Two Forms of Functional Reductionism in Physics

Lorenzo Lorenzetti
University of Bristol

December 15, 2022

Abstract

Functional reductionism characterises inter-theoretic reduction as the recovery of the upper-level behaviour described by the reduced theory in terms of the lower-level reducing theory. For instance, finding a statistical mechanical realiser that plays the functional role of thermodynamic entropy allows to establish a reductive link between thermodynamics and statistical mechanics. This view constitutes a unique approach to reduction that enjoys a number of positive features, but has received limited attention in the philosophy of science.

This paper aims to clarify the meaning of functional reductionism in science, with a focus on physics, to define both its place with respect to other approaches to reduction and its connection to ontology. To do so, we develop and explore two alternative frameworks for functional reduction, called *Syntactic Functional Reductionism* and *Semantic Functional Reductionism*, that expand and improve the basic functional reductionist approach along different lines, and make clear how the approach works in practice. The former elaborates on David Lewis' account, is connected with the syntactic view of theories, is committed to a logical characterisation of functional roles, and is embedded within Nagelian reductionism. The latter adopts a semantic view of theories, spells out functional roles mainly in terms of mathematical roles within the models of theories, and is expressed in terms of the related structuralist approach to reduction. The development of these frameworks has the final goal of advancing functional reductionism, to make it a fully-fledged alternative account for reduction in science.

Keywords: Reductionism; Functionalism; Thermodynamics; Statistical physics

Contents

1 Introduction

Functionalism is all about understanding things in terms of the roles they play. According to this view, theoretical terms are defined by the roles they have in theories, and properties are cashed out in terms of their causal roles or behaviour. Functional reductionism exploits functionalism to shed light on inter-level relations: finding a lower-level realiser for an upper-level functional role gives us a way to connect the two levels.[1].

Functional reductionism is a view with a venerable tradition in the philosophy of mind (e.g. Lewis (1972), Kim (1998, 2005), Morris (2020)), where it has been employed to relate phenomenal and mental states. If pain is that state "that tends to be caused by bodily injury, to produce the belief that something is wrong with the body [...]" (Levin (2021)) and so on, and we individuate a brain state that fills those roles, we can functionally reduce pain to that specific kind of physical state.

This account is growing in importance within the philosophy of science as well, especially in the philosophy of physics (e.g. Esfeld and Sachse (2007), Lam and Wüthrich (2018, 2020), Huggett and Wüthrich (2021), Butterfield and Gomes (2020a, 2020b), Robertson (2020), Lorenzetti (2022)).[2] In this context, functional reduction is primarily used to model theoretical reduction between scientific theories and represents a unique approach to reduction. It has been used for instance to model reductive relationships between thermodynamics and statistical mechanics, between classical and quantum mechanics, and between general relativity and quantum gravity theories. According to functional reductionism, the primary aim of reduction is to find the right lower-level realisers for the upper-level behaviour: reduction is secured if we find in the bottom-level theory some theoretical elements that play the functional roles described by the upper-level theory. For instance, let's say we can functionally define 'temperature' in terms of its role within thermodynamics, and we find out that 'mean kinetic energy' plays the role of temperature: in that case, we can functionally reduce temperature to mean

[1]Following Rosaler (2015, 2019) we interpret the high-level/low-level distinction in very loose terms. Instead of high and low levels we can for instance talk about the domain of the reduced theory as opposed to the domain of the reducing theory, or we can talk about broader-scoped and narrower-scoped theories. In this sense, our discussion of reduction is able to capture also cases of inter-theoretic reduction between what are arguably same-level theories, such as general relativity and Newtonian mechanics. Nothing we say about reduction presupposes the idea that reality is ordered in a hierarchy of levels

[2]The value of functionalism in physics has been also recently defended e.g. by Knox (2019) and Wallace (2012, 2021). Functional reductionism has been discussed also by Albert (2015). Moreover, Lewis' and Kim's models quoted above apply to science as well.

kinetic energy, and this can be regarded as a step in the reduction of thermodynamics to statistical mechanics.

The aim of this paper is to advance the literature on functional reductionism in science by making clear how this approach to reduction works, elaborating on the way in which reduction is exactly achieved according to the view, on its relationship with other standard approaches to scientific reduction, and on its connection with ontology. We focus in this essay on functional reduction in physics, and we take as our starting point and case study an instance of functional reduction recently advanced by Robertson (2020), concerning the reduction of thermodynamic entropy to statistical mechanics.

In order to fully develop functional reductionism, we first elaborate on and review in detail the only full-fledged functional reductionist account in the literature, introduced by Lewis (1970) and recently defended by Butterfield and Gomes (2020a). This framework, which we call *Syntactic Functional Reductionism*, is a form of functional reduction embedded within the syntactic view of theories, is committed to a logical characterisation of functional roles, and is a functionalist form of Nagelian reduction. We apply the view to our case study and present some possible shortcomings of this approach. We thus introduce a novel alternative framework for functional reduction, called *Semantic Functional Reductionism*. It is based on the semantic view of theories, spells out functional roles mainly in terms of mathematical roles within the models of the theory, and is expressed in terms of the structuralist approach to reduction. We then show how this framework can account for our case study, and find that it can overcome the issues faced by the Syntactic Functional Reduction in that respect. The primary results of this paper are thus to pose some potential issues for the standard approach to functional reduction and to put forward a new way to explicate functional reductionism. Overall, however, we point out that both two views remain viable, as they each enjoy particular strengths and weaknesses that the following discussion will bring about.

The divide between the syntactic and the semantic view – according to which theories are, respectively, sets of sentences and family of models – is an important distinction in this context, because the stance about theories one combines with functional reductionism heavily influences how we formalise the notion of functional role, which has overarching consequences for the whole account. However, for the purpose of the paper, we do not require a clear-cut distinction between syntactic and semantic, rather we are just mainly interested in distinguishing Syntactic Functional Reductionism and Semantic Functional Reductionism as approaches to functional reduction respectively focused on logical sentences and mathematical models.[3]

The upshots of this paper are important for several reasons. **First**, this essay takes a crucial step toward the establishment of functional reductionism as a fully developed alternative account of inter-theoretic reduction in science, and will therefore have an impact on both the specific literature on functional reduction and the general literature on reductionism in science. In fact, it clarifies the debate on functional reductionism by providing and analyzing two clear alternative frameworks according to which we can

[3]We therefore follow Wallace (2021), who broadly interprets the syntactic view as 'language-first' and the semantic view as 'math-first'.

articulate the view. As mentioned, the discussion of the frameworks leads us to show how we can develop in more detail the notion of functional role in different ways, sheds light on the relationship between ontology and reduction within functional reductionism, and makes clear the connection between functional reductionism and the other standard approaches to reduction. **Second**, relatedly, by discussing the thermodynamics case study, we show how functional reduction works in practice, and we deepen further our understanding of the view by analyzing both the advantages and disadvantages of each framework and the ontological bearings of each specific approach. The Syntactic framework delivers an approach to ontology that is very clear but also very rigid, whereas the Semantic one allows for a more flexible view of the ontological aspects of reduction. These aspects mirror the situation at the theoretical level: while the former approach adopts a very rigorous and logically-formulated view on the formulation of theories, the latter embraces a model-based account. The choice between the two frameworks hinges also on these features. Most importantly, the lack of flexibility characterising the Syntactic approach is problematic insofar as we want to freely choose which elements of the theory we want to be realists about, and the logical translation aspect of the view can work against the framework as well, whereas the Semantic framework works better in these respects. The present paper thus not only develops two alternative takes on functional reduction, but also assesses them via the analysis of a realistic case study and allows us to provide a thorough evaluation of each alternative. **Third**, proposing an alternative to the extent Syntactic account of functional reduction is an important improvement for the whole functional reductionist approach. Indeed, someone could find endorsing functional reductionism problematic just due to contingent issues related to the specific Lewisian account, as that is currently the only complete framework for the view. Providing an alternative, represented here by Semantic Functional Reductionism, makes functional reductionism much more resistant to this kind of risk and makes functional reduction more palatable overall. **Fourth**, more generally, the essay is intended to have a broader impact on the whole debate on theoretic reductionism, as we show that functional reduction can integrate either Nagelian or structuralist reduction and provide a revised and improved version of these approaches, embedded in the functional reductionist framework. Syntactic Functional Reductionism is indeed a form of Nagelian reduction in which bridge laws are not postulated as additional assumptions, and are thus less problematic, while Semantic Functional Reductionism improves on the rather broad structuralist account of reduction, providing a more precise version of it and a stronger justification for the approach. Thus we don't simply clarify the place of functional reduction with respect to other accounts of reduction, but we also argue that the two forms of functional reduction presented here can be considered to be improved versions of, respectively, Nagelian and structuralist reduction.

Section 2 reviews Robertson's functional reductionist proposal concerning thermodynamics, which will be the starting point of our discussion. Having presented how functional reduction works for a real example, in Sections 3 and 4 we discuss the two functional reductionist frameworks. Section 5 overviews the overall pros and cons of each approach.

2 A Case Study for Functional Reductionism

We review here the instance of functional reduction recently put forward by Robertson (2020). Being a state-of-the-art example of scientific functional reduction, it provides a good introduction to the approach, and a well-suited case study for our discussion about the relationship between functional reduction and ontology in Sections 3.3 and 4.2.

Robertson's aim is to reduce the thermodynamic entropy S_{TD} to some statistical mechanic quantity, as a step in the reduction of thermodynamics to statistical mechanics – in particular, to reduce the second law of thermodynamics, which can be expressed in terms of the behaviour of the thermodynamic entropy.[4] To do so, her goal is to find in statistical mechanics a realiser for the role of the thermodynamic entropy. We report here just the essential details. Let's start from the top-level theory, and in particular from thermodynamic entropy S_{TD}. This is a function of the state of a thermodynamic system, like pressure, and it is roughly said to measure the 'disorder' of the system.[5] Or, using Clausius' definition, entropy can be defined as the thing that increases by Q/T whenever heat Q enters a system at temperature T.[6] We can thus represent the change of entropy dS_{TD} in a system as:

$$\frac{dQ}{T} = dS_{TD}, \tag{1}$$

where dQ is the change in heat (the heat absorbed) and T is the temperature. Thermodynamic entropy can then be represented by integrating (1). In this way we represent the entropy difference between two states of the system, in this case state 0 and state B:

$$\int_0^B \frac{dQ}{T} = S_{TD}(B). \tag{2}$$

This quantity is crucial for modelling thermodynamic behaviour, and thus reducing it to statistical mechanics would be an essential step to reducing thermodynamics to statistical mechanics, as we can use this to formulate the second law of thermodynamics. The two central characteristic features of S_{TD} on which reduction is focused are related to how this function works in two kinds of situations.

On the one hand, let's look at the case of arbitrary quasi-static reversible cycles in the equilibrium space Ξ. A thermodynamic equilibrium state is a state in which no macroscopic change occurs in a system, and the equilibrium space is the space of those states. A quasi-static reversible cycle is a process in which the system moves through equilibrium states, thanks to the fact that it is evolving slowly. For these processes we expect the following to occur:

$$\oint \frac{dQ}{T} = 0. \tag{3}$$

[4] The second law of thermodynamics, according to Clausius' statement, says that it is impossible for heat to spontaneously move from a lower-temperature reservoir to a higher-temperature reservoir.

[5] We consider thermally isolated systems and reversible processes.

[6] Cf. Schroeder (1999).

That is, if a process P is a quasi-static reversible process, we can write:

$$\Delta S_{TD} = 0. \tag{4}$$

On the other hand, it can be proven that, if a process P (say, between state A and state B) is not quasi-static, the thermodynamic entropy is a quantity that cannot decrease:

$$S_{TD}(B) - S_{TD}(A) \geq 0, \tag{5}$$

that is:

$$\Delta S_{TD} \geq 0. \tag{6}$$

Functional reduction consists in finding a statistical mechanics realiser – a statistical mechanics function or quantity – for the roles of S_{TD} which are mathematically specified by (4) and (6). Using words to express that theoretical functional role, we can say that we have to "Find a statistical mechanics realiser which, for thermally isolated systems, is increasing in non-quasi-static processes, but non-increasing in quasi-static processes, such as those represented by curves in Ξ." (Robertson (2020), 21).

Let's thus move to statistical mechanics.[7] A key concept in statistical mechanics is that of canonical ensemble, which is used to represent the possible states in which a system can be. In particular, the canonical ensemble gives the probability that a system is in a specific state n:

$$p(n) = \frac{e^{-E_n/k_B T}}{\sum_m e^{-E_m/k_B T}}, \tag{7}$$

where E is the energy of each state and k_B is Boltzmann's constant. To simplify, we can introduce a new notation, and write $\beta \equiv 1/k_B T$ and $Z = \sum_n e^{-\beta E_n}$. We can thus rewrite (7) as:

$$p(n) = \frac{e^{-\beta E_n}}{Z}. \tag{8}$$

Moving to quantum statistical mechanics, we write (8) in a slightly different way. In quantum mechanics, a system can be described via a density matrix ρ. We can thus express the canonical ensemble for a given system as:

$$\rho = \frac{e^{-\beta \hat{H}}}{Z}, \tag{9}$$

where $\hat{H}$ is the Hamiltonian operator representing the energy.[8] In both classical and quantum statistical mechanics, the canonical ensemble can be used to represent thermal

[7] We slightly deviate here from Robertson's presentation, for simplicity of exposition. See e.g. Tong (2012) for an introduction.

[8] If we make a measurement of the energy of the system then the probability of finding ourselves in an energy eigenstate $|\phi\rangle$ is $p(\phi) = \langle\phi| \rho |\phi\rangle$, and for energy eigenstates this is indeed just (8).

equilibrium. What matters for us now is that ρ is important to the introduction of a new quantity, the quantum Gibbs entropy, since the canonical ensemble is said to maximise Gibbs entropy S_G:

$$S_G = -k_B Tr \rho \ln \rho, \tag{10}$$

where Tr is the trace over the density matrix. Having introduced S_G, we shall now gloss over a lot of details and just report here how the reduction of thermodynamic entropy is achieved through a functional reduction of S_{TD} to S_G. Briefly put, Robertson (2020, sect. 6) shows that, for quasi-static processes in quantum statistical mechanics, we can write:

$$\Delta S_G = 0 \tag{11}$$

On the other hand, for non-quasi-static adiabatic processes, with $t_1 - t_0 \approx 0$, we can derive:[9]

$$S_G[\rho_{can}(t_1)] - S_G[\rho_{can}(t_0)] > 0. \tag{12}$$

The presentation so far provides what we asked for: we have found a statistical mechanical function – i.e. the statistical mechanical entropy S_G – that is constant in quasi-static processes (11) and that increases in rapid non-quasi-static processes (12). Indeed, the statistical mechanical equations (11) and (12) for S_G mathematically mirror the thermodynamic equations (4) and (6) embedding S_{TD}. The result is that these equations display the functional similarities shared by the two quantities.

Finally, to strengthen the functional correlation between the two quantities, Robertson shows that, in the right parameter regime, S_{TD} and S_G evolve in a very similar way. First, take (1), and derive the following from the first law of thermodynamics $dE = TdS - pdV$, where V is the volume:

$$dS_{TD} = \frac{1}{T_{TD}}(dE_{TD} + p_{TD}dV). \tag{13}$$

On the other hand, within Gibbsian quantum statistical mechanics, given certain assumptions and approximations, we can derive Gibbs entropy as:

$$dS_G = \frac{1}{T}(d\langle E\rangle + \langle p\rangle dV), \tag{14}$$

where the brackets denote that we are taking the average value. All in all, we can conclude that the Gibbs entropy functionally reduce the thermodynamical entropy:

> The Gibbs entropy can play the right role, since it increases in non-quasi-static processes but is constant in quasi-static processes. Furthermore, S_G is connected to heat in the right way (Robertson (2020), 31).

[9] Provided that we have adopted a new canonical ensemble tailored to the process and different from the starting one, and we have abstracted away certain details.

To recap, Robertson's goal was to find a statistical mechanical reductive basis that could reduce a specific thermodynamic behaviour, which is codified by the evolution of the thermodynamic entropy S_{TD}. To do so, she exploited the functionalist idea that, in order to reduce the thermodynamic entropy to statistical mechanics we have to find a statistical mechanical quantity which – at least approximately – plays the role of S_{TD} in the upper theory. The functionalist model of reduction thus tells us here what we have to do if we want to establish reduction, that is we have to focus on finding something in the low-level theory which instantiates the right patterns of behaviour within the high-level theory. In this way, functional reduction provides a clear and plausible model for reduction, that we can use to find a statistical mechanical underpinning for thermodynamics. Indeed, as stressed by Robertson, formulating a reductionist account for the second law of thermodynamics is a notoriously difficult task,[10] and functional reduction provides the tools to do so.[11]

In Sections 3 and 4 we develop two alternative frameworks for functional reduction, and this case study will be very illustrative to discuss how the frameworks work. We see that Syntactic Functional Reductionism faces the problem of shoehorning the mathematical formalism used here into a logical formulation and the problem of accounting for the approximation required for the reduction. On top of that, the way in which the framework reformulates this example of functional reduction prompts a very specific, but also too restrictive, account of the ontological implications of the reduction. On the contrary, the formulation of functional reductionism provided by the Semantic framework is model-based and mostly mathematically formulated, and thus accommodates in a more straightforward way the case study as presented here. Framed in that way, the account also allows for a more flexible account of the ontological meaning of the reduction at stake.

3 Syntactic Functional Reductionism

The first functional reductionist framework we introduce is called 'Syntactic Functional Reductionism'. It is based on the functional reductionist account first put forward by Lewis (1970) and recently defended and improved by Butterfield and Gomes (2020a), and is currently the most developed functional reductionist account available in the philosophy of science literature.[12] According to this approach, reduction goes as follows. The first step is to write down the laws of the reduced theory in terms of the reducing theory. At that point, by appealing to functionalism, we derive the bridge laws between the theoretical terms of the two theories from the laws of the bottom theory alone. We are thus able to derive the upper-level laws from the bottom-level laws plus bridge laws.

[10] See Callender (1999) on this topic.

[11] For further discussions about the advantages of functional reductionism see for instance Lam and Wüthrich (2018), Butterfield and Gomes (2020a, 2020b), Lorenzetti (2022).

[12] E.g. Huggett and Wüthrich (2021) discuss the Lewisian approach in the context of the functional reduction of spacetime, and Lorenzetti (2022) applies this framework to the case study of Wave Function Realism, developing a functional reductionist account relating classical and quantum mechanics.

Since law-derivation via bridge laws is the essence of Nagelian reduction, this approach can be considered a kind of Nagelian reduction. However, within this functionalist form of Nagelian reduction, bridge laws are derived from the reducing theory, and not added as extra postulates, like it is in the standard Nagelian view. Because of this feature, the account can be regarded as an improved version of Nagelian reduction.[13]

The aim of this section is to present the most developed version of the account possible and then provide an assessment of the framework. In 3.1 we introduce the basis of Syntactic Functional Reductionism, i.e. David Lewis's account, show its connection with Nagelian reduction and with the syntactic view, and describe the link between theoretical reduction and ontological reduction within the account. In 3.2 we delve further into the account, showing how the basis can be improved with respect to two aspects: dealing with approximation and moving to a more local kind of reduction. In 3.3 we apply the case study of Section 2 to Syntactic Functional Reductionism and raise some related issues for the view.

3.1 The Core: The Lewisian Basis

This subsection introduces in more detail the core of Syntactic Functional Reductionism, as defended by Lewis, Butterfield, and Gomes. Since the view is a kind of Nagelian reduction, it is important to first briefly introduce the latter account. According to Nagel's (1962) classic model of reduction, a theory T_P can be said to be reduced to another theory T_F iff the laws of T_P can be deduced from the laws of T_F plus some auxiliary assumptions. In the (common) case in which the two theories do not share their theoretical terms we need also to postulate bridge laws connecting the two vocabularies. For instance, in the context of the reduction of thermodynamics to statistical mechanics, we can derive the Boyle-Charles law from statistical mechanics' laws given a bridge law stating that 'temperature' means 'mean kinetic energy' (cf. Dizadji-Bahmani (2021)).

The Lewisian approach provides a special Nagelian account of reduction that builds on functional reduction in order to obtain the required bridge laws. According to this view, inter-theoretical reduction essentially proceeds in three steps:

1. We write down the laws of theory T in logical terms, then we replace all the theoretical terms $\tau_1 ... \tau_n$ of the theory with open variables $x_1 ... x_n$, leaving just non-theoretical terms and connectives, i.e. we move from $T(\tau_1, ... \tau_n)$ to $T(x_1, ... x_n)$. We now build the Ramsey sentence of the theory by placing an existential quantifier in front of the sentence: $\exists x_1, ... x_n T(x_1, ... x_n)$. This says that there are certain xs which realise the theory. On the assumption that the theory is uniquely realised (i.e. there is only one set of $x_1 ... x_n$ that realises the theory), we can construct explicit functional definitions of the $\tau_1 ... \tau_n$ via the Ramsey sentence. These says e.g. that τ_i is 'that thing that occupies the x_i-role within the theory'.

[13] See Kim (2005), Lewis (1970), Butterfield and Gomes (2020a). See also the debate between Marras (2005) and Morris (2020) on the topic of bridge laws. See also Crowther (2018) for a recent discussion of Nagelian reduction in physics.

2. We find another theory T* embedding new theoretical terms $\rho_1...\rho_n$. Suppose that the following sentence is a theorem of T^*: $T[\rho_1...\rho_n]$. $T[\rho_1...\rho_n]$ does not contain τ-terms, and it says that the original theory T is realized by a n-tuple $\rho_1...\rho_n$, taken from T^*. In case T is uniquely realised by the n-tuple $\rho_1...\rho_n$, Lewis shows that we can functionally define the ρ_i as the occupiers of certain x-roles in T, and those functional definitions are theorems of T*.

3. Following step (2), we can derive theoretical identifications $\rho_1 = \tau_1, \ldots, \rho_n = \tau_n$ by transitivity of identity. These are *bridge laws* and they play the role of Nagelian bridge laws in the theory derivation of T from T*.

Summing up more informally, the process goes as follows: we first specify the theoretical roles of the theoretical terms within a theory T via the Ramsey sentence of a theory, i.e. we build functional definitions for the terms $\tau_1...\tau_n$ in the theory. Then, we find a second theory T*. This theory can realise the former theory T in terms of $\rho_1...\rho_n$, and so we show that it contains theoretical terms $\rho_1...\rho_n$ which play the roles of the entities $\tau_1...\tau_n$. Thus, on the assumption that the Ramsey sentence of theory T is uniquely realised, we deduce bridge laws between the two theories, i.e. we connect the vocabularies of the two theories. This happens because we have terms that fall under the same functional profile and thus they can be identified thanks to functionalism. If a term τ_i is identified to a term ρ_i in this way, we say that τ_i is functionally reduced to ρ_i. The functional reduction of the theoretical terms within different theories is thus a step in the full derivation of the laws of the reduced theory from the reducing theory's laws.

Lewisian reduction is therefore a special form of Nagelian reduction in which theory deduction is couched in terms of logical derivation and in which bridge laws are functionally derived and thus deduced, as opposed to postulated as additional empirical hypotheses. Given that theories are formulated as logical sentences within the account, and that reduction is a form of logical deduction, this functional reductionist view can be said to fit naturally within the syntactic view of theories. In a nutshell, as opposed to the semantic view of theories that conceives theories as sets of models, the syntactic view takes theories to be sets of sentences.[14] We can indeed say that for the syntactic view "the structure of a scientific theory is its reconstruction in terms of sentences cast in a meta-mathematical language" (Winther (2021)). Since a prerequisite of the Lewisian account is the idea of expressing scientific theories as sentences formulated in the language of second-order logic, the account arguably falls within the syntactic side. Moreover, within this approach to theories, inter-theoretical relations such as reductive relations between theories are formulated as deductions under a given class of logical relations, and the account of reduction presented above presents this feature. We can thus appreciate how the topic of the nature of scientific theories heavily influences functional reduction, since the stance we take on theories is crucially correlated to the way in which we cash out functional roles, which are here defined via the Ramsey sentence. We

[14] See Winther (2021), Halvorson (2019) for an introduction. See Halvorson (2013) and Lutz (2017) for a recent defence of the syntactic approach; and Suppes (1961, 1967), Suppe (1974), Van Fraassen (1980), Ladyman et al. (2007), Wallace (2021) for a defence of the semantic approach.

shall see that the same is true for Semantic Functional Reduction, which is couched in terms of the semantic view, according to which representing scientific theories requires mathematical tools and not predicate logic. It should be stressed that we are not attempting nor requiring here a complete and accurate reconstruction of the debate on the nature of theories, especially because many versions of each view are available, and the distinction between the two is often blurry. Rather, we just aim to present them in a way that can help us highlight the difference between the two brands of functional reductionism discussed in the paper.

Moving on, let's discuss the relationship between this account of inter-theoretic reduction and the ontology of the theories it concerns. Notice that here and in the rest of the paper we will be careful in distinguishing between formal mode and material mode, i.e. in discussing reductionism at the level of theories and at the level of ontology.[15] Functional reductionism, as discussed so far in this section, is clearly a form of reductionism about theories. However, in Lewis's account, this functional reductionism about theories, which leads to identity relations between theoretical terms, is meant to be a way to ensure functional reduction about ontology as well. Lewis makes this clear in several places, for instance:

> The T-terms have been defined as the occupants of the causal roles specified by the theory T; as *the* entities, whatever those may be, that bear certain causal relations to one another and to the referents of the O-terms. (Lewis, 1972, p. 255)

The passage from the formal mode to the material mode is thus straightforward here. On the assumption that the theoretical terms refer to actual entities, the theoretical functionalisation is just a means to codify in a scientifically accurate way the roles played by the worldly entities referred to by the theoretical terms. That is, functional reduction of theoretical terms can be a guide to functional reduction of entities. Lewis is explicit about this. For him, the theoretical term 'electron' is meant to refer to an actual entity, as he wants to maintain a clear form of scientific realism.[16] Thus, when we functionally define a theoretical term in the upper theory and we find some other theoretical term in the bottom theory with the same role, we should believe also that there is a bottom entity (referred to by the term ρ_i) to which the upper entity denoted by τ_i is reduced to.

Notice that this is a form of realiser functionalism, since the functionalised entity at the top is type-identified with its realiser at the bottom. This is actually a radical consequence of the account which should be stressed: on the ontological level, the Lewisian account leads us to maintain identity relations between the reduced and the reducing entity. When the entity belonging to the bottom level behaves in the right way, *that same entity* turns out to be the upper-level entity which is the target of the reduction,

[15] Cf. Ladyman et al. (2007).

[16] Lewis' framework is indeed part of the so-called 'Canberra plan', a naturalistic philosophical framework (see Braddon-Mitchell and Nola (2008), Ch. 1).

in virtue of the fact that it plays the role of the target entity.[17] However, it should be stressed that the step from functionalism about theory to functionalism about ontology is not automatic nor mandatory: *contra* Lewis, one is free to claim that theoretical terms should not be taken as designating actual physical entities. More on this in the next subsections.

This concludes the presentation of the core of Lewis' account. In Section 4.2 we discuss two ways in which the Lewisian basis can be expanded to block possible objections. In Section 4.3 we confront the framework with the case study introduced in Section 2, and highlight the possible shortcomings of Syntactic Functional Reductionism.

3.2 Improving the Lewisian Basis

The Lewisian account presented in the last subsection constitutes a fully-fledged framework that makes the broad functional reductionist approach more precise. It presents a formal way to spell out the notion of functional role at the theoretical level (via Ramsey sentence), embeds a specific approach to reduction (Nagelian reduction), and shows a close link with a specific view about scientific theories (the syntactic one). In this subsection we present how the Lewisian core can be expanded to account for two specific issues, thereby improving the view. This improved version can be taken as the real basis of the Syntactic Functional Reductionist framework. In particular, we discuss here (1) the move from a Nagelian to a Neo-Nagelian model, and (2) the move to a more local version of functional reductionism.

The first issue, and thus the first step to take in expanding the base view, concerns the role of approximations in reduction. We start by pointing out that the commonly adopted version of Nagelian reduction is not the classic model proposed by Nagel (1962) and introduced in the last subsection, but a more refined approach that has been put forward by Schaffner (1967) and recently by Dizadji-Bahmani et al. (2010)[18]. This 'Neo-Nagelian' account relaxes the derivability criterion and argues that, to ensure reduction, it is sufficient to derive laws that are *approximately* the same as the laws of the original theory T_P. More precisely, according to this view, T_F reduces T_P iff we can build a theory T_P^* – which is a corrected version T_P standing in a relation of 'strong analogy' with T_P – which is derivable from T_F given some appropriate auxiliary assumptions and bridge laws. The reason is that it is almost never the case that we can derive the exact laws of an upper theory (to be reduced) from a bottom theory. At most, we can recover the behavior described from the top theory in an approximate way and just in particular situations. Why is it important to point this out here? The reason is that the issue behind the introduction of the Neo-Nagelian approach affects the Lewisian account of functional reduction as well, *qua* Nagelian-based account, even though this approach does not require postulated bridge laws. In fact, the Lewisian process of reduction requires the deduction of the reduced theory's laws from the reducing theory's laws, and requires us to express the terms ρ_i as playing the role of the τ_i. However, if it is true that we

[17] Indeed, within the Lewisian account of mental states, realiser functionalism turns out to be just a peculiar version of the identity theory about mental and physical states.

[18] See also Schaffner (2012) and Dizadji-Bahmani (2021).

cannot ever exactly deduce the original reduced theory, but just an approximate version of it, then also the Lewisian view needs to be amended like the classic Nagelian approach. We should thus expect to replace the reduced theory with a strongly analogous version of it, and accordingly, we should expect the ρ_i to functionally realise some terms which are not strictly speaking our original τ_i but rather terms that behave approximately like them.[19] For instance, if we are dealing with the reduction of classical mechanics to quantum mechanics, we cannot expect quantum systems to behave exactly as classical systems, but only approximately so. Syntactic Functional Reduction should therefore be based on Neo-Nagelian reduction, and not on classic Nagelian reduction.

Moving to the second point, we can draw an important distinction within the Syntactic Functional Reductionist framework concerning the difference between a *global* and a *local* version of the view. In fact, inter-theoretic reduction can concern either the reduction of whole theories or the reduction of only specific laws or models.[20] In the present case, both the Lewisian basic account and the Neo-Nagelian version just introduced are primarily formulated in terms of global theory-to-theory reduction. However, they can easily be turned into a more local form. In fact, once we have logically expressed the theory and derived the Ramsey sentence, we are actually free to functionalise either every theoretical term in the theory or just some of them. In this second case, we can provide functional definitions just for one or some 'problematic' terms, and perform a functional reduction only for them.[21] The passage from formal to material mode then goes as usual: once a specific term is functionally reduced, we can take the formal functional reduction as representing an ontological functional reduction in the actual world.

An important motivation to prefer the local version comes from Newman's objection.[22] That is, being committed to the Ramseyification of a whole theory exposes one to the objection according to which providing a set of entities which can realise the Ramsey sentence is really a trivial matter, because the Ramsey sentence (in this case) can at most constrain the cardinality of the set. One possible line of response to the challenge is to resort to one of the different strategies that have been proposed against Newman's objection itself (e.g. French and Saatsi (2006), Saunders and McKenzie (2014), Bueno and Meier (2019), Ladyman (2020)). However, it should be noted that some responses rely specifically on adopting the semantic view, and so they would likely not be available here. On the other hand, a simpler route is available, which is exactly to adopt the local version of the functional reductionist account, dissolving the problem from the outset. Given this advantage of the local version of the Syntactic Functional Reductionist approach, we take this as a natural update of the Lewisian basis, even though the more global version remains available.

To sum up, in the last two subsections we have introduced the Lewisian basis of the Syntactic approach to functional reductionism, and then presented two ways to improve

[19]Lewis (1970) mention this issue talking about the 'near-realisation' of the reduced theory's terms. This discussion is meant as an elaboration of that idea, that clarifies the position of Syntactic Functional Reduction with respect to the broader literature on inter-theoretic reduction and the Nagelian approach.

[20]See Rosaler (2015) on the distiction between global and local approaches to reduction.

[21]This is the approach adopted by Butterfield and Gomes (2020b), when discussing their case studies.

[22]See e.g. Demopoulos and Friedman (1985).

the account: the move to Neo-Nagelian reduction and the passage to a local form of functional reduction. The next subsection confronts our case study with this framework.

3.3 Thermodynamics and Syntactic Functional Reductionism

Section 2 presented the reduction of the second law of thermodynamics via the functional reduction of thermodynamical entropy to Gibbs entropy. This is an example of a local functional reduction of the upper-level quantity S_{TD} to the lower-level quantity S_G. We employ here this case study to analyse Syntactic Functional Reductionism. We discuss the example both at the formal and the material levels, and raise three possible shortcomings of Syntactic Functional Reductionism: a translation issue, a challenge related to approximation, and finally an ontological problem.

Let's start from the theory-level. At the level of theoretical reduction, the aim is to functionalise the upper-level quantity S_{TD} and then to find out a bottom-level realiser that can approximately play the relevant upper-level role in a given system. Within this framework, the functionalisation process is carried out first by the translation of the theory from mathematics to formal (second-order) logic, and then via the construction of the Ramsey sentence. In principle, this is a consistent project, but we face here two challenges. First, the whole translation process is not merely a challenging and complex task, but it could be taken to be complex in a futile or avoidable way. This aspect of the framework comes from its appeal to the syntactic approach to theories, but the translation passage could be seen more as an unnecessary attempt to shoehorn the mathematical formalism into the language of second-order logic, than as a genuinely useful step within the reduction process. Thus, an alternative functionalisation strategy that does not presuppose this passage would be preferable, other things being equal. Second, we have seen in the previous section that the higher-level theory we are meant to logically translate is not really thermodynamics, but rather an approximate version of thermodynamics, or another theory standing in a relation of strong analogy with it. Building such a theory is not a trivial task, and thus this adds an additional burden to the process of functional reduction, above the logical translation. In particular, since we are here dealing with theories as logically formulated sets of sentences, we cannot simply directly appeal to mathematical notions of approximation between models, but rather we have to rely on a syntactic-based form of approximation.[23] We shall see in Section 4 that Semantic Functional Reductionism fares better than the Syntactic framework in both this latter respect and the previous one.

Moving now to the connection between theoretical functional reduction and ontology, a puzzle can be presented with respect to our case study, if one adopts a scientific realist attitude (as Lewis does). In fact, whereas S_{TD} can be interpreted as a property of an individual system, S_G is defined as a property of a probability distribution over possible micro-states, i.e. a property of an ensemble. In this sense, it is not clear if the step from formal to material is warranted. Even if we grant the success of functional reduction

[23]Furthermore, if one wants to embrace scientific realism, in this case they would need to provide a notion of 'approximate truth', which is a notoriously difficult task. See Wallace (2021, pp. 8-9).

at the theoretical level, it is *prima facie* difficult to see how to translate the functional reduction from theoretical quantities to physical properties, since we are supposed to reduce a property of an individual system to a property of an ensemble, which looks more like a mathematical construct than a real physical property. The problem is exacerbated by the fact that the account entails type-identities between the reduced and the reducing quantities, which for Lewis reflect type-identities in the world. The puzzle is thus how a property of an individual system could be identical to an ensemble property.

This objection is specific to our particular case study,[24] but this problem is arguably symptomatic of a more general potential issue for the Syntactic Functional Reductionist framework. That is, the connection between theoretical functional reductionism and ontology is here very straightforward, but at the same time very strict. Functional roles are logically formulated using the Ramsey sentence and thus are expressed via logical predicates: x is that thing that plays a certain role, where playing a role is to satisfy certain predicates that connect that x with other kinds of theoretical terms in the network. If we adopt a scientific realist attitude, the way in which functionalism is connected with the world is very direct: the theoretical term, defined via the functional role, directly refers to the actual property that plays the roles represented by the theory. There is, allegedly, a 1:1 correspondence between theoretical terms and actual entities, which naturally matches the standard Quinean approach to the ontological commitments of theories.[25] The fact that in the case of thermodynamics this strict correspondence is challenged by the puzzle presented above is thus a reason to think that the formal-material link embedded in this framework is too strong.

In other words, the case study raises the following dilemma for Syntactic Functional Reductionism. On the one hand, one can reject scientific realism, thereby employing functional reductionism merely at the theoretical level. On the other hand, one can respond that theoretical functional reduction is a guide to ontological functional reduction, but maintain that we should not take the link as a straightforward entailment like the one pictured by Lewis. Holding up this second option would not be easy though. For instance, we need a reason why the functional reduction of a term entails a functional reduction in the world only in certain situations and not in others. And, in general, we would need a novel story about the theoretical-ontological link within the account, different from the Quinean approach which underlies the Lewisian picture.

4 Semantic Functional Reductionism

Semantic Functional Reductionism constitutes an alternative to the Syntactic framework in providing a model for functional reduction. It combines the general functional reductionist approach to inter-theoretic reduction with a model-based stance on scientific theories, and a structuralist conception of reduction as a relation between models, usually expressed mathematically. Reduction is thus characterised in terms of the func-

[24] For example, one could argue that when moving to the quantum domain, the canonical ensemble is represented by a density matrix and thus can represent a single quantum state.

[25] Cf. Quine (1948), Bricker (2016).

tional realisation of certain – mostly mathematical – roles in the upper-level theory's models by the theoretical elements in the lower-theory's models. This approach is arguably a (functionally-based) improved version of the structuralist account of reduction. Especially within physics, the framework takes the mathematical formalism in which theories are expressed at face value, and uses maths and mathematical models to specify the functional roles. Because of this, the view does not run into the issues raised previously against the Syntactic approach, since it (i) does not require logical translation of the mathematical formalism, (ii) accounts for approximation using a notion of approximation between models, and (iii) allows for a more flexible approach to ontology by replacing the Quinean stance on ontological commitment with a more flexible relation characterised in terms of representation between the models and the world.

In Section 4.1 we introduce Semantic Functional Reductionism, building on the semantic view of theories and the related structuralist account of reduction, in particular in the form defended by Rosaler (2015, 2019). In Section 4.2 we further develop the view by showing how it can account for Robertson's case study of functional reduction and by discussing the advantages of the view.

4.1 Introducing Semantic Functional Reductionism

As mentioned, the Semantic Functional Reductionist approach that we are going to introduce here importantly draws on the semantic view of theories and the related account of structural reductionism. It is thus necessary to introduce these two notions here.

According to the semantic view, representing scientific theories requires mathematical tools and not predicate logic, in contrast with the syntactic view. In the context of physics, which is the focus of the essay, theories are constituted by sets of models which are mainly mathematically formulated, in the sense of 'model' which is employed by physicists.[26] For the purpose of this paper, the two main features of the semantic approach are the following,[27] quoting Ladyman et al. (2007, 118): "(a) The appropriate tool for the representation of scientific theories is mathematics; (b) The relationships between successive theories, and theories at different scales whether spatio-temporal or energetic, are often limiting relations and similarities of mathematical structure (formally captured by structure-preserving maps or morphisms of various kinds), rather than logical relations between propositions". In saying this, they follow Suppes' famous slogan according to which "philosophy of science should use mathematics, and not meta-mathematics" (cf. van Fraassen 1980, 65). As we shall see, the Semantic Functional Reductionist framework follows these principles, by characterising functional roles in mathematical terms and functional reductive relations as structural relations between models. Doing this, it can thus be classified as a kind of semantic-based structural reductionism.

[26]By characterising the semantic view in this way we thus follow the approach of Van Fraassen (1980).

[27]As said in Section 3, we don't claim this description of the semantic view to be exhaustive, given the complexity of the debate, rather we just need a characterisation of the view that can distinguish it from the syntactic view in a way that is salient enough for our purpose of discussing functional reductionism.

More precisely, according to the structuralist account, reduction obtains by virtue of relations of (approximate) instantiation between theoretical structures belonging to different models. The view has been endorsed by Suppes, who claimed for instance that "the thesis that psychology may be reduced to physiology would be for many people appropriately established if one could show that, for any model of a psychological theory, it was possible to construct an isomorphic model within physiological theory." (Suppes 1967: 59).[28] The relation of isomorphism has been considered to be too strong in the subsequent literature, but the notion of reduction as a model-model mathematical relation has remained the hallmark of the approach. For example, in the passage quoted above concerning the semantic view, Ladyman et al. (2007) talk about reduction as a link between mathematical structures in terms of structure-preserving mappings or 'morphisms', whereas Wallace (2021, p. 16) argues that "reduction is something like *instantiation*: the realizing by some substructure of the low-level theory's models of the structure of the higher-level theory's models". An approach to reduction focused on model-model relations is also endorsed by Rosaler (2015, 2019). Notice that this view characterises reduction as a *primarily local* relation, that takes place between specific models. Global theory-to-theory reduction is thus derivative of local model-to-model reductions. Moreover, those relations of instantiation or morphism between the models will not be exact but approximate – for the same reasons expressed earlier in Section 3.2 – but approximation is here a relation between the models standing at different levels, and not between the higher-level theory and its own corrected version.

Having broadly presented the semantic view and structural reductionism, we start introducing Semantic Functional Reductionism by first looking more closely at a specific structuralist account of reduction, proposed by Rosaler, and show how we can build on this to develop a novel form of functional reduction alternative to Syntactic Functional Reductionism. One reason to pick this particular view is that it is arguably one of the most fully-fledged structuralist accounts of reduction available. Another reason is that, since we aim to show that the Semantic version of functional reductionism can be regarded as an improved version of the general structural reductionist view, the way in which Rosaler's account is formulated makes it particularly suitable for this goal.

Rosaler (2015, 2019) elaborates on the model-based view advanced by Ladyman et al. and Wallace, and proposes a notion of reduction as 'domain subsumption', where:

> Subsumption of the domain of some high-level description by a low-level description requires that any real behaviour that is accurately represented by the high-level description be represented more accurately and in at least as much detail by the low-level description. (Rosaler (2019), 272)

To formalise this notion he relies on model-model relations as follows:

> Theory T_h reduces$_T$ to theory T_l iff for every system K in the domain of T_h – that is, for every system K whose behavior is accurately represented by

[28]See also Suppes (1961).

some model M_h of T_h – there exists a model M_l of T_l also representing K such that M_h reduces$_M$ to M_l (Rosaler (2015), 59),

where a low-level model reduces a high-level model if "the low-level model accounts for the success of the high-level model at tracking the behavior of the system in question". E.g. for dynamical systems reduction this boils down to finding mathematical mappings between the low-level state space and the high-level state space which approximately tracks the evolution in the high-level space from the evolution in the low-level one.[29]

To briefly present an example of structuralist-type reduction within this proposal, a semi-classical model for a point-particle system can be mathematically matched with a quantum model of the same system, under the right conditions. In fact, thanks to Ehrenfest theorem, it can be shown that we can derive Newton's law from the Schrödinger equation for the system, if the quantum state is highly localised in space. This means that, within the quantum mechanical model, the centre of the localised wavepacket has a trajectory in configuration space which is (to a very high approximation) identical to the trajectory in configuration space of a point particle of mass m within classical mechanics (in the Hamiltonian formulation). Thus, the trajectory of the wavepacket can be practically considered as a solution to the classical dynamic equation for a classical particle, and we can draw a map between the quantum and the classical models defined over the respective state spaces.[30]

Taking stock of this, our aim now is to show how functional reductionism, as formulated according to the Semantic Functional Reductionist framework, can build on and improve this view. To begin with, recall that theoretical functional reductionism can be broadly characterised as the view according to which a theory T can be reduced to another theory T* in virtue of the fact that theory T* embeds theoretical elements (e.g. quantities, systems) which can play the theoretical (formal) roles of the theoretical elements belonging to T. The thesis of Semantic Functional Reductionist is then that to establish theoretical functional reduction we need to find lower-level mathematical structures, variables, or quantities playing approximately the same theoretical roles (i.e. roles in the models) of upper-level mathematical structures, variables or quantities in the upper-level model. This is a mainly mathematically-couched way of expressing the idea that functional reduction proceeds by identifying the lower-level realiser for an upper-level role: in our case, the role is spelt out in mainly mathematical terms as a role in the models, and the realiser is a piece of mathematical structure.

Semantic Functional Reductionism thus shares Rosaler's structuralist notion of reduction as domain subsumption – which focuses on reduction as the recovery of upper-level real behaviour from the lower-level, and cashes this out in terms of model-model relations – and reformulates this in functional reductionist terms. A reason for advocating for functional reductionism here is that it gives a justification to Rosaler's structuralist approach, or at least it exposes an implicit assumption: the reason why, to reduce a classical model to a quantum model, all we need to do is providing a formal account of

[29] Rosaler (2019, 4.2.2) generalises this notion for non-dynamical systems as well. Notice that this approach fits particularly well within the version of the semantic view defended by Van Fraassen (1980).

[30] See Rosaler (2015, 63) for more details.

how the quantum model can represent the behaviour described by the classical model, is given by the functionalist thesis that 'being a classical system' just means to perform certain roles within the model of classical mechanics. Indeed, consider again the example of structural reduction state above. We have claimed that, for certain kinds of physical systems, i.e. highly localised systems, we can build mathematical mappings between the state spaces for the same system within the bottom and the upper theories' models. That account of reduction is structuralist in the sense that we provide an asymmetrical inter-levels link between the two models. But, it may be asked why this particular mapping ensures reduction, i.e. why the mapping provides a reason to believe that we can recover the classical system from the quantum one. Adopting a functional reductionist version of the structuralist approach to reduction provides us with the justification: the condition for being a classical system is to play a certain role in the classical models, and the mathematical mapping at stake shows exactly that the quantum system can indeed evolve like the classical one. It is in this sense that Semantic Functional Reductionism is an improved form of structuralist reduction.

4.2 Thermodynamics and Semantic Functional Reductionism

We reconsider in this section the case study presented in Section 2. In Section 3.3 we analysed this case study with respect to Syntactic Functional Reductionism. We show here how the Semantic Functional Reductionist framework can accommodate the case study, in order to present in more detail how this approach works, and its differences from the Syntactic framework. In particular, we go through the same points that we have discusses in Section 3.3 with respect to the Syntactic approach to functional reduction. We thus consider first how the account works at the formal level, and second we discuss the theoretical-ontological link within this framework.

Let's start with the formal side of the account. One sharp contrast between a Semantic Functional Reductionist version of the case study and a Syntactic one concern the way in which the functional role itself is characterised within the account. With respect to the latter approach, functional reduction starts with the logical translation of the higher-level theory and the formulation of the S_{TD} theoretical role via Ramseyification in terms of formal logic. On the contrary, the core tenet of Semantic Functional Reductionism is to get rid of those steps. The privileged tool for representing scientific theories is maths, and we should just mainly stick with maths in specifying the functional roles. Thus, the approach of Semantic Functional Reductionism in dealing with our case study is to maintain that the functional roles we have to identify thermodynamic entropy with are just the mathematical roles that appear in Section 2, i.e. thermodynamic entropy is that quantity which 'for thermally isolated systems, is increasing in non-quasi-static processes, but non-increasing in quasi-static processes', as formally represented by the equations of thermodynamics introduced earlier. In other words, thermodynamic entropy is that mathematical function that bears those mathematical relations within thermodynamic models of thermally isolated systems, in thus-and-so conditions. Therefore, any bottom-level piece of mathematical structure in the lower theory's models that can fulfil those mathematical relations is said to functionally realise – and reduce – the thermodynamic entropy. Of course, exact fulfilment is not required: approximate realisation of the mathematical role is enough, and that very approximation is mathematically expressed. In other words, we just need the bottom-level model to approximate the upper-level one at the mathematical level.

The upshot is that there is a crucial difference between the Syntactic and the Semantic frameworks considered here. First, there is no need for logical translation within the present account: as far as the mathematical presentation of the functional reduction carried out in Section 2 is clear enough to show how reduction works, we should just take this at face value. That is, we can account here for the case study at stake simply by reading off functional reduction from the mathematical presentation provided by Robertson. This renders reduction comparatively easier to demonstrate. Second, the Semantic framework deals with approximation and approximate reduction in a straightforward way, directly exploiting the mathematical way of representing approximation which is employed by physicists. This distinguishes this framework from the Syntactic one, and makes it preferable to the latter, as it achieves the same goal in a simpler way.[31]

We move on now to the ontological aspects of the framework. We argue that the Semantic framework fares better than the Syntactic one in this respect, as it is more flexible. To recap, Section 4.3 elaborated on the link between theoretical reduction and ontology within Syntactic Functional Reductionism. That account embeds a Quinean view concerning the ontological commitments of theories, which raises a potential problem when confronted with cases of functional reduction like the thermodynamics one. That account indeed predicates a direct correspondence between theoretical terms and entities, and also entails identity relations between the upper functionalised entity and the bottom entity realising its role, and thus arguably provides a too strict account of the ontological committments of functional reduction.

In contrast, Semantic Functional Reductionism allows for a more flexible approach. At the theoretical level, this framework adopts a model-based account of reduction, and frames functional reduction as a realisation relation between models: theoretical elements in the reducing theory's models, such as mathematically-formulated quantities, are taken to functionally realise certain patterns of behaviour described by the reduced theory's models. Given its reliance on models, the view is less restrictive concerning the relationship between theory and ontology, and concerning the ontological consequences of theoretical reduction, because models are merely required to *represent* the world. Adapting Wallace's (2021, 7) words to our context, we can stress that while, within the Syntactic framework, "The relations that a good theory's empirical statements have to the facts are those familiar from ordinary-language semantics: truth, reference, satisfaction", within Semantic Functional Reductionism "a theory makes contact with empirical data by modelling them. [...] The theory/world relation here is *representation*, more akin to the relation between map and territory than that between word and object". Thus, when we claim that e.g. Gibbs entropy plays the role of thermodynamic entropy,

[31]Furthermore, if we endorse scientific realism, we don't face any problem related to the fact that the models are just approximate: since the relation between models and world is just one of representation, we don't need to deploy a notion of *approximate truth* (see ft. 23).

we are not committed to the claim that there's a specific property of a system denoted by thermodynamic entropy that is realised by a specific property denoted by Gibbs entropy and that they turn out to be ontologically identical properties – rather, the Semantic approach just claims that the physical systems represented by the models of statistical mechanics can be modelled accordingly to the thermodynamics models under the right conditions. Because of this, we are not committed to any specific view about the kind of relations between the physical systems – and thus we are not forced to endorse identity relations between them – and we are also free to be selectively realists about the type of entities our theories represents. In this way, we have an account which is both flexible enough to let us decide case by case, and which does so for a principled reason, and thus does not run into a dilemma like the one we raised against Syntactic Functional Reductionism.

Before concluding, it is worth pointing out that the flexibility and the reliance on models that make the Semantic approach to functional reduction more liberal and more straightforward, as it does not require the regimented translation into logic on which Syntactic Functional Reductionism builds on, could also prompt possible drawbacks for the approach. First, the model-world relation of representation is left largely under-specified, and thus the ontological implications of inter-theoretical reduction are less clear within this approach than in the Syntactic one. Second, as admitted by Wallace (2021, 10) while discussing the advantages of the semantic view of theories in dealing with issues such as approximation, "it's certainly true that we have a fuller and more carefully developed philosophical analysis of the semantic notions appealed to in the language-first [i.e. syntactic] view than we do of math-first-style [i.e. semantic-style] scientific representation".

5 Gauging the Two Forms of Functional Reductionism

	Syntactic Funct. Red.	**Semantic Funct. Red.**
Theories	Syntactic	Semantic
Theor. Reduction	Nagelian: Logical derivation	Structural: Model-model
Functional Roles	Logically formulated	Mainly math. formulated
Scope	Local or Global	Primarily Local
Ontology	Quinean approach	Representation relations

This section wraps up what we have presented so far, in order to compare the two forms of functional reductionism we introduced, and gauge the general advantages and drawbacks of each view beyond the particular discussions that we carried out around the case study of the paper.

To recap, Syntactic Functional Reductionism, which builds on the only full-fledged functional reductionist account in the literature, models the account as an improved and *sui generis* form of Nagelian reduction, cast within the syntactic approach to theories, where functional roles are logically formulated. Semantic Functional Reductionism, which is a novel introduction, provides a functionalist and upgraded version of structuralist reduction, where scientific theories are based on models and functional roles are accordingly (mostly) mathematically formulated. Both views can be classified as either local or global approaches to reduction, even though the Semantic one is markedly more locally scoped. The table above sums up these features.

Discussing the case study of thermodynamics has allowed for a more detailed presentation of these frameworks and a confrontation between them. In a nutshell, Syntactic Functional Reductionism enjoys the following advantages: it provides a rigorous framework; embeds a clear link with ontology and with ontological reduction; can be easily endorsed by those that already support the syntactic view and/or Nagelian reduction. At the same time, it presents the following drawbacks: it requires a translation of theories into logic; faces issues related to approximation; the link it imposes between theoretical reduction and ontology, and the one-to-one mapping between terms in different theories, are very rigid and demanding. On the other hand, Semantic Functional Reductionism has the following merits: it is a more liberal framework with respect to the ontological implications of reduction; it does not require logical translations, as functional roles are directly extracted from the models and functional reduction is mainly mathematically formulated; makes dealing with approximations easier; can be easily endorsed by those that already support the semantic view and/or structuralist reduction. However, we also concede that the high flexibility of the account (both in terms of theoretical reduction and link with ontology) can be considered a drawback, compared with the rigour of the Syntactic framework.

Thus, overall, the discussion of the last sections brings to light the strengths and weaknesses of each approach and we believe that each approach is ultimately viable. In this sense, rather than completely ruling out one approach or the other, the paper has been focusing on delving deeper into the Lewisian approach while also proposing a possible alternative, eventually providing a classification that can bring clarity to the current and future literature on functional reduction.

A final note is in order, concerning the relationship between Syntactic Functional Reduction and Semantic Functional Reduction as represented in the paper. That is, we grant that the distinction between the two accounts is more of a matter of degree, rather than a clear-cut distinction. Indeed, we sketch here a way in which the Lewisian basis underlying the Syntactic approach could be embedded into the Semantic approach, and a sense in which Semantic Functional Reductionism can be said to incorporate syntactic aspects. On the one side, the Lewisian approach involving functional definitions and cross-theoretical identifications could also be predicated in mainly mathematical terms along the lines of the Semantic approach. That is, instead of defining terms via a logically formulated Ramsey sentence, we could formulate the Lewisian definitions in terms of the roles in models, thus via mathematics. However, this is not the way in which the view has been presented in the literature, and thus we take the characterisation of Section 3 as more representative of the actual debate. On the other side, as pointed out by Wallace (2021, 7), it could be argued that "extracting coherent content from physics requires substantial reconstruction along language-first [i.e. syntactic] lines", and thus the Semantic approach to functional reduction could be said to shade into the Syntactic one in this respect. That being said, the classification defended in the paper is still useful and important as it carves important differences between two possible approaches to functional reduction and it can be a helpful starting point for future literature. In particular, even if the distinction is not a clear-cut one, the novel approach alternative to Syntactic Functional Reduction that we introduced still presents important differences with the latter and bears crucial advantages on it, as we argued, and could thus be welcomed by those defenders of functional reductionism that do not want to be committed to a kind of Syntactic Functional Reduction.

6 Conclusion

Functional reductionism is a candidate account for scientific reduction, that provides an alternative to more standard approaches like Nagelian and structuralist reductionism, and that has recently been fruitfully applied to inter-theoretic reductive relations between theories like thermodynamics and statistical mechanics, classical mechanics and quantum mechanics, and general relativity and quantum gravity theories. Its potential value is demonstrated by those applications, but the view is still underdeveloped in several respects. This paper offers a thorough analysis of this approach and develops two kinds of frameworks that provide two fully-fledged alternative models of functional reductionism, that improve and clarify the view. The paper thus brings the functional reductionist approach to theoretical reduction to a higher level of clarity, and provides a more complete picture of how this account works with respect to both theoretical and ontological reduction, thereby contributing to making functional reduction a viable account for reduction.

References

Albert, D. Z. (2015). *After physics.* Cambridge: Harvard University Press.

Braddon-Mitchell, D. and R. Nola (2008). *Conceptual Analysis and Philosophical Naturalism.* Bradford.

Bricker, P. (2016). Ontological Commitment. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 ed.). Metaphysics Research Lab, Stanford University.

Bueno, O. and T. Meier (2019). Structuralism, empiricism, and newman's objection. *Principia: an international journal of epistemology 23*(1), 53–67.

Butterfield, J. and H. Gomes (2020a). Functionalism as a species of reduction. *http://philsci-archive.pitt.edu/18043/.*

Butterfield, J. and H. Gomes (2020b). Geometrodynamics as functionalism about time. *http://philsci-archive.pitt.edu/18339/.*

Callender, C. (1999). Reducing thermodynamics to statistical mechanics: The case of entropy. *Journal of Philosophy 96*(7), 348–373.

Crowther, K. (2018). Inter-theory relations in quantum gravity: Correspondence, reduction and emergence. *Studies in History and Philosophy of Modern Physics 63*, 74–85.

Demopoulos, W. and M. Friedman (1985). Bertrand russell's the analysis of matter: Its historical context and contemporary interest. *Philosophy of Science 52*(4), 621–639.

Dizadji-Bahmani, F. (2021). Nagelian reduction in physics. In *The Routledge Companion to Philosophy of Physics*, pp. 499–511. Routledge.

Dizadji-Bahmani, F., R. Frigg, and S. Hartmann (2010, 11). Who's afraid of nagelian reduction? *Erkenntnis 73*, 393–412.

Esfeld, M. and C. Sachse (2007). Theory reduction by means of functional sub-types. *International Studies in the Philosophy of Science 21*(1), 1–17.

French, S. and J. Saatsi (2006). Realism about structure: The semantic view and nonlinguistic representations. *Philosophy of Science 73*(5), 548–559.

Halvorson, H. (2013). The semantic view, if plausible, is syntactic. *Philosophy of science 80*(3), 475–478.

Halvorson, H. (2019). *The Logic in Philosophy of Science*. Cambridge University Press.

Huggett, N. and C. Wüthrich (2021, January). Out of nowhere: The emergence of spacetime in quantum theories of gravity. *http://philsci-archive.pitt.edu/18612/*.

Kim, J. (1998). *Mind in a Physical World: An Essay on the Mind–Body Problem and Mental Causation*. MIT Press.

Kim, J. (2005). *Physicalism, or something near enough*. Princeton University Press.

Knox, E. (2019, 8). Physical relativity from a functionalist perspective. *Studies in History and Philosophy of Science Part B - Studies in History and Philosophy of Modern Physics 67*, 118–124.

Ladyman, J. (2020). Structural Realism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2020 ed.). Metaphysics Research Lab, Stanford University.

Ladyman, J., D. Ross, with John Collier, and D. Spurrett (2007). *Every Thing Must Go*. Oxford University Press.

Lam, V. and C. Wüthrich (2018). Spacetime is as spacetime does. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 64*, 39–51.

Lam, V. and C. Wüthrich (2020). Spacetime functionalism from a realist perspective. *Synthese*.

Levin, J. (2021). Functionalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2021 ed.). Metaphysics Research Lab, Stanford University.

Lewis, D. (1970). How to Define Theoretical Terms. *The Journal of Philosophy 67*(13), 427.

Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy 50*(3), 249–258.

Lorenzetti, L. (2022). Functionalising the wavefunction. *Studies in History and Philosophy of Science Part A 96*, 141–153.

Lutz, S. (2017). What was the syntax-semantics debate in the philosophy of science about? *Philosophy and Phenomenological Research 95*(2), 319–352.

Marras, A. (2005). Consciousness and reduction. *British Journal for the Philosophy of Science 56*(2).

Morris, K. (2020). Does functional reduction need bridge laws? a response to marras. *The British Journal for the Philosophy of Science*.

Nagel, E. (1962). The structure of science: Problems in the logic of scientific explanation. *Philosophy 37*(142).

Quine, W. V. O. (1948). On what there is. *Review of Metaphysics 2*(5), 21–38.

Robertson, K. (2020). In search of the holy grail: How to reduce the second law of thermodynamics. *British Journal for the Philosophy of Science*.

Rosaler, J. (2015). Local reduction in physics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 50*, 54–69.

Rosaler, J. (2019). Reduction as an a posteriori relation. *British Journal for the Philosophy of Science 70*(1), 269–299.

Saunders, S. and K. McKenzie (2014). Structure and logic. In L. Sklar (Ed.), *Physical theory: Method and interpretation*, pp. 127–162. Oxford University Press Oxford.

Schaffner, K. F. (1967). Approaches to reduction. *Philosophy of Science 34*(2), 137–147.

Schaffner, K. F. (2012). Ernest nagel and reduction. *The Journal of Philosophy 109*(8/9), 534–565.

Schroeder, D. V. (1999). An introduction to thermal physics.

Suppe, F. (1974). *The Structure of Scientific Theories*. Urbana, University of Illinois Press.

Suppes, P. (1961). A comparison of the meaning and uses of models in mathematics and the empirical sciences. In *The concept and the role of the model in mathematics and natural and social sciences*, pp. 163–177. Springer.

Suppes, P. (1967). What is a scientitic theory? *Philosophy of science today*, 55–67.

Tong, D. (2012). Statistical Physics. *http://www.damtp.cam.ac.uk/user/tong/statphys.html*.

Van Fraassen, B. (1980). *The scientific image.* Oxford University Press.

Wallace, D. (2012). *The Emergent Multiverse.* Oxford University Press.

Wallace, D. (2021, December). Stating structural realism: mathematics-first approaches to physics and metaphysics.

Winther, R. G. (2021). The Structure of Scientific Theories. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021 ed.). Metaphysics Research Lab, Stanford University.