no longer—in general—make sense of the pure positions on their own. However, the pure positions need not be so entangled that they can only be made sense of all together. For each pure position p there is a logical property L_p and a unique plurality of pure positions pp with exactly the same essence as p. While we have to generate the positions pp simultaneously, the pp can be generated independently of any disjoint plurality of pure positions.

4.4 Haecceities

The *haecceity* of an object a is the property of being that very object. Many structuralists have held that indiscernible objects fail to have haecceities. An example:

> because haecceities are intrinsic to each individual, the permutation [=automorphism] of individuals always results in a new situation, whereas permuting exactly structurally similar individuals in a mathematical structure results in exactly the same structure. (Leitgeb and Ladyman 2008, 394)

As we understand haecceities here, the existence of haecceities is compatible with what Ladyman and Leitgeb say. In particular, we may accept that there is a generative operation H that takes any object a and gives us a property H_a such that it is essential to H_a that a and only a has H_a. Formally, $\Box_H \forall x \exists! P(H(x, P) \wedge \Box_P \forall y(Py \leftrightarrow x = y))$: haecceities in our sense are generated from the objects whose haecceities they are, and so they depend on them. Thus any automorphism of two indiscernible objects will induce an automorphism of their haecceities.

This is not just special pleading on behalf of haecceities: the phenomenon of indiscernible generated and dependent objects is widespread. Set-formation, e.g., is a generative operation and it can be applied to indiscernibles like i, and $-i$ too. The unique member of $\{i\}$ is i; and $-i$ is the unique member of $\{-i\}$. But this fact does not show that i and $-i$ are, after all, discernible; rather, since i and $-i$ are indiscernible, this just shows that $\{i\}$ and $\{-i\}$ are themselves indiscernible.[53]

[53] Wigglesworth (2021) has proposed that we can carry out abstraction on non-rigid systems by first "rigidifying" them by adding, for each object a in the domain of the system, the property of being that object. If the above comments are correct it is far from clear that this can be done in

We thus see that allowing indiscernible objects to have haecceities does not restore factorizability. What we have in the rigid case is that for every pure position p there is a property E_p such that (i) E_p depends only on the abstraction operation A; (ii) it is essential to E_p that it applies to exactly one object; and (iii) it is essential to p that E_p applies to it. The haecceity of i fails the first condition since it depends on i.

4.5 The identification problem in the non-rigid setting

As in the rigid case (§ 3.3) the generativist's proposed solution to the Identification Problem turns on the representational role of the pure structures. Given that there is no longer a unique abstraction operation the claims have to be reformulated as follows. Let $S = \langle D_S, R_S \rangle$ be a system and let $\mathcal{B}$ be a property of operations such that for all $T \cong S$ and all C such that $\mathcal{B}(C)$ we have $D_S = \{p \colon \exists a(a \in D_T \wedge C(T, a) = p)\}$. The system S represents, *relative to* $\mathcal{B}$, the systems isomorphic to it in the following sense. An element $p \in D_S$ represents—relative to each C such that $\mathcal{B}(C)$—those (T, a) such that $C(T, a) = p$. And the holding of the relation R_S between the elements p_0, p_1 of S represents—relative to $\mathcal{B}$—the holding of R_T beween those a_0, a_1 such that $C(T, a_0) = p_0$ and $C(T, a_1) = p_1$, for each C such that $\mathcal{B}(C)$.

The conditions (i)–(iv) on non-arbitrary representation are as before except that we now relativize to a property of $\mathcal{B}$ of operations. It follows from the generativist's general commitments that the property $\mathcal{A}$ defined by (CIC), (CAC), and (CEC), and the pure structures defined from $\mathcal{A}$ are the only property of operations and the only systems meeting (the collective versions of) conditions (i)–(iv). In the non-rigid case, too, the generativist solution to the Identification Problem is that what

a way that makes the objects discernible. Due to considerations of space, we cannot discuss his proposal further here.

privileges the pure structures is that they are the only systems the essences of which are exhausted by their representing the systems isomorphic to them.

5 (Meta)Semantic Issues

While the generativist's view is not a view about meaning, in the rigid case it is naturally paired with one. We can introduce an operator "A", where the meaning of "A" is implicitly defined by holding that the sentence (EA) be true. The generativist then takes terms like "$A(S, a)$" to be canonical names[54] of the objects they denote. The generativist holds that we can come to understand terms like "$A(S, a)$" and come to understand that they refer without having independent access to the referents of these terms.[55] Let us assume that such an account is defensible in the rigid case. (Arguing for it will take us too far afield here.) Since there is no unique abstraction operation, this account cannot work in the non-rigid case; however, an analogous view is available.

We illustrate it using the pure graph with two elements. For each mere graph $G = \langle V, R\rangle$ and each object a in V, the generativist introduces a term $|a|_G$ that is meant to refer to a result of abstracting on a (in G). Such terms refer *relative to abstraction operations*; the reference of $|a|_G$ relative to A is, naturally, just $A(G, a)$. A term has *determinate reference* if it has the same referent relative to every abstraction operation.

[54]In this paper, expressions like "A" are officially relations. But we can amend the theory and let expressions like "A" be functors.

[55]The semantic and metaphysical takes on the abstraction principle are orthogonal. One could hold that the *operations* are really implicitly defined by the abstraction principle, while holding that it is no part of the meaning of the operator that the principle holds for it. Conversely, one could hold that it is *analytic* of the operator that the principle hold, but deny that this tells us anything about the nature of the *operation*. (This position is laid out with care in Rosen 2003.)

We may go on to define truth relative to an abstraction operation. The important case is that an identity sentence $t = s$ is true relative to A if the reference of s relative to A is the same as the reference of t with respect to A. (We prescind from going into the rest of the theory; the details are fairly obvious.) A sentence is determinately true (false) if it is true (false) relative to every abstraction operation.

Consider the graph G_0 with two vertices a_0, b_0 and no edges introduced above. Neither $|a_0|_{G_0}$ nor $|b_0|_{G_0}$ determinately refers. However, since any abstraction operation A will map a_0 and b_0 to distinct pure positions the terms $|a_0|_{G_0}$ and $|b_0|_{G_0}$ determinately have distinct referents. Thus the sentence $|a_0|_{G_0} \neq |b_0|_{G_0}$ is determinately true. If, on the other hand, we consider terms that are drawn from different graphs the resulting identity statements are no longer guaranteed to be either determinately true or determinately false. For instance, since there are abstraction operations that map a_0 (in G_0) and a_1 (in G_1) to the same pure position and also operations that map them to distinct objects, the identity-statement $|a_0|_{G_0} = |a_1|_{G_1}$ is neither true nor false.

The account sketched here is, of course, formally a supervaluationist account. (Any admissible assignment of referents to the terms $|a_0|_{G_0}$ and $|b_0|_{G_0}$ will have to assign them different referents, but there are admissible assignments that assign $|a_0|_{G_0}$ and $|a_1|_{G_1}$ the same referent, and some that do not.) But the generativist's version of supervaluationism is distinctive in that she can explain where the admissible assignments come from (she has a metasemantics for "admissible assignment"): an assignment v of referents to the terms $|a_0|_{G_0}$ and $|b_1|_{G_1}$ is admissible iff there is an abstraction operation such that $v(|a_0|_{G_0}) = A(G_0, a_0)$ and $v(|b_1|_{G_1}) = A(G_1, b_1)$.

This view about how terms like $|a_0|_{G_0}$ function has similarities with the semantic relationalist view of Fine 2007. In Fine's framework one could think of the terms $|a_0|_G$ and $|b_0|_{G_1}$ as free variables such that: (i) they can each take as their value either one of the pure positions p_0, p_1; (ii) they are *anti-coordinated* in that whatever value $|a_0|_G$ takes, $|b_0|_G$ has to take the other value. What the generativist provides is an account—a metasemantics—of how the patterns of coordination come to be: the patterns of coordination are exactly those that are induced by the abstraction operations.[56]

Unlike in the rigid case, the generativist does not give an account of how we *refer* to the individual pure positions: there is nothing we can do to refer to one of $i, -i$ as opposed to the other. But there is still a sense in which we can refer to the pure positions. For suppose we have a plural term forming operator "[ ]" that takes a list of terms $t_0, t_1, \ldots$ and gives us a plural term $[t_0, t_1, \ldots]$. The reference of a plural term $[t_0, t_1, \ldots,]$ relative to A is, naturally, the plurality of the referents of $t_0, t_1, \ldots$ relative to A. This has an interesting consequence: while it is indeterminate whether $|a_0|_{G_0} = |a_1|_{G_1}$ and it is indeterminate whether $|b_0|_{G_0} = |b_1|_{G_1}$ the plural identity statement $[|a_0|_{G0}, |b_0|_{G_0}] = [|a_1|_{G_1}, |b_1|_{G_1}]$ is determinately true.

This may offer some solace to neo-Fregeans. They have often held that to be an object is to be a possible referent of a singular term.[57] Non-rigid structures seem to pose a fatal problem to this view—as argued forcefully by Brandom (1996) (see also Shapiro 2008, 291). But the possibility of determinate plural reference shows that this might be premature: while it is true that it is impossible to fashion a term that determinately refers to one of the roots of -1, it is possible to fashion a plural term that determinately refers to the roots of -1. One might perhaps say that while to be is not to be a possible referent of a term, to be is to be some possible referents of some terms.

[56] The present account may also be of use for Horwich (2005). Terms like "i" and "$-i$" pose problems for a use theory of meaning. For how could the (semantically relevant) use of "i" be different from the use of "$-i$"? In response to this Horwich suggested that while "i", "$-i$" individually do not have a definite use, they *collectively* have a definite use. One may take the present account to flesh out this idea.

[57] See amongst many others Dummett 1991; Wright 1983; Hale and Wright 2001.

6 An Application: Positionalism

The techniques behind collective abstraction have applications elsewhere in metaphysics, especially in cases where we need to account for the natures of indiscernible yet distinct objects. The techniques are applicable to the problem of accounting for the nature of (indiscernible) possibilia, giving an account of fictional objects, and giving an account of qualitativism. Because of its interest for the generativist we will here consider an application to the problem of converse relations.[58]

[58]It must be stressed that the techniques of collective abstraction are not applicable to every case involving indiscernible objects. A case in point might be structuralism about physics, in particular, the family of views that goes under the name of *Ontic Structural Realism* or OSR. While this is not the place to go into the details of this vast literature, four comments are in order. First, let us distinguish between moderate and radical OSR. Radicals like French and Ladyman (2003) and Ladyman (1998) hold that, fundamentally, there are only physical *structures*. Radicals either adopt an eliminativist stance according to which there are (properly speaking) no physical particles or else they adopt a reductionist stance according to which particles are logical constructions from the structures. Moderates—like Esfeld (2004) and Esfeld and Lam (2008)—hold that while there are fundamental physical particles they are mutually dependent on the physical structures of which they are constituents. Second, the generativist agrees with McKenzie (2014) that the moderate's dependence claims are best understood as claims about essence; where the present paper contributes something beyond McKenzie's is in the use of higher-order resources. Third, whether generativism gives a plausible account of OSR depends on whether one is a moderate. The generativist's is an account of, well, generated objects—objects the existence of which is grounded in the existence of that from which they are abstracted. But for the moderate: from what would the fundamental particles be abstractions? Moderate ORS is more akin to mathematical *ante rem* structuralism. Fourth, while generativism looks more promising as an account of radical OSR, the account of collective abstraction developed here is not immediately applicable. That account shows how we can generate pure positions from objects in systems. But presumably a defender of radical OSR would not want to take as basic such *object*-containing systems! So a generativism appropriate for

According to the standard view of relations a relation—e.g., the SHORTER-THAN relation—applies to its relata in an order. When SHORTER-THAN applies to Bob and Suzy in that order, we get the state of affairs that Bob is shorter than Suzy. If SHORTER-THAN applies in the opposite order we get the state of affairs that Suzy is shorter than Bob. On the standard view every non-symmetric relation has a distinct *converse*. For instance, the SHORTER-THAN relation has the TALLER-THAN relation as its converse.

The standard view—as pointed out by Fine (2000a)[59]—miscounts states of affairs. Intuitively, the state of affairs that Bob is shorter than Suzy is the very same state of affairs as the state that Suzy is taller than Bob. But how can this be if one state involves the relation SHORTER-THAN and the other the relation TALLER-THAN?[60]

For present purposes what is particularly important is that converse relations give rise to an Identification Problem. Is the pure structure of the natural numbers the one where the numbers are ordered by the pure LESS-THAN relation or rather the one where they are ordered by the pure GREATER-THAN relation? Related to this metaphysical worry there is—as Williamson (1985) pointed out—a metasemantic worry: in virtue of what could our expression "less-than" pick out one of the relations LESS-THAN and GREATER-THAN as opposed to the other?

OSR would have to be more radical, allowing us to generate objects and structures from systems that are not isomorphic to the resulting generated structures. Developing such an account of collective abstraction is obviously of considerable interest.

[59]For different problems see Williamson 1985 and Dorr 2004.

[60]Here we are, of course, relying on the principle that the states of affairs that result from completing distinct relations have to be distinct. Recently, many have taken the moral of the so-called Russell-Myhill paradox to be that this is misguided (see e.g.. J. Goodman 2017; Fairchild 2017; Dorr 2016; Uzquiano 2015). This is not the place to discuss this paradox; we just note that the problem of converse relations arises even once we have avoided the Russell-Myhill paradox—say by adopting a ramified theory of states of affairs.

The Positionalist proposes to get around these problems by changing our conception of how relations apply to their relata. Each relation *R* is equipped with a number of *positions* (or argument-places) and the relation applies to its relata relative to an assignment of the relata to those positions. To illustrate, the SHORTER-THAN relation is equipped with two positions—SHORT and TALL. Upon assignment of Bob to SHORT and Suzy to TALL we get the state-of-affairs that Bob is shorter than Suzy. On the positionalist view relations do not have converses; we thus do not have to distinguish the state of affairs that Bob is shorter than Suzy from the state of affairs that Suzy is taller than Bob.

Unfortunately, strictly symmetric relations present the positionalist with a seemingly fatal problem. A strictly symmetric relation is a relation *R* such that the state *Rab* is identical to the state *Rba*. Consider the state of Bob's being next to Suzy. This is the same state as Suzy's being next to Bob. (The NEXT-TO relation is clearly strictly symmetrical.) But how can the positionalist account for this? Since NEXT-TO is a binary relation, it has two positions—NEXT and NIXT. But there is a difference between assigning Bob to NEXT and Suzy to NIXT and doing the reverse. This leads the positionalist to hold that there are in fact *two* states corresponding to the sentence "Suzy is next to Bob"!

Initially, the positionalist can respond by holding that while the NEXT-TO relation has two positions, the positions are *indiscernible* in the sense that the same state results if we assign *a* to NIXT and *b* to NEXT as if we assign *b* to NIXT and *a* to NEXT (Leo 2008).[61] This is correct as far as it goes, but what *explains* that the same state results from assigning *a* to NEXT and *b* to NIXT as from assigning *b* to NEXT and *a* to NIXT?

[61] Dixon (2018) provides a different solution to the problem presented by NEXT-TO, but his solution does not deal with the problem posed by the relation *R* that holds between objects a, b, c, d if they are arranged in a circle in that order (Fine 2000a, p. 17n10). A different—and very interesting—suggestion is presented in Donnelly 2016. Due to considerations of space we cannot discuss it further here.

One view is that nothing explains this—it is simply a brute fact about the NEXT-TO relation. Leo (2008, 351-352) suggests a different view: the positions NEXT and NIXT are abstractions from the states involving the NEXT-TO relation. This view would have the advantages that we do not have to posit positions as fundamental entities and that, rather than taking the identity and distinctness facts involving positions to be brute, we could explain them in terms of how the positions are abstracted (cf. Fine 2000a, 16).

Unfortunately, Leo does not offer an account of abstraction vindicating this view. The theory of collective abstraction provides the vindication needed. Here is a sketch of the construction.

Let us take as given a collection of states. Each state S has a number of *occurrences* of objects, O_S, where we allow an object to have many occurrences in a state.[62] We take for granted that it makes sense to substitute objects for occurrences of objects in states. More precisely, a substitution is a function σ from occurrences of objects to objects. We write $\sigma(S)$ for the state that results from applying the substitution σ to S; if S and T are such that $T = \sigma(S)$ we say that S, T are *co-relational*. We assume that there is a bijection $f\colon O_S \to O_{\sigma(S)}$ such for all $a \in O_S$, $f(a)$ is an occurrence of $\sigma(a)$. We call such an f an *isomorphism* between S and $\sigma(S)$. If $\sigma(S) = S$ we say that f is an automorphism.[63]

We define the notion of a substitution profile exactly like the notion of an abstraction profile. That is, a *substitution profile* is a relation $\approx$ between state-occurrence pairs satisfying the conditions in Definition 4.1. An automorphism bundle is, as before, a collection of automorphisms with exactly one for each state. To abstract the positions of a relation we need to lay down an abstraction principle

[62] *Fundamentally*, this is not a positionalist view. At the fundamental level we have a collection of states with occurrences of objects. Positions are then generated by abstraction from these states.

[63] The assumption that there is a *bijective* f elides some complications involving *coalescence* of occurrences; see Leo 2010, 147-148, 168.

for the property $\mathcal{A}$ that applies to exactly the abstraction operations that takes the occurrences of objects in a state S and gives us a position in the relation R of which S is a completion. The principle will be exactly like it is for the case of abstraction operations.

So take the state of the Bob's being next to Suzy. The present approach ensures that there are two abstraction operations and two positions NEXT and NIXT in the NEXT-TO relation. One abstraction operation maps Bob to NEXT and Suzy to NIXT; the other operation does the opposite. Structurally, this is the same situation as with the trivial graph of two elements. We thus have an account of how we can generate two distinct yet indiscernible positions in the NEXT-TO relation.

7 Closing

What we may call the generativist program holds that abstract objects are the values of generative operations. In this paper we have furthered this program in several ways. We have made the notion of a generative operation precise using a higher-order logic of essence; we have showed how we can treat structural abstraction for rigid systems as a generative operation; we defined pure structures and pure positions and developed a non-eliminative structuralist solution to Benacerraf's Identification Problem. The main contribution of the paper, however, lies in the development of the method of collective abstraction. What distinguishes this method is that instead of defining a single abstraction operation we simultaneously define a collection of operations. This allows us to generate indiscernible objects by structural abstraction on non-rigid systems thereby solving a long-standing problem for non-eliminative structuralists. By using the essentialist framework we investigated what pure positions and structures depend on; strikingly, the dependence facts differ between rigid and non-rigid structures, with non-rigid structures showing a higher degree of “entanglement”. We also showed how the account of collective abstraction allows us to develop an account of reference to indiscernible pure positions. Finally, we applied the account of collective abstraction to the metaphysics of relations, developing a novel version of positionalism.

References

Antonelli, Aldo G. (2010). “Notions of Invariance for Abstraction Principles”. In: *Philosophia Mathematica* 18.3, 276–292.

Benacerraf, Paul (1965). “What numbers could not be”. In: *Philosophical Review* 74, 47–73.

Boccuni, Francesca and Jack Woods (2020). “Structuralist Neologicism”. In: *Philosophia Mathematica* 28.3, 296–316.

Brandom, Robert B. (1996). “The significance of complex numbers for Frege’s philosophy of mathematics”. In: *Proceedings of the Aristotelian Society* 96, 293–315.

Burgess, John P (1999). “Book Review: Stewart Shapiro. Philosophy of Mathematics: Structure and Ontology”. In: *Notre Dame Journal of Formal Logic* 40.2.

Button, T. and Sean Walsh (2018). *Model Theory and Philosophy*. Oxford University Press.

Button, Tim (2006). “Realistic Structuralism’s Identity Crisis: A Hybrid Solution”. In: *Analysis* 66.3, 216–222.

Cook, Roy T. and Øystein Linnebo (2018). “Cardinality and Acceptable Abstraction”. In: *Notre Dame Journal of Formal Logic* 59.1, 61–74.

Correia, Fabrice (2000). “Propositional Logic of Essence”. In: *Journal of Philosophical Logic* 29.3, 295–313.

— (2005). *Existential dependence and cognate notions*. Philosophia Verlag.

Correia, Fabrice (2006). “Generic Essence, Objectual Essence, and Modality”. In: *Noûs* 40.4, 753–767.

— (2012). “On the Reduction of Necessity to Essence”. In: *Philosophy and Phenomenological Research* 84.3, 639–653.

— (2014). “Logical Grounds”. In: *Review of Symbolic Logic* 7.1, 31–59.

— (2017). “Real Definitions”. In: *Philosophical Issues* 27.1, 52–73.

Correia, Fabrice and Alexander Skiles (2019). “Grounding, Essence, and Identity”. In: *Philosophy and Phenomenological Research* 98.3, 642–670.

Dasgupta, Shamik (2014). “The Possibility of Physicalism”. In: *Journal of Philosophy* 111.9/10, 557–592.

— (2017). “Constitutive Explanation”. In: *Philosophical Issues* 27.1, 74–97.

deRosset, Louis (2021). “Hollow Truth”. In: *Philosophical Review* 130.4, 533–581.

Ditter, Andreas (2020). “The Reduction of Necessity to Essence”. In: *Mind* 129.514, 351–380.

— (n.d.). “Essence and Necessity”. In: *Journal of Philosophical Logic* ().

Dixon, T. Scott (2018). “Plural Slot Theory”. In: *Oxford Studies in Metaphysics* 11, 193–223.

Donaldson, Thomas (2017). “The Foundations of Arithmetic?” In: *Noûs* 50.4, 775–801.

Donnelly, Maureen (2016). “Positionalism Revisited”. In: *The Metaphysics of Relations*. Ed. by Anna Marmodoro and David Yates. Mind Association Occasional Series. Oxford University Press. Chap. 5, 80–99.

Dorr, Cian (2004). “Non-Symmetric Relations”. In: *Oxford Studies in Metaphysics* 1, 155–92.

— (2016). “To Be F Is To Be G”. In: *Philosophical Perspectives* 30. Ed. by John Hawthorne and Jason Turner, 39–134.

Dummett, Michael (1973). *Frege: Philosophy of Language*. London: Duckworth.

Dummett, Michael (1991). *Frege: Philosophy of Mathematics*. London: Duckworth.

Esfeld, Michael (2004). “Quantum Entanglement and a Metaphysics of Relations”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 35.4, 601–617.

Esfeld, Michael and Vincent Lam (2008). “Moderate Structural Realism About Space-Time”. In: *Synthese* 160.1, 27–46.

Fairchild, Maegan (2017). “A Paradox of Matter and Form”. In: *Thought: A Journal of Philosophy* 6.1, 33–42.

Fine, Kit (1991). “The study of ontology”. In: *Noûs* 25.3, 263–294.

— (1994a). “Essence and Modality”. In: *Philosophical Perspectives* 8, 1–16.

— (1994b). “Senses of Essence”. In: 53–73.

— (1995a). “Ontological Dependence”. In: *Proceedings of the Aristotelian Society*. New Series 95, 269–290.

— (1995b). “The Logic of Essence”. In: *Journal of Philosophical Logic* 24.3, 241–273.

— (1998). “Cantorian Abstraction: A Reconstruction and Defense”. In: *Journal of Philosophy* 95.12, 599–634.

— (2000a). “Neutral Relations”. In: *Philosophical Review* 109, 1–33.

— (2000b). “Semantics for the Logic of Essence”. In: *Journal of Philosophical Logic* 29.6, 543–584.

— (2002). *The Limits of Abstraction*. Oxford: Oxford University Press.

— (2005). “Class and Membership”. In: *The Journal of Philosophy* 102.11, 547–572.

— (2007). *Semantic Relationism*. Blackwell Pub.

— (2010a). “Some Puzzles of Ground”. In: *Notre Dame J. Formal Logic* 51.1, 97–118.

Fine, Kit (2010b). “Towards a Theory of Part”. In: *Journal of Philosophy* 107, 559–589.

— (2012a). “Guide to Ground”. In: *Metaphysical Grounding*. Chap. 1, 37–80.

— (2012b). “The Pure Logic of Ground”. In: *The Review of Symbolic Logic* 5.1, 1–25.

— (2015). “Unified Foundations for Essence and Ground”. In: *Journal of the American Philosophical Association* 1.2, 296–311.

French, Steven and James Ladyman (2003). “Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure”. In: *Synthese* 136.1, 31–56.

Goodman, Jeremy (2017). “Reality is Not Structured”. In: *Analysis* 77.1, 43–53.

Goodman, Nelson (1958). “On Relations That Generate”. In: *Philosophical Studies* 9.5-6, 65–66.

Hale, Bob and Crispin Wright (2001). *The Reason’s Proper Study*. Oxford: Clarendon Press.

Hellman, Geoffrey (2007). “Structuralism”. In: *The Oxford Handbook of Philosophy of Mathematics and Logic*. Ed. by Stewart Shapiro. Oxford University Press. Chap. 17, 536–562.

Horwich, Paul (2005). *Reflections on Meaning*. Oxford University Press, Clarendon Press ;

Keränen, Jukka (2001). “The Identity Problem for Realist Structuralism”. In: *Philosophia Mathematica* 9, 308–330.

Koslicki, Kathrin (2012). “Varieties of Ontological Dependence”. In: *Metaphysical Grounding: Understanding the Structure of Reality*. Ed. by Fabrice Correia and Benjamin Schnieder. Cambridge University Press, 186.

Krämer, Stephan (2013). “A Simpler Puzzle of Ground”. In: *Thought* 2.2, 85–89.

Ladyman, James (1998). “What is Structural Realism?” In: *Studies in History and Philosophy of Science Part A* 29.3, 409–424.

— (2005). “Mathematical Structuralism and the Identity of Indiscernibles”. In: *Analysis* 65.287, 218–221.

Leitgeb, Hannes (2020). “On Non-Eliminative Structuralism. Unlabeled Graphs as a Case Study, Part A”. In: *Philosophia Mathematica* 28.3, 317–346.

— (2021). “On Non-Eliminative Structuralism. Unlabeled Graphs as a Case Study, Part B”. In: *Philosophia Mathematica* 29.1, 64–87.

Leitgeb, Hannes and James Ladyman (2008). “Criteria of Identity and Structuralist Ontology”. In: *Philosophia Mathematica* 16.3, 388–396.

Leo, Joop (2008). “The Identity of Argument-Places”. In: *Review of Symbolic Logic* 1.3, 335–354.

— (2010). “Modeling Occurrences of Objects in Relations”. In: *Review of Symbolic Logic* 3.1, 145–174.

Lewis, David (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Linnebo, Øystein (2006). “Sets, properties, and unrestricted quantification”. In: *Absolute Generality*. Oxford University Press, 149–178.

— (2008). “Structuralism and the Notion of Dependence”. In: *The Philosophical Quarterly* 58.230, 59–79.

— (2009). “Bad Company Tamed”. In: *Synthese* 170.3, 371–391.

— (2011). “Some Criteria for Acceptable Abstraction”. In: *Notre Dame Journal of Formal Logic* 52.3, 331–338.

— (2018). *Thin Objects*. Oxford: Oxford University Press.

Linnebo, Øystein and Richard Pettigrew (2014). “Two Types of Abstraction for Structuralism”. In: *Philosophical Quarterly* 64.255, 267–283.

Litland, Jon Erling (2015). “Grounding, Explanation, and the Limit of Internality”. In: *Philosophical Review* 124.4, 481–532.

— (2020). “Prospects for a Theory of Decycling”. In: *Notre Dame J. Formal Logic* 61.3, 467–499.

Lovett, Adam (2020). “The Puzzles of Ground”. In: *Philosophical Studies* 177.9, 2541–2564.

McGee, V. (1996). “Logical operations”. In: *Journal of Philosophical Logic* 25.6, 567–580.

McKenzie, Kerry (2014). “Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis”. In: *British Journal for the Philosophy of Science* 65.2, 353–380.

Nodelman, Uri and Edward N. Zalta (2014). “Foundations for Mathematical Structuralism”. In: *Mind* 123.489, 39–78.

Normand, Scott (2018). “Criteria of Identity and the Hermeneutic Goal of Ante Rem Structuralism”. In: *Synthese* 195.5, 2141–2153.

Parsons, Charles (2004). “Structuralism and Metaphysics”. In: *Philosophical Quarterly* 54.214, 56–77.

Prior, Arthur N. (1971). *Objects of Thought*. Clarendon Press.

Rosen, Gideon (2003). “Platonism, Semiplatonism and the Caesar Problem”. In: *Philosophical Books* 44.3, 229–244.

— (2015). “Real Definition”. In: *Analytic Philosophy* 56, 189–209.

— (2020). “Abstract Objects”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2020. Metaphysics Research Lab, Stanford University.

Shapiro, Stewart (1997). *Philosophy of mathematics: structure and ontology*. Oxford: Oxford University Press.

— (2000). *Thinking About Mathematics*. Oxford: Oxford University Press.

— (2008). “Identity, Indiscernibility, and ante rem Structuralism: The Tale of i and -i”. In: *Philosophia Mathematica* 16.3, 285–309.

Teitel, Trevor (2019). “Contingent Existence and the Reduction of Modality to Essence”. In: *Mind* 128.509, 39–68.

Thompson, Naomi (2016). “Metaphysical Interdependence”. In: *Reality Making*. Ed. by Mark Jago. Oxford: Oxford University Press, 38–56.

Uzquiano, Gabriel (2015). “A Neglected Resolution of Russell’s Paradox of Propositions”. In: *Review of Symbolic Logic* 8.2, 328–344.

Wigglesworth, John (2018). “Grounding in Mathematical Structuralism”. In: *Reality and Its Structure*. Ed. by Graham Priest and Ricki Bliss. Oxford University Press. Chap. 11, 227–236.

— (2021). “Non-Eliminative Structuralism, Fregean Abstraction, and Non-Rigid Structures”. In: *Erkenntnis* 86.1, 113–127.

Williamson, Timothy (1985). “Converse Relations”. In: *The Philosophical Review* 94.2, 249–262.

— (2003). “Everything”. In: *Philosophical Perspectives* 17, 415–465.

— (2013). *Modal Logic as Metaphysics*. Oxford University Press.

Wilson, Jessica (2020). “Essence and Dependence”. In: *Metaphysics, Meaning, and Modality: Themes from Kit Fine*. Ed. by Mircea Dumitru. Oxford University Press. Chap. 14, 283–300.

Woods, Jack (2018). “Emptying a Paradox of Ground”. In: *Journal of Philosophical Logic* 47.4, 631–638.

Wright, Crispin (1983). *Frege’s Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press/Humanities Press.

— (2020). “Replies to Part I: Frege and Logicism”. In: *Language, Logic, and Mathematics*. Ed. by Alexander Miller. Oxford University Press, 279–253.

Yablo, Stephen and Gideon Rosen (2020). “Solving the Caesar Problem—with Metaphysics”. In: *Logic, Language and Mathematics. Themes from the Philosophy of Crispin Wright*. Ed. by Alexander Miller. Oxford University Press. Chap. 5, 116–134.

Zanetti, Luca (2020). “Grounding and Auto-Abstraction”. In: *Synthese* 198.11, 10187–10205.