Exploratory Modeling and Indeterminacy in the Search for Life

Franklin Jacoby

Abstract
The aim of this article is to use a model from the origin of life studies to provide some depth and detail to our understanding of exploratory models by suggesting that some of these models should be understood as indeterminate. Models that are indeterminate are a type of exploratory model and therefore have extensive potential and can prompt new lines of research. They are distinctive in that, given the current state of scientific understanding, we cannot specify how and where the model will be useful in understanding the natural world: in this case, the origin of life on Earth. The purpose of introducing indeterminacy is to emphasize the epistemic uncertainty associated with modeling, a feature of this practice that has been under emphasized in the literature in favor of attempts to understand the more specific epistemic successes afforded by models.

Hear and attend and listen; for this befell and behappened and became and was.

—Rudyard Kipling, *Just So Stories*

1 Introduction

Interest in scientific modeling has recently given prominence to exploratory models, a type of model that may not have a target that it represents, but which still has epistemic value because it allows scientists to explore something, be it a phenomenon, a theory, a hypothesis, a concept, or something else (Gelfert, 2018; Massimi, 2019). Part of what makes these models interesting is that accounts of models based on representation (van Fraassen, 1980; French, 2003; Giere, 2004; Giere, 2010; van Fraassen, 2008; Suarez, 2010; Suárez, 2004; Suárez, 2003) do not seem to capture what exploratory models are doing because those accounts presume models have targets, whereas exploratory models may not. Many details surrounding these models need clarification, for instance whether lacking a target is enough to make a model exploratory and whether exploratory models are the same as minimal models (Fumagalli, 2015; Fumagalli, 2016).[1]
Another clarificatory need, one that this article will confront, is to establish what the epistemic value of an exploratory model might be given its lack of target. What do we get from exploration via modeling? This article seeks to bring clarity to this issue by arguing that some exploratory models—such as the GARD model from the origin of life studies—have an epistemic value that stems from the very fact that we do not know *whether* it has a target. Consequently, we need to be careful about what we expect the model to do and how we evaluate it. It may be that exploratory models like the one discussed in this article will give us genuine knowledge of actual-world phenomena, but they also might not and, crucially, it may be very difficult to establish this. We can capture some of this difficulty by appealing to what I will call indeterminacy.

The proposal developed in this article builds on some common elements shared with Gelfert's treatment of exploratory models (Gelfert, 2018), but stands in contrast to other attempts to understand these models using modality. Two specific attempts to understand the epistemic value of models in this way are 1) to appeal to analogical reasoning and 2) to appeal to how-possibly explanations. These attempts lie at the intersection of exploratory and fictional modeling literatures. I will argue that these approaches do not work here. To show why this is the case, I will present an influential exploratory model used in the origin of life studies (2). I will then show that it does not support analogical reasoning or how-possibly explanations (section 4). Instead we should think of models like the GARD as exploratory tools that provide a context for interpreting evidence and which help establish research priorities. If this is right, the GARD model from the origin of life studies shows us that exploratory models can have very indefinite, or indeterminate, connections to the actual world; this indeterminacy is where the epistemic value of this model lies. Indeterminacy is valuable because it provides motivation and an interpretive structure, both of which prompt new lines

1A number of authors discuss models without targets and how such models relate to various accounts of modeling is intensely complex (Frigg and Nguyen, 2016; Gelfert, 2018; Massimi, 2018; Massimi, 2019; Knuuttila, 2005; French, 2003; French, 2014; Giere, 2010; Poznic, 2016; Suárez, 2003; Contessa, 2011) of research and a context for evaluating new evidence. Scientists can use indeterminate models in this way because of the uncertainty over the target: were it known whether there was a target, the model would not have the same epistemic value. The view defended here is compatible with, but emphasizes differently, some previous discussions of exploratory models.

2 A Model in Search of Life

The study of the origin of life is intensely difficult because, like cosmology, paleontology, and many other studies, the subject matter is so removed from our time and place; direct forms of investigation, such as through laboratories, observations, or experiment, is impossible. Modeling has consequently proven a useful avenue for researching where life has come from on Earth and how life might in principle develop (say on other planets). One model that scientists have used to gain traction with origin-of-life questions is the GARD model. The following sections will lay out one problem those studying the origins of life face, show how the model works, and then clarify how the GARD model makes headway with solving the problem.

2.1 The problem: deciding among potential protocells

There are a number of ways that life, we think, might have emerged, but there are certainly three promising options (Lanier and Williams, 2017). In entertaining one or more of these options, scientists are taking a contemporary function found in living organisms, isolating it, and exploring whether such a function could provide the foundation for life's origins. These three candidates are not logically exhaustive: there are other ways one could study the origin of life and other ways life could have emerged from physical systems. However, one has to start somewhere and there is good reason to think these three candidates are the most promising, given what we now know. Those three are as follows:

1. The precursors to life were self-replicating molecules that store information (replication first);
2. The precursors were molecules that could maintain themselves (metabolism first);
3. Life emerged from structured aggregates of molecules (compartmentalization first).

Even assuming that one of these options can help us understand life's origins, deciding among them is not at all easy, especially since we are so far away from life's origins on Earth and we have not found life extra-terrestrially, nor indeed the spontaneous generation of life on Earth. Nevertheless, one way of breaking down the problem is to select what sort of organic molecules may have been involved in one or more of the three options. Lipids and RNA are two popular proposals (the lipid-first and RNA-first hypotheses). Ultimately, either molecule (or something very much like it) would need to perform all three of the functions listed above. Accounting for RNA's self replication is relatively straightforward,[2] but it is a complex molecule that is not particularly stable, making its prebiotic formation difficult to explain given the conditions that may have been present on Earth when life formed.[3] Lipids, in contrast, are often much simpler molecules that can spontaneously form structured aggregates (addressing the compartmentalization mentioned above) (Segré et al., 2001), but lipids do not obviously store information and replicate the way RNA does. If lipids are to be a promising precursor to life, scientists must account for how aggregates of lipids could metabolize and how (generally like RNA) they could store and pass on information. This has been an active research project for the past 20 odd years, which makes the lipid world hypothesis relatively new compared with the nearly 60-year-old RNA world hypothesis (Neveu, Kim, and Benner, 2013).

One hypothesis about how lipids might store and pass on information involves autocatalysis, a process whereby aggregates of molecules self-sustain and replicate one another. The challenges, then, are to try and establish what conditions are required for autocatalytic aggregates of lipids and whether such an aggregate is plausible. The GARD model attempts to make some progress with these challenges.

2.2 The GARD model

The lipid-first hypothesis assumes that precursors to life emerged in a “primordial soup,” a liquid in which there were a number of different types of organic models, particularly lipids, in high concentrations (Segré et al., 2001). These lipids, when the conditions were right, spontaneously formed structured assemblies, assemblies that later evolved (perhaps through Darwinian evolution, but perhaps not) into life. An initial challenge, however, is to assess whether such lipids may have existed, whether they could form assemblies, whether the assemblies could reproduce and pass on information, and whether the conditions were suitable for the

2 See Neveu, Kim, and Benner 2013 for an overview of the advantages of the view

3 See Bernhardt 2012 for an overview of some of the problems with the RNA-first hypothesis.

formation of those assemblies.

There is much debate about whether the primordial soup existed, but Kahana, Schmitt-Kopplin, and Lancet (2019) argue that recent observations of conditions on Saturn's moon Enceladus indicate that a primordial soup on ancient Earth is more possible than critics imagined. Assuming such a soup was present on Earth some 3.7 billion years ago, and assuming there were sufficient lipids of the right kind in that soup, could molecules form assemblies that self-replicate and evolve? The GARD model (Graded Autocatalysis Replication Domain) is an attempt to explore this possibility, specifically the possibility of the reproduction and evolution of assemblies of lipids (Segré, Pilpel, and Lancet, 1998b). How does it explore this possibility? Lancet, Zidovetzki, and Markovitch (2018, page 4) describe the GARD model as a toy that offers a "proof of concept", a proposal that I will return to in section 4. The model's purpose is to show the plausibility of lipids forming self-replicating assemblies that can pass on information. An assembly of this kind is a sort of protocell, not a living cell, but possibly a precursor to it.[4]

The model can be expressed in the following differential equation: [5]

$$\frac{\mathrm{d}D_{ij}}{\mathrm{d}t} = [k_{ij}M_iM_j - k_{-ij}D_{ij}]\left[1 + \sum_{p,q,=1}^{F}\beta_{\mu(i,j)v(p,q)}D_{pq}\right] - \lambda D_{ij}$$

D is the
number of dimer molecules that form from two random (*i* and *j*) lipid or lipid-like monomers, of which there might be many species. The monomers (M) can freely pass through a membrane that encloses the assembly, but the dimers cannot. The set of M is *F*. To the right of the equality relation, the first k term describes the rate at which two monomers (*i* and *j*) spontaneously form dimers within the assembly. That rate depends upon which two monomers form the dimer. The second, negative *k* term describes the decomposition rate of those same dimeric species into monomers: the rate at which species within the assembly spontaneously decompose. Both *k* rates are thought to be quite slow, in part because it is stipulated that 1) the protective membrane allows for the passage of monomers, but 2) no dimers can cross the membrane, and 3) the monomer concentrations are constant within the enclosed vesicle.

The beta term is a matrix of catalytic relationships between the dimeric species forming that assembly. Each species has a different probability for catalyzing each species (including itself), i.e. for any two given species, there is a probability that they will enhance one another's formation (mutual catalysis). Additionally, a species might promote its own formation (autocatalysis). If all the species are arranged along both the x- and y-axes of a matrix, then scientists can assign probability values for catalysis at the intersection of each row and column. We might imagine a chess board with monomers arranged along the top and one side; each square of the board, rather than having a chess piece, has a number, say at the 3F position. That number reflects the probability that the monomer species at the top of the F column will catalyze the monomer species at row 3. Each monomer has both a row and a column, so the matrix can represent any given monomer's role as catalyst and as base. So the matrix captures all the unique catalyzing relationships among all the species of monomers. The sum of all those catalyzed reactions is captured in the brackets of the differential equation.

Generally speaking, this differential equation describes the rate of growth of an assembly of lipids, which is the product of the spontaneous composing/decomposing rate and the catalysis. The lambda is a volume measure, so the greater the assembly volume (keeping number of species constant), the slower the rate of growth.

An important feature of this model is that it can be said to describe an assembly that stores and transmits information without DNA or RNA. The information is stored in the type and relative abundance of the dimeric species forming the assembly. So unlike modern life where individual molecules store information in variable base-pair sequences, the GARD assemblies store information in the variable type and number of dimeric species. That is, there is a pattern to the number and type of dimeric species in a vesicle described by the GARD model; that pattern is largely unchanging as the vesicle evolves, grows, and splits into two vesicles.

4 Whether and to what extent it provides this proof is open to question. See Lancet, Zidovetzki, and Markovitch 2018 for some discussion of the criticisms of the model and how those criticisms might be ameliorated.
5 There are many ways to adjust this equation and many versions of it. Choice of version depends in part on what assumptions and what levels of precision researchers are interested in exploring. This particular version is used by Segré et al. (1998a). I discuss it here because it is the original and the first demonstration of a compositional genome. For a brief history of the model, see Lancet, Segrè, and Kahana (2019). A more recent version can be found in Lancet, Zidovetzki, and Markovitch (2018), a version which links the model dynamics more closely with entropy, a topic beyond the scope of this article.

There are a number of approaches one can take to solving the equation (see 1998b, pages 562-563), though one common strategy is to use a Monte Carlo type of simulation technique (as described in fig. 1 caption in Segré, Ben-Eli, and Lancet, 2000). By doing so, several interesting things can be shown.

1. With the right starting conditions, lipids can assemble.
2. The assembly is an autocatalytic entity (an entity that self-sustains) in a steady state.
3. While in a steady state, lipid assemblies can reproduce and pass on information that persists, thus having the potential to undergo a kind of evolution.

So what the model here has shown is that it is theoretically conceivable that assemblies of lipids, formed by certain lipids and under certain conditions, can be said to carry information in their composition. This information is "stored" in the concentration ratios of the components of the assembly, ratios that remain constant through an increase in volume of the assembly Lancet, Zidovetzki, and Markovitch (2018, page 3). This suggests that one could conceive of a kind of protocell made entirely of lipids that has some of the functions of modern day life, specifically the carrying of information, very generally construed, and the self-maintenance (metabolism). To return to the three possible avenues for emerging life (section 2.1), the GARD model promises to make headway with all three (lipids can compartmentalize, store information, and metabolize). The suggestion is quite significant: a kind of primitive life that stores information in composition, but not in sequences like modern DNA and RNA. Remarkable though the model is, we must be careful about what we take it to show because it is very easy to imagine that it tells us how life actually or could form, but this is not where the value of the model lies. That is, although the model does involve possibility (given relevant assumptions of lipids and starting conditions), the possibility is not very closely aligned with the actual world and therefore does not support analogical reasoning or how-possibly explanations philosophers often ascribe to models without targets (more on this in section 4.1).

2.3 Model Assumptions and Uses

Although this model has proven fascinating, there is reason to be cautious about what precisely it does for us given the assumptions underpinning it. Though the fact that there are assumptions is trivial: all models make assumptions. Let's examine a few of those associated with the GARD model. The purpose of this exercise is not to undermine the value of the model, but ultimately to show that its value is neither as a tool for analogical reasoning nor as a tool for how possibly explanations. Some crucial assumptions are these:

- The different dimers have similar kinematic properties.
- Those properties are conducive to the kinds of spontaneous reactions specified in the model, such as mutual catalysis and forming bonds to form dimers in the assembly.
- The monomers have specific concentrations inside and outside the membrane-bound assembly.
- The environment, apart from supplying monomers, stays stable at the right conditions for dimers to form as specified in the model.
- Monomers and dimers are the only molecules present.

These are substantial assumptions; in fact, it is hard to imagine how such a system could exist at any time or place, whether that be a present day lab, 3.7 billions years ago on earth, or elsewhere in our solar system. Now there has been some work to address the first assumption about lipid behavior as well as environmental assumptions and assumptions about dimer concentrations. Armstrong *et al.* re-used the GARD model using kinetic parameters more closely aligned with some existing lipids (2011). They also varied dimer concentrations and tried to make more realistic assumptions about the environment. However, this work making the model more "realistic" amounts to tweaking and manipulating the model. It does not tell us anything direct about the context in which life emerged nor what early protocells were like. As a first pass, here are some more specific things it strikes me that the model might tell us:

1. The hypothetical behavior of actual lipids
2. The possible conditions from which life emerged
3. The possibility of non-DNA information carrying assemblies of molecules

4. How to prioritize future research

Option 1 cannot work as it stands because it is not clear whether the lipid behavior described by the model is in fact hypothetical or actual. If there were or is an assembly of lipids as specified in the model, then it is not hypothetical, but we do not know whether this is the case. We know lipids exist and we know some lipids can catalyze, but does this amount to having modal knowledge about whether lipids *could* form assemblies and autocatalyze? It does not, because the model does not tell us anything modal about actual, real lipids. Some authors are quite explicit about the indeterminacy of the lipids described by the model (Kahana, Schmitt-Kopplin, and Lancet, 2019, page 1272). So I am not certain whether option 1 can be given a definitive answer because a definitive answer would require knowing whether it is possible that there exist lipids that behave along the lines of the GARD model and we do not know whether this is possible.

Option 2 is not viable either because it demands too much from the model, which has not allowed for any inferences about life whatsoever. Option 3 bears more promise, but as with option 1, it is not clear how to read the possibility. We do not know how far the possible world is that has lipid assemblies as specified in the model, which does not mean the model cannot be analyzed using possible worlds. But it does mean that comparisons with the actual world are tenuous. Indeed, it could be that the actual world has such assemblies, but given the topic, it is unclear how to establish whether this is the case. We are still left with options 3 and 4. Some users of the model claim that the model demonstrates a kind of conceptually possibility, or proof of concept, for mutually catalytic sets, which touches on option 3. In their (1998b, page 9), Segré, Pilpel, and Lancet claim this:

> The GARD model further allows a rigorous kinetic analysis of the behavior of the mutually catalytic components, providing a quantitative assessment of their self-replication capacity. We demonstrate that components of a mutually catalytic set, none of which is necessarily autocatalytic, might resemble in their collective behavior the kinetics of a single autocatalyst. The GARD model provides a natural and simple quantitative definition of self-replication capacity of members of catalytic sets.

This suggests to me the model allows scientists to make claims about the behavior of a set of “mutually catalytic components”, which we might conceive of as lipid dimers. That set might have the property of resembling an entity that autocatalyzes. That is, even though no one member of the set can reproduce itself, the set as a whole can reproduce itself. So the model establishes the theoretical possibility of information-carrying molecules that are not DNA or RNA, which is an important contribution. So if we read option 3 as a theoretical or conceptual (not modal) claim, the option might be a good characterization of what the GARD model can deliver. What about option 4?

Another claim scientists make about the model is that it has the potential, despite its abstraction and many assumptions, to be made “more realistic” (Armstrong et al., 2011; Lancet, Zidovetzki, and Markovitch, 2018). I think we should interpret this as the claim that if we can provide better justification for making the assumptions found in the GARD model, then we will have more confidence in a lipid-first hypothesis. Prompting that search for justification is, I think, another point where the real value of the model lies (option 4), not in telling us anything directly or analogically about the actual world. More generally we might say that the GARD model points toward new lines of research, identifies what sources of evidence we need, and specifies where the gaps in our knowledge lie. This is incredibly important. The model might to help make future research more productive and to help, ultimately, increase our confidence in the sequence of steps that we think proto-life followed.

However, I think we need to be careful here what our confidence is about and what it is not about. While generally it may be true that better justification for a model would increase our confidence in the sequence of steps by which life formed, finding such justification would not establish that the hypothetical lipids posited by the GARD model actually did or even could exist. In other words, better justification for model assumptions does not entail that we have better justification that modeled entities actually exist or existed. We just do not know whether they do or did. Our epistemic position could, in principle, change. We might travel back in time 3.7 or 3.8 billion years an examine primordial soup on Earth. We might find primordial soup today (on Earth or elsewhere). We might create an artificial lipid-based mutually catalytic assembly in the laboratory. We have not managed any of these things and thus far this is not what is valuable about the GARD model.

Because the value of the GARD model does not involve representation and instead involves exploring evidence, justification, and theoretical assumptions, it is a good candidate for an exploratory model. What does it mean to be an exploratory model and what is distinctive about it?

3 Exploratory Models

The GARD model is a good candidate for an exploratory model, but fully understanding its epistemic value requires further developing our philosophical account of these models. Traditionally, philosophers have thought models have targets and the former are valuable in virtue of the representation relation they have with the latter, but recent discussions of modeling have indicated that some models, particularly exploratory models, do not have targets. A number of philosophers, such as Massimi (2018; 2019), Gelfert (2018), and Fisher (2006), have gone down this road. These accounts stress that exploratory models do not need to have a target to be epistemically valuable and in many cases either have no target or may not have a target. What makes these models valuable is their capacity to allows scientists to explore, in some sense that needs further development.

Despite this common emphasis on exploration, there is an open question about how to characterize the epistemic value of a model with no target and why the lack of target might be important for the epistemology. As we'll examine further in section 4.1, Massimi takes exploratory models to essentially involve the exploration of modal properties and to treat exploratory models as a type of fictional model. This does not seem applicable in the GARD case because the model has not offered us modal knowledge. So this characterization of exploratory modeling applies more neatly to a different class of models associated with less epistemic uncertainty and more specific knowledge outputs (i.e. the ascription of modal properties). We might say that scientists know too little in the GARD case to learn something specific about the origin of life, whereas for instances when scientists do have greater knowledge initially, they are able to use models to make more specific claims about the natural world. In such instances, Massimi's analysis seems correct.

The notion of exploration that Gelfert uses is more applicable to the GARD case because he allows for a different range of epistemic uses, including some involving greater epistemic uncertainty (i.e. uncertainty about what the target phenomenon is and how to study it). He argues, using the example of Turing's model, that some models that initially seem to have no target can still be fruitful and might have a target in the future. Such a model might allow scientists to give a proof of concept or demonstrate the suitability of a method, for example. Such epistemic tasks (and others) may not require that a model have a target, though such a model may in the future have a target. Gelfert argues that rather than understanding models as representational, we should understand them as epistemic tools that allow scientists to explore (2018, pages 254 to 255), a proposal that synergizes neatly with the GARD model. This proposal does not involve the ascription of modal or other properties to actual world phenomena. Put another way, scientists are in too epistemically weak a position to learn much specific about the actual world.

Gelfert is claiming that exploration, even if it does not provide us much knowledge, should be valued because an exploratory model may later become empirically confirmed or relevant. That is, the possibility that a model might one day have a target indicates that the model is worth using because the risk is outweighed by the possible benefit; the risk would be that the model does not have a target (and so we learn nothing, or at least we only learn that there was no target) and the possible reward is that the model does have a target after all and we are then able to learn something about that target.

This works quite well for a fairly abstract mathematical model, like Turing's, that can in principle be applied to a variety of patterns. For such cases, Gelfert rightly argues that exploratory fruitfulness should be a criterion whereby some models are evaluated. However, we must be careful about how we should use this criterion and what we should conclude about a model that satisfies it. What we should not conclude is that based on explanatory fruitfulness, a model will later be empirically vindicated. In other words, an exploratorily fruitful model sits quite far away from giving us knowledge. Giving us knowledge is not their central purpose, but the possibility that they might is one reason we continue to find such models interesting. On the other hand, it is also very difficult to reject a model on empirical grounds if that model is exploratory and hence not very answerable to empirical constraints. Hence, we should be very restrained in our expectations about what we expect exploratory models to give us, a point that Gould made about explanations with weak empirical constraints (1978).

So even though the GARD model certainly does involve some of the characteristics that Gelfert observes, it should prompt us to place emphasis on the epistemic status of the target that is not explicit is his account. We need this emphasis because we might wonder why prefer a model without a target and what is it about having no clear target that makes an exploratory model epistemically valuable? What seems interesting to me about the GARD case is not the possibility that it might one day acquire a target (which is partly what is valuable about the cases Gelfert considers), but rather that it is unclear what we are going to learn from the model and what relationship it will have with the world in the future. This indefiniteness has not been emphasized very much in the literature, even in the exploratory modeling literature, and I think it is crucial to understanding not only the GARD model, but probably other exploratory models as well. The indefiniteness is a direct result of the fact that it is currently impossible to know whether the model has a target.

3.1 Indeterminate Modeling

Other discussions of exploratory models have touched on uncertainty briefly, but I would like to confront the indeterminacy of exploratory models more directly and place greater emphasis on it. Doing so will allow us to better understand the epistemic uncertainty associated with origin of life studies and with the use of models that do not have a clear application to the actual world, i.e. that do not provide analogical reasoning or how-possibly explanations (types of explanation that we will return to in 4.1). The purpose of appealing to indeterminacy is not to establish a new type of model that is distinct from existing classifications of models, it is to place emphasis where it is needed. By doing so, we can gain insight into the role of targets in exploratory modeling and we can understand better what expectations we should have for exploratory models. Let me clarify what I mean by a model being indeterminate before describing its value.

A model is indeterminate iff:

- The model is without a clear target and
- We are not in a position to rule out the possibility that it can have an application to the real world one day and
- The model does not allow for modal or analogical reasoning.

These are quite negative claims and suggest strong limits on what an indeterminate model can do. Indeed, some of the more obvious kinds of epistemic value we associate with modeling appear precluded by these characteristics of indeterminate models as I have defined them. What, then, is their value?

3.2 The Value of Indeterminacy

The value of an indeterminate model lies in way scientists can use it to structure future research and by doing so, provide a structure to the space of uncertainty that might be associated with understanding a particularly obscure phenomenon (such as the origin of life). Using a model to structure research and uncertainty can happen in a variety of ways. It can point to specific pieces of information that are especially important, it can identify types of data that would be particularly valuable, it can recast old evidence to offer new relevance, it provides a context or a motivation for future experiments, it can be used to draw connections between seemingly disparate fields or disciplines, and on and on. These are crucially important tasks that can be particularly hard to perform when we know very little about the phenomenon in question and do not have a clear path forward to address that uncertainty. Also crucially, these tasks are best accomplished through indeterminate modeling because if we knew more about the target, more specific kinds of reasoning about that target would be possible. That is, if we knew the model had a target, we would be in a position to establish a more robust and specific line of research. Given that we cannot travel back in time to observe the emergence of life on Earth, how can we best study it? Establishing a research program about such a topic is not straightforward, especially when faced with limited funds, time, and uncertainty about not only what it is we are studying, but also how to study it. The GARD model exemplifies this by helping scientists confront uncertainty in four ways.

> 1. The GARD model offers an avenue for constraining uncertainty by suggesting how the problem area (the origin of life) might be conceptualized. In conceptualizing a problem area in a particular way, the model helps scientists identify what further evidence is needed and where to look for that evidence.
> 2. The model offers ready interpretive resources that can be brought to bear on discoveries or new results that may otherwise not have appeared relevant to the origin of life. This is particularly relevant to the Enceladus flyby that provided so much data bearing applicability to the primordial soup question (discussed more below).
> 3. The GARD model's ability to identify and open lines of research is what I think Lancet et al. meant when they claimed the GARD model offered a “proof of concept,” an expression that cannot be literally interpreted. There is no actual proof involved in the GARD model and it is a vast overstatement to say that using the GARD model somehow proves that life emerged from lipids or that lipids form protocells. Instead, we should take a light reading of “proof” and interpret this model as a tool that provides justification for thinking that a lipid world research program is fruitful. This does not amount to endorsing lipids as the origin of life, but rather the potential fruitfulness of assuming that life emerged from lipids. If origin-of-life researchers make this assumption with the understanding it may not be turn out to be true, a number of research avenues then become interesting. Without the model offering some plausibility to the lipid world hypothesis, those avenues of research and lines of interpretation would not appear interesting or relevant for our understanding of the origin of life. It gives an affirmative answer to the question “is this line of research worth pursuing?” This is related to, but distinct from, the first way the model helps scientists listed above because the proof of concept lies at a more general level of research: one must establish the plausibility of a project before considering what data to pursue.

Indeterminacy lies at the heart of how the GARD model points scientists forward. The model shows us that we need to be looking elsewhere for further information that could help us establish whether lipids were the precursors to life, i.e. we need to establish *whether* this model has a target. If we already knew the model had a target, then there would be no point in looking for evidence to establish that it does. There would also be no point in motivating the line of research because it would already be clear what to study. On the other hand, if we knew the model had no target, then there would also be no point in looking for more evidence because there would be no point in trying to find support for a rejected hypothesis. Indeterminacy, therefore, is essential to understanding how this model stimulates research and offers epistemic value. To flip Gelfert's expression around, this is not a model in search of a target, but a model built and used intentionally to help researchers determine if there is a target.

There is support for this emphasis on indeterminacy in the scientific practice. For many years it seemed unlikely that a primordial soup was present in Earth's distant past, which is a prerequisite for life to develop from lipids. This seemed catastrophic for the lipid-world hypothesis, suggesting any model that supposed lipids formed protocells would have no target. However, Kahana, Schmitt-Kopplin, and Lancet (2019) recently noted that Saturn's moon Enceladus probably has liquid water beneath its ice cap, possibly filled with a variety of organic molecules; these may be signs of a primordial soup. Enceladus therefore provides very distant support for the lipid-world hypothesis because it now seems more plausible that there could have been a primordial soup on Earth (because if the soup has formed once in our galaxy, it may be more likely to have formed twice). If there were a primordial soup on Earth, the conditions may have been right for lipids to self-assemble. If this is born out, that Earth did once have a primordial soup, then it is possible the GARD model does have a target, which would be protocells in that soup. Crucially, it may turn out that the model is representational; the possibility that the model represents provides motivation, the promise of results, and it indicates that the model is not merely heuristic.[6]

What, one may ask, does this add to the discussion about exploratory models? I think it shows several things. The first is that there is nuance we may not have expected when it comes to models without targets. That is, one cannot easily lump models with targets into one category and models without them into another. Some models might have targets and the scientific project of attempting to establish *whether* they do is the entire point of the model. This epistemic difficulty is what can lead scientists to build that model in the first place.

The second thing this discussion shows is that given the nature of the GARD model, it is unlikely we would have much basis for rejecting it because the empirical constraints on it are very weak. It may, for instance, require confirming a rival hypothesis, such as the RNA world hypothesis, before we could really be in a position to evaluate the GARD model. But given the difficulty of this task, such an evaluation would be most difficult indeed. In other words, indeterminacy suggests we may not ever really be able to confirm or reject the potential explanations this model can offer, hence the connection between Gould's just-so explanations and exploratory modeling: we should not mistake exploratory interest for knowledge of phenomena. This is not necessarily a criticism of the practice that uses the GARD model, but a note about what expectations we might place upon it.

Finally, this discussion shows with some more precision how modeling can help in cases with a high degree of uncertainty. This stands in contrast to cases of modeling that involve answering fairly specific research questions or calculating results. Such cases involve some uncertainty of course, but not very much and models can help close fairly small gaps in our understanding. In more contexts with vast degrees of uncertainty, the function of modeling is a bit different. To show this contrast more sharply, let's examine why the GARD model is not easily captured by some important accounts of modeling without targets that are better suited to cases with more certainty.

4 Modeling Without Representing

Thinking of the GARD model as indeterminate and exploratory captures its epistemic value well because of the epistemic uncertainty associated with the phenomenon of interest: the origins of life. But there are lingering questions about how an exploratory model that is indeterminate is related to several other types of models discussed in the literature. In this section, I would like to clarify the similarities and differences between what philosophers have said of other models and what I have here said about exploratory models like the GARD model. There are two key points. First, some extreme examples of exploratory models, because of the epistemic uncertainty associated with them, do not allow for analogical or modal reasoning, yet. This stands in contrast to thinking that it is the representation relationship that allows us to classify the model. But instead, it is the uncertainty, or the epistemology of the phenomenon of interest, that distinguishes types of models. Second, the metaphysics of exploratory models can reasonably be captured by thinking of them as artifacts, though whether they are also fictions is less clear.

6 The model is not heuristic if by the term we mean a tool that is merely convenient for achieving some task, but which we think has no truth or real-world significance behind it. However, heuristic models are not well defined, so there is scope for taking a different position on this matter if one were to take an alternative stance toward what it means to be heuristic.

4.1 Modality and Analogical Reasoning

Philosophers have more recently recognized that some models are not best understood in terms of target systems; Massimi (2019), Knuuttila and collaborators (Knuuttila, 2009; Knuuttila, 2011; Knuuttila and Loettgers, 2017; Gelfert, 2018; Frigg and Nguyen, 2016), Gelfert (2018), and Bokulich (2011; 2012) are just a few good examples. Such models, setting aside the discussion of exploratory modeling for now, resemble the GARD model because of this lack of target. Some of these analyses are most applicable to models that tell us something about the modal properties of real world target systems. So even though the model may be heavily fictitious and not directly represent, the model is meant to allow us to conceive how the world might be and in so doing to ascribe modal properties to the features of the actual world. Models of this kind therefore give us modal knowledge and modal knowledge is very useful. Scientists can use it to then design more thoughtful and promising experiments, constrain theories, etc. Massimi writes:

> Fictional models ultimately deliver important modal knowledge about what is causally possible concerning a phenomenon under study [...], by attributing to a fictional entity some essential properties and relations [...], which via suitably keying are translated into essential properties and relations [...] imputed to the phenomenon (Massimi, 2019, page 5).

Crucial to the use of fictional models of this sort is a phenomenon. Even though a fictional model has a non-existent target, the use of the model nevertheless allows scientists to attribute real properties to actual, real-world phenomena. There are strong parallels between this and the indirect representation that Godfrey-Smith (2009) and Frigg (2010) describe. They argue that fictional models represent directly a fictional model system, which in turn represents some real world target. The model consequently represents a target indirectly via a fictional system. We might then say that for this general conception of fictional models, a model allows us to make claims about real world things, to solve real world problems.

The GARD model is not of this type because it does not offer much help, at this point, with any real world phenomena: it is unclear what sort of modal knowledge we would get from it and thus how we could use it to ascribe modal properties to phenomena. Nevertheless, there are two candidates for real world phenomena that one might suggest the GARD model helps us understand (thus supporting the thesis that the GARD model is a fictional model that supports analogical reasoning). One phenomenon might be life or its origin; the other might be lipid behavior. The thinking here would be that the GARD model allows scientists, through analogical reasoning, to attribute real properties (perhaps modal properties) to the actual origin of life or to actual lipids.

These two candidates are not well-suited to the GARD model because the model has not actually allowed for the attribution of properties to the origin of life. Furthermore, it is not certain that life ever emerged from lipids at all, so the first option (that the GARD model allows for reasoning about the origin of life) is very tenuous. The second option is little better. It is not clear whether lipids as described by the model could have any of the properties indicated by the model. Rather than tell us modal properties of real entities (as the fictional model literature suggests for some models), the GARD model tells us about hypothetical properties of hypothetical entities. The case would be different if we had stronger reasons for thinking that lipids with properties similar to those specified in the model exist, or if the model allowed scientists to attribute properties to actual lipids. Neither is the case. It is possible that someday we will have these reasons and the model will provide via analogical reasoning modal knowledge of either life, its origins, or the behavior of lipids. We are not there yet.

Nevertheless there is a general sense in which the analyses discussed in this section are relevant. Exploratory models of the kind Massimi and Knuuttila examine serve to help scientists to conceive what is or is not possible. The GARD model in *very general* terms does this, but it is less specific than some of the aforementioned philosophers' case studies because upon using the model considered here, scientists do not necessarily have a clear conception of what is and is not physically possible. Much further work would be needed to establish such a conception, if it is indeed possible to do so for a phenomenon as obscure as the origin of life.

There is another type of epistemic value associated with modeling that is closely related to analogical reasoning. The epistemic value is a type of explanation called a *how-possibly explanation*, in contrast to *how-actually explanations*. The difference hinges on whether the explanation explains a phenomenon (how actually) or whether it merely offers a possible explanation (how possibly). Philosophers and scientists have argued about this distinction and whether how possibly explanations are really worthwhile. Stephen Jay Gould originally criticized what might have been how-possibly explanations by showing they are really just-so stories: stories that we make up that have little empirical constraint and therefore cannot be easily refuted (1978). Philosophers took up a defense of these types of explanations in the context of evolution in biology (Brandon, 1990; O'Hara, 1988). More recently, Forber (2010) and Bokulich (2014; 2011) have argued that it is not just evolution that lends itself to how-possibly explanations; Bokulich argues models give us how-possibly explanations and such explanations are valuable because they tell us something about how or why a phenomenon might have come about. Forber claims that how-possibly explanations help scientists identify a set of possible explanations for a phenomenon, which they can then use to focus their research. More precisely, this type of explanation gives us modal information about the target phenomenon, as Verreault-Julien has argued (2019). Because how-possibly explanations ascribe epistemic value to models in terms of modality, it is not different in kind from the analogy accounts discussed previously: both accounts suggest models allow us to ascribe modal properties to target phenomena and this does not seem to be what the GARD model is doing at this point in time.

There is another way of thinking about how-possibly explanations that matches what the GARD model does somewhat more accurately, but which does not capture the epistemic uncertainty. Drawing from Forber and Brandon, Reydon (2012, page 307) argues that Brandon's how-possibly explanations are not really genuine explanations, but instead heuristic. When one gives an " heuristic explanation" of this kind, one offers speculative initial conditions and then uses laws or generalizations to generate predictions.[7]

So one does not provide an explanation of a phenomenon (because we are unsure of the initial conditions), but one does provide a kind of investigation of an explanatory landscape by establishing, given the relevant laws, what would happen if such-and-such initial conditions were actual. Applied to the GARD case, this framework suggests scientists are offering speculative initial conditions (that lipids of the relevant kind existed in a primordial soup) and then generating predictions using generalizations, laws, and a Monte Carlo simulation.

Although this is generally the correct picture, the GARD model is quite far removed from allowing scientists to provide an exploration in the way Reydon discusses. The reason is that not only are scientists unsure of the initial conditions, they are unsure of the explanandum. That is, the actual phenomenon to be explained is the origin of life, whereas the GARD model provides an exploration of assemblies of lipids. The assembly could be a precursor to the origin of life, but it is not an actual phenomenon as far as we know. Therefore, how-possibly explanations à la Reydon capture some of the logical structure, but not the epistemic uncertainty, associated with using the GARD model.

This general discussion of modality and analogical reasoning with models suggests that there is a complex taxonomy of targetless modeling that is structured by uncertainty and epistemology: models without targets have variable epistemic value and use depending upon the uncertainty associated with the phenomenon scientists are trying to understand. For cases with more certainty, scientists are able to use models to reason analogically (as the above authors have suggested). For cases with less certainty (such as for the GARD model), scientists might know too little to ascribe modal properties to a phenomenon of interests and they must instead use models to manage uncertainty and structure future research.

4.2 Metaphysics and Models as Artifacts

There is another way of thinking about models that places less emphasis on modality and the target phenomenon: the artifactual approach (Knuuttila, 2010; Knuuttila, 2011; Knuuttila and Loettgers, 2013a; Knuuttila and Loettgers, 2013b; Knuuttila and Loettgers, 2016; Knuuttila, 2017; Knuuttila and Koskinen, 2020; Knuuttila, 2005). This approach is fruitful and makes important advances over the imaginary approach, particularly in terms of ontology and in terms of the manipulatability of models. But there is in my mind a tension in thinking of the GARD model as as artifact, at least if one is also committed to the fictionality of models. This suggests that the GARD-model-as-artifact is a useful analysis, but we might need to say something else—or something more precise—about the fictionality of this model.

The artifactual approach holds that we should treat models as artifacts, i.e. man-made objects that have physical manifestation, regardless of whether it is manifested in a computer hard-drive, on a sheet of paper, or in three dimensional space. The medium of a given model has important consequences, and places different limitations upon, the ways scientists can manipulate and use models.

> From the artifactual perspective models are like any other artifacts in that they are human-made objects intentionally produced for some purposes within the sphere of particular human activities (Knuuttila, 2017, page 11).

Models are like tools that scientists can use to do a variety of things. Sometimes a model can be used for a particular task for which it was built, but it can also be used in new ways not envisioned by the initial creators (Morrison, 1999; Morrison and Morgan,

7 Reydon and the other authors here are working within a deductive-nomological framework (Hempel and Oppenheim, 1948). In this framework, one uses initial conditions and covering laws to generate predictions. One achieves an explanation for a phenomena when one is able to use those initial conditions and laws to accurately predict that phenomenon.

1999). This account downplays the importance of representation and places greater value on the materiality and upon the use of models.

In many ways this approach fits very nicely with the GARD model, which is used as a sort of tool for several epistemic tasks. It also physically manifests, such as in the form of an equation, and that manifestation shapes how scientists can and cannot use it. It can, for instance, be used in mathematical simulations, but cannot be measured using rulers. One could write an article about the GARD model as artifact and in so doing describe in greater detail the various ways scientists manipulate it, but this is a separate project because our interest here lies in the relationship this model has to the actual world, i.e. its epistemic value. The relationship is often understood in terms of fictions, but we should have a reservation about this use of fiction.

The connection between artifacts and fictions is not obvious, but Knuuttila appeals to Thomasson's fictions-as-artifacts account (1999)—instead of the appeal others (Frigg and Nguyen, 2016; Toon, 2012; Toon, 2016; Salis and Frigg, 2020; Godfrey-Smith, 2009) have made to Walton's theory of fiction (1990). Knuuttila writes:

> When a model functions as a fiction in scientific research, its truth (or falsity) is bracketed away from consideration. On the other hand, in some other uses, the very same model can be regarded as a representation of some specific target system and employed to generate empirically testable inferences concerning it. But then it is not treated as a fiction anymore. No model is a fiction in and of itself, it gains a fictional status in those uses in which the primary interest of scientists is to study the model system apart from any determinable representational ties to external target systems (2010, page 18).

A model is a fiction if, and when, it is used as a fiction, i.e. without a link to the actual world. A consequence of this is that truth and falsity are not relevant (is bracketed), which sits well with the GARD model. But we might have a worry about how fictional and artifactual analyses would work together to give us an understanding of models. The metaphysics and pragmatics appear ill-suited to one another. For example, literary fictions do not have an ever-changing identity that varies with use, i.e. that something is a fiction does not depend upon how I use it. Literary fictions are always fictions and in the quotation above, Knuuttila suggests a given model may or may not be a fiction, depending on use. Regarding ontology, are models artifacts, are they fictions, how can they be both, and how can they be only fictional some of the time? The original point in characterizing models as fictions was to help us understand their ontology (Godfrey-Smith, 2006), but this is also the main task for artifacts. In non-philosophical discourse, we make a distinction between artifacts and fictions. The fact we do so does not mean that artifacts and fictions cannot be used together to help us understand models in philosophical discourse, but the ontology is quite complicated and the issue needs more attention.

Models without targets, we expect, might one day have real relevance to the actual world, but we would never have this same type of expectation about a fiction. I have never expected to see Anna Karenina walking down the street, but this epistemic pessimism about seeing Anna Karenina is a point of difference between some models and fictions. I will never, at any point in the future, expect to see Anna Karenina walking down the street. However, some types of model that we think currently lack a target may someday turn out to have a target.

One further point I would like to attend to here is the status of the model. For the artifactual approach, the status as fiction can change depending upon how the model is used. If the model is used to represent a target, it may not be best understood as fictional. If scientists are not using the model representationally, then it may have fictional status. Now the implication here is that the use of the model tells us definitively whether the model has fictional status, i.e. whether it has a target and therefore what its connection to the world is. The reason I mentioned above that the fictional element to the artifactual approach does not help for this model is that it is not clear whether the GARD model has a target; this is part of the reason why it is interesting. Therefore, we do not know whether it is fictional or not, even if it is an artifact, put in terms of this framework.

5 Concluding Remarks

This article used a case study from the origin of life studies to further our understanding the epistemic value of exploratory models. That case study, focused on the GARD model, suggests that in cases of high uncertainty associated with a target phenomenon and the methods used to study it, scientists use exploratory models to manage uncertainty. That management involves structuring future research, offering a mode for conceptualizing the target phenomenon, offering an interpretive framework for existing and future evidence, and demonstrating the promise of an uncertain line of research. Some exploratory models are able to offer these benefits because it is unclear whether they have a target; it might turn out that the model does have a target, but it might turn out otherwise. Such models with uncertain targets I call indeterminate. It is essential, not accidental, to the workings of such models that it is unclear whether they have a target: the value stems from their indeterminacy.

Because of the extreme uncertainty, analogical reasoning and how-possibly explanations are not the right tools for understanding this highly speculative project. Artifactual and exploratory modeling are better lenses for thinking about the GARD model, but this case prompts us to place an emphasis on the uncertainty, or indeterminacy, of modeling in the origin of life studies that has not been emphasized elsewhere. I introduced this idea in 1 and developed it further in 3.1. Although this distinction between models with and without targets is sharp when it comes to the GARD model (described in 2 and 4), there may be many vague cases. More generally, we might say that for models without targets, it is unclear what status they should have and this lack of clarity can come in degrees. Depending upon the uncertainty, different epistemic tasks will be possible. Science is an inquiry into the natural world and scientific models should be evaluated based in part on how effectively they allow us to learn about that world. Those models that have not yet done so are not necessarily useless, indeed they may be very useful for other epistemic tasks, such as directing future research and providing interpretive frameworks on new evidence. For these kinds of tasks, indeterminacy, the uncertainty associated with a model's target, can be an asset. Such cases suggest that the modeling abilities of scientists have outrun the observational and experimental aspects of science and these two latter aspects need time and investment to catch up, a process that indeterminate models might accelerate. Will further work on observation and experiment vindicate the GARD model in particular? Time may tell, but we do not know if it will.

References

Armstrong, D., Markovitch, O., Zidovetzki, R., and Lancet, D. (2011). "Replication of simulated prebiotic amphiphile vesicles controlled by experimental lipid physicochemical properties." *Physical Biology,* 8(6), 66001.

Bernhardt, H. (2012). "The RNA world hypothesis: the worst theory of the early evolution of life (except for all the others)." *Biology direct* 7(23).

Bokulich, A. (2011). "How scientific models can explain." *Synthese,* 180(1), 33--45.

Bokulich, A. (2012). "Distinguishing explanatory from nonexplanatory fictions." *Philosophy of Science,* 79(5), 725--737.

Bokulich, A. (2014). "How the Tiger Bush Got Its Stripes: How Possibly' vs. How Actually Model Explanations." *Monist,* 97(3), 321--338.

Brandon, R. (1990). *Adaptation and Environment.* Princeton University Press.

Contessa, G. (2011). "Scientific models and representation." in French, S. and Saatsi, J. (ed.s) *The Continuum Companion to the Philosophy of Science*. Continuum International Publishing Group.

Fisher, G. (2006). "The autonomy of models and explanation: anomalous molecular rearrangements in early twentieth-century physical organic chemistry." *Studies in History and Philosophy of Science Part A,* 37(4), 562--584.

Forber, P. (2010). "Confirmation and explaining how possible." *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences,* 41(1), 32--40.

French, S. (2003). "A model-theoretic account of representation (or, I don't know much about art but I know it involves isomorphism)." *Philosophy of Science,* 70(5), 1472--1483.

French, S. (2014). *The structure of the world: Metaphysics and representation.* Oxford University Press.

Frigg, R. (2010). "Fiction and scientific representation." Frigg, R. and Hunter, M. (ed.s) in *Beyond mimesis and convention: Representation in art and science.* 97–138. Springer

Frigg, R. and Nguyen, J. (2016) "The fiction view of models reloaded." *The Monist* 99(3), 225--242.

Fumagalli, R. (2015). "No learning from minimal models." *Philosophy of Science,* 82(5), 798--809.

Fumagalli, R. (2016). "Why we cannot learn from minimal models." *Erkenntnis,* 81(3), 433--455.

Gelfert, A. (2018). Models in Search of Targets: Exploratory Modelling and the Case of Turing Patterns. In: Christian, A., Hommen, D., Retzlaff, N., Schurz, G. (eds) *Philosophy of Science. European Studies in Philosophy of Science,* 9.

Giere, R. (2004). "How models are used to represent reality", *Philosophy of science*, 71(5), 742--752.

Giere, R. (2010). "An agent-based conception of models and scientific representation." *Synthese,* 172(2), 269.

Godfrey-Smith, P. (2006). "The strategy of model-based science", *Biology and philosophy,* 21(5), 725--740.

Godfrey-Smith, P. (2009). "Models and fictions in science", *Philosophical studies,* 143(1), 101--116.

Gould, S. (1978). "Sociobiology: the art of storytelling", *New Scientist,* 80(1129), 530–533.

Hempel, C. and Oppenheim, P. (1948). "Studies in the Logic of Explanation", *Philosophy of science,* 15(2), 135--175.

Kahana, A., Schmitt-Kopplin, P., and Lancet, D. (2019). "Enceladus: First observed primordial soup could arbitrate origin-of-life debate", *Astrobiology,* 19(10), 1263--1278.

Knuuttila, T. (2005). *Models as epistemic artefacts: Toward a non-representationalist account of scientific representation* in Philosophical Studies from the University of Helsinki, 8.

Knuuttila, T. (2009) "Some consequences of the pragmatist approach to representation", in *EPSA epistemology and methodology of science,* 139--148.

Knuuttila, T. (2011) "Modeling and representing: An artefactual approach to model-based representation", *Studies in History and Philosophy of Science Part A,* 42(2), 262--271.

Knuuttila, T. (2017). "Imagination extended and embedded: artifactual versus fictional accounts of models", *Synthese*, 198, 5077–5097.

Knuuttila, T. and Koskinen, R. (2021). "Synthetic fictions: turning imagined biological systems into concrete ones." *Synthese, 198*(9), 1--18.

Knuuttila, T. and Loettgers, A. (2013). "Synthetic biology as an engineering science? analogical reasoning, synthetic modeling, and integration." in *New Challenges to Philosophy of Science,* 163--177.

KKnuuttila, T. and Loettgers, A. (2013). "Synthetic modeling and mechanistic account: Material recombination and beyond." *Philosophy of Science,* 80(5), 874--885.

Knuuttila, T. and Loettgers, A. (2016). "Contrasting Cases: The Lotka-Volterra Model Times Three" in *The philosophy of historical case studies. Springer,* 319, 151-178.

Knuuttila, T. and Loettgers, A. (2017). "Modeling as indirect representation? The Lotka-Volterra model revisited", *The British Journal for the Philosophy of Science,* 68(4), 1007--1036.

Lancet, D., Zidovetzki, R., and Markovitch, O. (2018) "Systems protobiology: origin of life in lipid catalytic networks", *Journal of The Royal Society Interface,* 15(144), 20180159.

Lancet, D., Segrè, D., and Kahana, A. (2019). "Twenty Years of Lipid World: A Fertile Partnership with David Deamer", *Life,* 9(4), 77.

Lanier, K. and Williams, L. (2017) "The origin of life: models and data." *Journal of molecular evolution,* 84 (2-3), 85--92.

Massimi, M. (2018). "Perspectival modeling." *Philosophy of Science,* 85(3), 335--359.

Massimi, M. (2019). "Two kinds of exploratory models." *Philosophy of Science* 86(5), 869--881.

Morrison, M. (1999). "Models as autonomous agents." *Ideas in Context,* 52, 38--65.

Morrison, M. and Morgan, M. (1999). *Models as mediators : perspectives on natural and social sciences.* Cambridge University Press.

Neveu, Marc, Kim, Hyo-Joong, and Benner, Steven A. (2013). "The strong RNA world hypothesis: Fifty years old." *Astrobiology*, 13(4),391--403.

O'Hara, R. (1988). "Homage to Clio, or, toward an historical philosophy for evolutionary biology." *Systematic Zoology*, 37(2), 142--155.

Poznic, M. (2016). "Representation and Similarity: Suárez on Necessary and Sufficient Conditions of Scientific Representation."

Journal for General Philosophy of Science.

Reydon, T. (2012). "How-possibly explanations as genuine explanations and helpful heuristics: A comment on Forber." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences,* 43(1), 302--310.

Salis, F. and Frigg, R. (2010). "Capturing the scientific imagination." *The scientific imagination,* 17--50.

Segré, D., Lancet, D., Kedem, O., and Pilpel, Y. (1998). "Graded autocatalysis replication domain (GARD): kinetic analysis of self-replication in mutually catalytic sets." *Origins of Life and Evolution of the Biosphere,* 28(4-6), 501--514.

Segré, D, Pilpel, Y, and Lancet, D. (1998). "Mutual catalysis in sets of prebiotic organic molecules: Evolution through computer simulated chemical kinetics." *Physica A: Statistical Mechanics and its Applications* 249(1-4), 558--564.

Segré, D., Ben-Eli, D., and Lancet, D. (2000). "Compositional genomes: Prebiotic information transfer in mutually catalytic noncovalent assemblies." *Proceedings of the National Academy of Sciences,* 97(8), 4112 LP -- 4117.

Segré, D., Ben-Eli, D., Deamer, D., and Lancet, D. (2001). "The Lipid World", *Origins of life and evolution of the biosphere : the journal of the International Society for the Study of the Origin of Life,* 31, 119--145.

Suárez, M. (2010). "Scientific Representation." *Philosophy Compass,* 1(5), 91--101.

Suárez, M. (2003). "Scientific representation: Against similarity and isomorphism." *International Studies in the Philosophy of Science,* 17(3), 225--244.

Suárez, M. (2004). "An inferential conception of scientific representation." *Philosophy of science,* 71(5), 767--779.

Thomasson, A. (1999). *Fiction and metaphysics.* Cambridge University Press.

Toon, A. (2012). *Models as make-believe: Imagination, fiction and scientific representation.* Springer.

Toon, A. (2016). "Imagination in scientific modeling", in Kind, A. (ed.). *The Routledge Handbook of Philosophy of Imagination.* Routledge.

Verreault-Julien, P. (2019). "How could models possibly provide how-possibly explanations?" *Studies in History and Philosophy of Science Part A,* 73, 22--33.

Walton, K. (1990). *Mimesis as make-believe: On the foundations of the representational arts.* Harvard University Press

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

van Fraassen, B. (2008). *Scientific Representation.* Oxford University Press.