

THE ROLE OF TARGET SYSTEMS IN SCIENTIFIC PRACTICE

Alkistis Elliott-Graves

A DISSERTATION

in

Philosophy

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2014

Supervisor of Dissertation

---

Michael Weisberg

Associate Professor of Philosophy

Graduate Group Chairperson

---

Michael Weisberg, Graduate Chair Philosophy Department

Dissertation Committee

Cristina Bicchieri, Patterson Harvie Professor of Social Thought and Comparative Ethics

Elisabeth Camp, Associate Professor of Philosophy Rutgers University

Daniel J. Singer, Assistant Professor of Philosophy

THE ROLE OF TARGET SYSTEMS IN SCIENTIFIC PRACTICE

COPYRIGHT

2014

Alkistis Tania Elliott-Graves

This work is licensed under the

Creative Commons Attribution-

NonCommercial-ShareAlike 3.0

License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/2.0/>

## ABSTRACT

### THE ROLE OF TARGET SYSTEMS IN SCIENTIFIC PRACTICE

Alkistis Elliott-Graves

Michael Weisberg

Scientists often construct simplified and idealized models in order to study complex phenomena. Yet they do not model a phenomenon in its entirety but *target* only the aspects of the phenomenon which they consider relevant. Hence, the model is said to describe the target system and not the whole phenomenon. The term ‘target system’ has become popular in the philosophy of science, yet most authors do not provide a definition or analysis of the concept. The result is that the term is used ambiguously, which has undermined its potential value and usefulness for scientific practice. The aim of this dissertation is to provide a cogent account of target systems and their importance in science, with examples taken from case studies in ecology. The central issue I explore in my dissertation concerns the nature of target systems. What are target systems? How are they specified? How can they be evaluated? In my dissertation I give an account of target systems as real parts of systems in the world, which are specified through a process of partitioning and abstraction. I also provide a tentative theory of target system evaluation based on the notion of aptness for a particular scientific purpose.

A deep understanding of nature and function of targets can resolve problems in science. I use the term ‘target system analysis’, to denote the specification of target systems of one enquiry and the comparison of targets across enquiries. The last part of the dissertation is devoted to the application of the theory of target system specification and evaluation to a case study from actual scientific practice, invasion biology. Target system reveals that a scientist constructing a unificatory framework in invasion biology faces a tradeoff between generality and predictability. A truly unified framework must incorporate a multitude of different causes of invasion, yet the causes of each invasion are unique. Hence, invasion biology can have a unified theory, based on the *process* of invasion, yet this theory will be of little use to predicting particular invasions.

## Table of Contents

List of Illustrations	v
Chapter 1 Introduction .....	1
Chapter 2 What is a Target System? .....	17
Chapter 3 Target System Metaphysics - Part 1 <i>Partitioning</i> .....	39
Chapter 4 Target System Metaphysics - Part 2 <i>Abstraction</i> .....	61
Chapter 5 Evaluating Target Systems - <i>Aptness</i> .....	84
Chapter 6 Case Study - Can there be a Unified theory of Invasion Biology? .....	112
Chapter 7 Integration and Generality - Analysis of Target Systems in Invasion Biology .....	139
Chapter 8 Conclusion .....	156
Bibliography .....	165

## List of Illustrations

1.1	Per capita population growth rate and Allee effect	3
1.2	Demonstration of the Allee effect in the VIM	4
1.3	Logistic Growth Equation	5
1.4	Demographic composition of historical and recent VIM study populations	6
1.5	Map of Vancouver Island	7
2.1	Partitions	20
2.2	Abstractions	22
2.3	Target System Specification	23
2.4	Exponential Growth Equation	28
3.1	Spatial representation of Natural Kinds views	47
6.1	Merging the various types of invasion studies	122
6.2	The Synthetic Meta-Framework (SIM) for invasion biology	125
6.3	The Hierarchical framework for invasion biology	128
6.4	The unified framework for invasion biology	131
7.1	Correlation between invasion success, species richness & resource variability	146

# 1. Introduction

## Marmots, Models and Targets

### 1.1 The Case of the Vancouver Island Marmots

The marmots of Vancouver Island (*Marmota vancouverensis*) are classified as critically endangered. It is estimated that their population has dropped 80%-90% since the 1980's and currently consists of roughly 200 individuals (Brashares *et al.* 2010). These rodents, which are closely related to squirrels, are quite special, as they are highly social. They live in groups of 5-15 individuals, have an intricate pattern of social interactions and a variety of alarm calls. One of these alarm calls has given rise to their nickname 'whistle pigs' by locals of Vancouver Island. They spend most of their time socializing, as they forage communally, taking it in turns to look out for predators, and share in the upkeep of their burrows. When a marmot out foraging encounters another marmot, they greet each other by touching their noses together, while the youngsters, who stay on with the family until they are around 2 years old, spend most of their time play-fighting.

Usually, the cause of a population's dramatic drop in size, is human activity. This can include disturbance of the habitat, active extermination of the population or other indirect effects, such as the introduction of invaders (predators or competitors). In this case, the situation seems more complicated. While there is some disturbance of their habitat, through logging, this actually led to a temporary increase in the growth of the population, as clearings are favored by marmots. This is because the absence of thick roots makes it easier to build their burrows, and also there is less vegetation which provides cover for predators. However, it seems that after this temporary increase, the population dropped drastically and kept on dropping.

Another reason why populations can go into decline is an increase in predators, or a drop in reproductive rate. In the case of the Vancouver Island marmots, there has been a small increase in predation rate and a small decrease in the rate of reproduction. This is quite strange, since a population is expected to grow when there are abundant resources for it to consume. Thus, the increase in rate of predation should be counteracted by an increase in the rate of reproduction, something has not taken place. Yet there is also something else. The population of marmots has not been decreasing steadily, but the rate of decline itself has grown over the years. This shows that there is probably another factor which is compounding the effect of the increase in predation, which itself leads to the decrease in reproduction rate.


The task of the scientists studying these marmots, is two-fold. First, they must determine the cause of this increase in the rate of decline, and second they must find a way to reverse it. In 2010, a group of scientists published a paper based on a study which involved observing the marmots over a period of three years (2002-2005) and a comparison of these observations with data from an earlier study, which took place between 1973 and 1975 (Brashares, Werner, & Sinclair, 2010).

The scientists hypothesized that the Vancouver Island marmots (VIM) were exhibiting an *Allee effect*, a positive correlation between the density of the population and its growth rate. Populations typically have a negative correlation between density and growth rate. As an environment becomes more densely populated, the growth rate of the population drops, as there are not enough resources to sustain an exponentially growing population. However, with an Allee effect, a population will decline despite an abundance of resources. This is because resources are not the limiting factor on a population but another factor is causing the a negative growth rate in the population. Allee effects are extremely dangerous for small populations as they have an increased chance of extinction (Courchamp, Clutton-Brock, & Grenfell, 1999). Allee effects can have many different causes, such as mate limitation and diminishing of cooperation or inbreeding depression (Courchamp et al., 1999; Stephens & Sutherland, 1999).

The Allee effect can be shown through particular deviation from a general model of population growth: the logistic growth model. This model describes the growth of a population which inhabits an environment with a limited supply of resources. Given this limitation, as the population size increases, the population's rate of growth decreases. In the case of an Allee effect, the growth rate of the population will decline even when the population size is also declining. Figure 1.1 shows the difference between simple

logistic growth, a weak allee effect (dotted line) and a strong allee effect (dashed line), in the case of growth per capita.<sup>1</sup>

**Figure 1.1 Per capita population growth rate (logistic, weak Allee effect, strong Allee effect)**



The difference between standard, logistic growth (unbroken line), and population growth incorporating Allee effects. The per capita (or 'specific') rate of population growth indicates whether the population is likely to increase, decrease, or remain constant. When this is zero, the population is stable, and thus where the growth curves cross the abscissa, the population is at an equilibrium point. When the specific growth rate is positive, the population will increase, and when it is negative, the population will decrease. The dotted line shows growth with an Allee effect that is evident over a large range of population sizes, but is relatively weak. The dashed line shows growth with an Allee effect which is evident over a smaller range of population sizes but is much more severe.

This hypothesis rested on the idea that social interactions are essential to the VIM survival, as socializing gives the marmots increased vigilance against predators, and also makes their foraging and resting time more efficient, as they can rely on others to be vigilant while they are foraging or resting. It also could affect the establishment and distribution of colonies, as a certain proportion of the teenagers in each generation leave their own family, meet teenagers from a different family and found a new colony in a


<sup>1</sup> The per capita growth rate is simply the population growth rate divided by the number of individuals in the population. It is not usual to represent logistic growth in the per capita form, and most people are more familiar with the sigmoid curve graph. However, the per capita growth curve is typically used in cases where there is an Allee effect, so I am keeping with tradition and illustrating with the per capita growth curves.



different place on the island. Thus the scientists named the situation leading to the VIM decline a ‘social meltdown’ (Brashares et al., 2010).

Yet in the actual study, Brashares and colleagues had to determine whether the VIM population was exhibiting an Allee effect, before they could investigate the causes of this Allee effect. That is, they had to determine whether there really was a positive relationship between density dependence and growth. As it turned out, the VIM were exhibiting an Allee effect, and the first part of the hypothesis was corroborated. Figure 2 shows the Allee effect in the VIM population (again the graph shows per capita population growth), with a very clear positive relationship between density and growth, which means that the VIM population is dropping despite the abundance of resources on Vancouver Island.

**Figure 1.2 Demonstration of the Allee effect in the VIM**



Annual counts of free-living V. I. marmots reveal strong inverse density dependence in per capita population growth from 1970 to 2007. Trend line represents least-squares quadratic fit ( $R^2 = 0.855$ ). Data exclude animals introduced from captivity. (Reprinted from Brashares et al 2010).

When philosophers of science describe a particular scientific study in detail, it is usually because they want to investigate the model or experimental setup which is being used, or because they think that the

model, experiment or the their results can be used to support a theoretical point. My interest in this case is rather different. Instead of focusing on the model itself (in this case the modified logistic growth model), I want to investigate the relationship between the model and the marmots that it represents. In doing so, I will look at the preliminary work carried out by the scientists work which was necessary for the application of the model to the actual population. In other words, I want to understand the process by which a scientist goes from the observation of a living breathing marmot, to a model which shows that the population of marmots exhibits this positive relationship with density.

The relationship between the organisms and the model is not entirely straightforward. The model contains idealizations and simplifications. For example, the upper bound of resources, known as the carrying capacity is assumed to be set, yet in nature it varies with environmental conditions. Also, the model does not consider effects of biotic interactions with other species, such as competition or predation. Finally, the model includes variables which do not have a direct correlate or analog in the real population.

One of the most important of these is  $r$  the intrinsic growth rate of the population. This is the growth rate of the population independent of resource limitations, competition, or predation. In other words, it is the average number of offspring that an individual has, when the population is at low density. In the graph, this is depicted as the slope of the curve, yet it role can be seen more clearly in the logistic growth equation (Equation 1.3). Here, the growth rate of the population ( $rN$ ) is limited by the density of the population itself, given the total number of resources, or carrying capacity ( $K$ ) contained in the system.

### **The Logistic Growth Equation**


$$\frac{dN}{dt} = rN\left(1 - \frac{N}{K}\right) \quad 1.3$$

The interesting point for my purposes, is the relationship between  $r$  and the actual VIM population. Strictly speaking, a growth rate  $r$  is not a part of the real-world system, as it is not a property of actual marmots, but a statistical variable, derived from data collected from parts of the system in the world. Marmots are born, have offspring and die, they do not have growth *rates*. How then do scientists calculate variables such as  $r$  from the marmot individuals and the events in their lives (i.e. births and deaths)?

The standard procedure for calculating a growth rate in ecology, is the construction of a *life table*, a tabular summary of the birth rates, fecundity and death rates of a population, divided into age groups (Ricklefs & Miller, 2000). These rates are calculated from the observation of the birth and death events of the actual organisms. Then, the values from the life table are then used to calculate the intrinsic growth rate of the population.

The construction of a life table might be conceptually straightforward, yet it is quite laborious and time consuming. A very interesting point is that this work is considered ‘preliminary’ work for the purposes of the model, even though it takes up most of the time of the entire study. To illustrate, the observations by Brashares and colleagues, were conducted between 2002 and 2005, while the paper was not published until 2010. Even allowing for the time it took to work on the model, and any mishaps which could have delayed the publication, the gap shows that a lot of time must have been spent on this preliminary work. Yet very little of this preliminary work appears in the actual publication (in the form of explaining the the methods of observation). Importantly, there is no mention of the intermediary steps between the observations and the input of the variables into the model.

**Figure 1.4 Demographic composition of historical and recent VIM study populations**



Fortunately, in this case, there is some mention of this preliminary work in the supplementary materials of the paper, which can be accessed online. In particular they show a graph of the demographic composition of the VIM populations at two different times: 1973-75 and 2002-4 (see Figure 1.4).

Yet there are other aspects to this preliminary work, which are usually mentioned even less explicitly, even though they are integral parts of the study itself. In general terms, these aspects are the way in which the problem itself is conceptualized, the designation of the exact spatio-temporal location in which the phenomenon takes place and the determination of the factors which are relevant for the investigation of the phenomenon. This case is rather exceptional, as some of these aspects are mentioned, even explained in the paper. First, the authors provide a figure which shows the domain in which the phenomenon is situated, shown here in Figure 1.5.

**Figure 1.5. Map of Vancouver Island from Brashares *et al.* (2010)**



**Fig. 2.** Map of Vancouver Island, Canada with stars indicating the general location of V. I. marmot subpopulations included in this study.

Second, the authors mention some of the phenotypic and behavioral traits which are important for the study of the phenomenon, within the domain. The explicit mention of some of the traits of the VIM

gives some indication of which traits the authors think are relevant for the investigation of the phenomenon in question.

The species is distinct from other marmots in skull morphology and pelage ... , behaviour (Heard 1977), and vocalizations (Blumstein 1999). VIM are diurnal, consume a wide variety of grasses and forbs (Martell & Milko 1986) and use underground burrow systems for rest and escape during the 5-month active season (May–September) prior to hibernation (October–April). ... There is a single breeding season ... one to seven pups (mean 3-4) are born after a gestation of 30–32 days.

The upshot of this example is that the preliminary work involved in setting up a model is both laborious and time-consuming, yet neither scientists nor philosophers accord it the importance that it merits. The main aim of this dissertation is to provide the conceptual space for the analysis of this preliminary work. The secondary aim of the dissertation is to show how the analysis of the preliminary work can actually help to solve some of the problems associated with scientific practice.

The conceptual tool which encompasses the preliminary work for applying a model to a real-world system is a ‘target system’ (Suarez 2003, Giere 2004, Knuuttila 2005, Frigg 2009, Godfrey-Smith 2009, Weisberg 2013). It is called a target because it encompasses those aspects of the real-world phenomenon which will appear in the model. That is, it includes the parts of the real-world phenomenon that the modeler has *targeted* as being important for the construction of the model. Thus, models represent target systems, not natural systems in their entirety.

However, just as scientists often relegate their preliminary work to the background, philosophers have spent little time characterizing target systems and their role in the relationship between models and the world (called model-world relations). This dissertation is aimed at providing a full characterization of target systems, by showing that specifying a target system involves precisely the preliminary work which scientists conduct so that a model can be applied to a real-world system.

In a nutshell, my view of target systems is that they are parts of real-world systems. The specification of a target system can be divided into four conceptually distinct parts, the *phenomenon of interest*, the *domain of study*, *partitioning* and *abstraction*. Identifying the phenomenon of interest is the determination of the precise phenomenon under investigation. In the case of the VIM, it is population growth, more specifically, a particular instance of population growth, the Allee effect. Identifying the domain of study is the location of the spatiotemporal region in which the phenomenon is taking place, in

this case Vancouver Island. Partitioning involves the categorization of the domain of study into units, given the scientific purpose designated by the phenomenon of interest and the contextual framework within which the scientist is operating. In this case, Vancouver Island was partitioned into organism-sized units (marmots and other organisms) each of which has a set of properties, such as fur length, fur color, eye color, size, sex and metabolic rate.

Yet many of the parts and properties of the partitioned system will not be relevant for the study of the phenomenon through the particular model, hence they will be abstracted, or omitted from the system. It is these parts and properties which are used to calculate variables which are found in the model and which have no counterpart in the real world system. Thus, specifying a target system of a scientific investigation is an essential aspect of the preliminary work which needs to be conducted so that a model can be applied to the real-world system in question.

In philosophical terms, a target system stands in a representational relationship to a model, as the model represents the target system. Yet it is also important to understand that it stands in a part-whole relationship with the real-world system, as it is part of that system. This means that target systems are real parts of the world. This is an important point, as it is essential for a coherent account of target systems, which allows the analysis of target systems to reach its full potential as a philosophical and scientific tool. In order to explicate this point, I will give an overview of the debate concerning model-world relations, and situate the notion of target systems and their role within the debate.

## 1.2 Model-world Relations

Modeling is an indisputably important aspect of scientific practice. Scientists create models of phenomena in the natural world when they cannot manipulate them or study them directly. Models can be found across many disciplines in the sciences and social sciences, with many diverse phenomena being studied through the use of models. In philosophy of science models and modeling have become increasingly important, with the increase in popularity of the semantic view of scientific theories. On this view, models are central to scientific theorizing, as theories themselves can be understood as clusters of models (Giere, 1988; Weisberg, 2013).

The existing literature on models and simulations can be divided roughly into three categories, though it should be noted that these categories cannot always be neatly pulled apart. Nonetheless, each

category is aimed at answering a distinct question about models and modeling. The first question is ‘what is a model?’ and focuses on understanding the nature of models. The second question is ‘what is the relationship between models and the world?’ and aims to understand how models relate to natural phenomena. The third question is ‘when is a model successful?’ and aims to determine how models should be evaluated.

These three questions are, of course, related and most accounts of modeling attempt to answer all three at once. For example, when a view takes a stand on the nature of a model, this affects what the model-world relation could be, which in turn affects how models could be evaluated. Still, pulling apart these three questions at least at the conceptual level is important, because doing so reveals there is a major gap in the answers given to one of these questions. I am referring to the issue of ‘model-world relations’, where most views do not actually give a full account of how models relate to the world.

Before I turn to the question of model-world relations, I will give an overview of the literature on the nature of models. The class of entities to which the term ‘model’ is usually attributed is quite large and heterogeneous, which includes amongst others physical systems, mathematical systems and computational systems (Giere, 2004; Weisberg, 2013). Models are often defined as some type of system or structure which can potentially represent a system in the world (Giere 1998, Godfrey-Smith 2009, Suarez 2002, Weisberg 2013). One view of models is that they are interpreted structures (Weisberg 2013), that is, they are systems which have some kind of structure and are intended by scientists to represent another system. Another view holds that models are fictional entities or imaginary systems, which are analogous to systems in the world (Frigg, 2009; Godfrey-Smith, 2008). All these views agree on the fact that models contain idealizations, and are therefore distorted analogues of the systems they represent. There are many reasons to idealize and many types of idealization, yet all idealizations have something in common. When a representation of a system contains idealizations, it will not be identical to the system itself. Models, therefore, are not identical to the systems they represent. This mismatch between model and phenomenon complicates the task of evaluating the usefulness of the model for representing a particular phenomenon. In order to deal with this complication, philosophers of science have tried to give a precise account of the relationship between a model and the phenomenon it represents.

Most accounts of the practice of modeling define models partly by their function, that is their capacity to represent other systems, usually systems in the world (Giere 1998, Godfrey-Smith 2008, Suarez

2003, Weisberg 2013). Understanding the nature of this representational relationship is the subject of ‘model-world relations’. This, in turn determines when a model is considered to be successful. Most of the accounts of model-world relations fall in the camp of the *semantic view* of models and theories. Proponents of this view think that all scientific theorizing involves modeling, as theories can be reconstructed as sets of models (Weisberg, 2007b). In many of these accounts, models are thought to be isomorphic with respect to the real-world phenomena they represent (Van Fraassen, 1980). That is, a model has structures that are analogous to the ones in the real-world system and preserves the same relations between those structures that appear in the real world system. To take a very simple example, a square table top is isomorphic to a Euclidian square because its sides are of equal length and the angles between them are right angles, just as in the Euclidian square (van Fraassen, 2008).

The problem is that these views of model-world relations do not seem to account for idealization. Idealized models do not have exactly the same structures as real-world systems nor do they always preserve the relations between those structures. This means that proponents of this view are often forced to reject models for the study of particular phenomena, because they are not isomorphic to the system they are modeling, even though they may actually be quite useful at explaining and predicting the phenomenon in question (Weisberg 2013). For example, a simple harmonic oscillator is not isomorphic to a clock’s pendulum, because even though it describes the oscillation observed in the pendulum, it does not take into account the effect of friction on the real-world system.

In order to deal with this issue, some proponents of the semantic view have relaxed the constraints slightly and view model-world relations as *partially* isomorphic (Costa & French, 2003). Here, some aspects of the model are isomorphic with respect to the real-world system but others are not. Thus, the idealized features of the model are not problematic, as *they* do not need to be isomorphic to the real world system. The fit between model and real system is preserved by those aspects that *are* isomorphic. The problem with this move is that very often it is precisely those idealized features that are supposed to represent aspects of the real-world system (ibid). This is especially apparent in models from economics and the social sciences that explain human behavior through the analysis of idealized agents. Neither the agents nor their behavior are isomorphic with respect to real humans, yet it is precisely that behavior that should be isomorphic if the model is to fit the phenomenon.



An alternative account of model-world relations (still within the semantic camp) can be found in Giere's account of scientific explanation (Giere 1988). On Giere's account, scientific theories are simply families of models, while models are idealized structures that stand between scientific statements and real systems. Giere thinks that theoretical statements or equations do not describe real systems directly, but define models. The models, in turn, represent real world phenomena and stand in a relation of *similarity* to these phenomena. If a model is sufficiently similar to the real world phenomenon, then it is a good representation of that phenomenon. For example, a pendulum is a real world system which moves in a particularly interesting way (it oscillates). The theoretical model of a simple harmonic oscillator is similar to this system, because it is abstract and idealized (it does not take into account friction, for instance). Nonetheless, it is a good representation because it captures the oscillatory motion, and we can use the model to indirectly study real world systems such as pendulum clocks.

I agree with proponents of the semantic tradition that the practice of modeling is a very important aspect of science. I also agree that it is useful to develop a precise account of model-world relations. However, I think that these accounts of model-world relations are not entirely satisfactory. More specifically, the criticisms leveled against isomorphism and partial isomorphism views are correct and I do not think they can be overcome. Also, while Giere's view of similarity is promising, he does not have a developed account of how similarity ought to be measured. Without a detailed metric, model-world relations will remain rather obscure. After all, everything can be similar to everything else; we need a way of measuring *relevant* points of similarity between models and real systems.

There is one view of model-world relations which does get more precise (Weisberg, 2013). Drawing on Tversky (1970) Weisberg states that a model is similar to a real-world system when it shares certain important features of the system and does not lack many important features contained in the system (ibid). Weisberg thinks of systems and models as sets of features  $\Delta$ , which can be divided into attributes and mechanisms. The former are the properties and patterns of a system and the latter are the underlying mechanisms that generate those properties. With this distinction in place, we can identify the common attributes and the common mechanisms between the model and the real system. We can also identify the differences (in terms of attributes and mechanisms) between the model and the system. Thus a model is more similar to its system when it shares many attributes and mechanisms and less similar when it contains additional or different attributes and mechanisms. Moreover, the various attributes and mechanisms of

systems and models will differ in terms of importance. For example, a real pendulum has a certain color, which is not present in the model or the model description. However, this is not as important as the fact that the model does not take into account friction. Weisberg thinks that we can capture this by differentially weighting subsets of features.

The details of how the similarity account works are not relevant for my dissertation, as my main objective is to provide the final piece in the puzzle of model-world relations. Weisberg does give a full account of how similarity can be assessed, along with a non-subjective method for weighting the various features of models and systems.

There is, however, a residual issue in these accounts of model-world relations. In the previous section I suggested that models do not represent entire natural systems but target systems. Yet in order to have a full account of model-world relations it is not sufficient to describe the relationship between models and targets, but also to characterize the relationship between targets and the world. While it may seem that target systems are conceptually simple notions, there is no real consensus on the nature of targets. I think that target systems should be understood as parts of the world, yet they could also be understood as fictional entities or models, or even all three. The issue is that each of these views of target systems has different implications for the characterization of model-world relations, and the value of target system analysis.<sup>2</sup> This dissertation is an argument for the conception of target systems as real parts of the world, a view which is coherent and which allows target system analysis to reach its full potential, by helping to solve genuine conceptual and methodological issues in scientific practice.

In chapter 2, I present my view of how target systems are specified, through a process of *partitioning* and *abstraction*. Put simply, partitioning is the division of the domain of study into parts. In other words, when partitioning the domain of study, a scientist is deciding what sorts of things count as objects or units. Abstraction is the process by which the scientist decides which aspects of the partitioned domain of study are relevant for studying the phenomenon of interest and hence, which parts of the entire system will be used in the model. Neither partitioning nor abstraction change the ontological status of the natural system, hence target systems are real systems in the world, albeit not entire natural systems.

In chapter 3, I look more closely at the process of partitioning. On my view, there are many potential ways that a natural system can be partitioned. For example, an ecologist can partition a natural

---

<sup>2</sup> By target system analysis I mean the precise specification of the target's partition and abstractions.

system into individual organisms, yet an evolutionary biologist can partition the same system into smaller units, alleles. Both partitions are possible in principle and selecting a partition depends on the type of phenomenon being investigated. I will also address a worry which is often associated with pluralistic views of partitioning, namely that if there are no rules that govern partitioning, then the resulting partitions cannot be real. I will adopt a version of *promiscuous realism*, which has hitherto been used to argue that there are countless ways of taxonomizing the world into kinds (Dupré 1993). On this view, even though many taxonomies of the natural world are possible, this does not mean that taxonomies are artificial. In addition, there are principled ways of deciding which of the many possible taxonomies should be used. This view has been influential in philosophy of biology as it provides one way of solving the debate concerning the nature of biological species. I will also show that a pragmatic conception of natural kinds is compatible with my view of partitioning.

I then turn to the notion of abstraction (chapter 4). There are several accounts of abstraction in philosophy of science. The received view, which I call the *material view*, is that an object is abstract when it is not concrete (Cartwright, 1994). If this were the case, then my argument that target systems are real parts of the world would break down. If target systems are aspects of the world they must be concrete, yet abstraction results in objects that are not so. Nonetheless, there is another notion of abstraction which I call the *omission view*, where abstraction simply means the omission of irrelevant parts and properties of a system (Jones, 2005). As it stands, this notion suffers from the criticism that abstraction can be understood as a type of idealization. I offer a refined account of the omission view of abstraction which overcomes the criticisms leveled against it. On my version of the omission view, abstract systems are merely incomplete and can therefore also be concrete. Target systems are an example of abstract systems that are concrete and not idealized.

There are a number of reasons which explain why target systems and their relation to natural phenomena have been overlooked. Philosophers and scientists alike seem to think that identifying which aspects of the world are actually represented by the model are not that important. In some cases, this is because scientists think of this as preliminary work which does not get published. The model itself, along with its explanations and predictions, are what are considered important. In other cases, scientists simply do not need to worry about which aspects of the system they are studying are relevant for their model, because they are determined by pre-existing conventions in a discipline. In the philosophical literature, the

overlooking of target-world relations usually occurs because the process of identifying what is relevant for the model and what is not, is subsumed under the notion of idealization, hence it is not thought to merit any special treatment.

In philosophical terms, the main reason why target systems are important is because they are a fundamental part of modeling, hence understanding their relationship to natural systems provides a complete picture of model-world relations. In addition, establishing the place for target systems in modeling provides us with an easy way to distinguish between the world and the model. This division is founded on the distinction between abstraction and idealization, which already has support in the philosophical literature (Cartwright 1994, Jones 2005). On my view, idealization (and any problems pertaining to idealization) should be analyzed at the level of the model, whereas issues concerning partitioning and abstraction should be examined at the level of the target.

As targets are not idealized models but real parts of the world, the methods we use for evaluating models, such as isomorphism and similarity, will not apply to targets. In chapter 5, I propose an alternative method of evaluation based on the notion of *aptness*. What we usually mean when we say that a subject or object is apt, is that it is suitable or appropriate given a particular purpose or set of circumstances. In the case of target systems, aptness is determined by the appropriate partitioning and level of abstraction. An appropriate partition is *useful* for the construction of a model, while an appropriate abstraction contains all and only the *relevant* factors which give rise to a phenomenon. In this chapter I explicate my notions of usefulness and relevance and develop a general method for determining whether or not a factor is relevant. To this end, I incorporate aspects of the *kairitic* account of causal explanation which provides a way of identifying the actual causes of an event from a wider web of potential factors (Strevens 2004).

In chapters 6 and 7, I will examine the use of target systems in scientific practice. I will consider the case of invasive species research, a relatively new field in biology characterized by a multitude of approaches. In chapter 6 I will give an overview of the history of invasion biology highlighting the fact that its diversity is problematic because of its failure to predict invasions. According to invasion biologists this problem can be solved by the creation of a more general unified framework for studying invasions (Gurevitch, Fox, Wardle, Inderjit, & Taub, 2011). However, I will show that a scientist constructing a unificatory framework in invasion biology faces a tradeoff between generality and predictability. A truly unified framework must incorporate a multitude of different causes of invasion, yet the causes of each

invasion are unique. I will show that it is possible to have a unified theory of invasion based on the *process* of invasion, yet this framework cannot be predictive as it does not focus on the *causes* of invasion. I will then present an alternative conceptual framework which is predictive and integrative but not unificatory as it can only achieve mid-level generalizations. In chapter 7 I will give a deeper analysis of target systems in invasion biology. First, I will examine the scientists' own questioning of the usefulness of the discipline's theoretical framework and show that it can be understood as a case of target system analysis. I will then show how this preliminary target system analysis can be extended to determine the extent to which integration is possible and useful, as well as the extent to which results can be generalized. The discussion in these chapters demonstrates the importance of target systems in scientific practice. In the first instance my accounts shows that target system analysis is already part of the theoretical aspect of scientific inquiry albeit in an informal and imprecise way. It then goes on to show that a more precise and philosophically grounded target analysis can help to determine which theoretical issues are solvable and can suggest ways in which they can be solved. It is then that target system realise their full potential as tools in the scientific toolkit.

The importance of target systems is not limited to cases of modeling in ecology. In the concluding chapter I outline some ways in which target system analysis can be extended to other disciplines and methods. For instance, it is quite common in philosophy of science to distinguish between modeling and experimentation as two radically different aspects of science. This distinction has sparked a debate concerning the relative virtues of the two. However, there has been a recent shift in focus, as some authors have started looking for commonalities between modeling and experimentation so that the two aspects of science can become more integrated . Paying attention to the specification of target systems can help with this integration as experiments also have targets. Just like a modeler, a successful experimenter must choose a particular framework in which to work and identify the factors relevant to the phenomenon being studied.

Target system analysis can be instrumental in the apt choice of frameworks.

## 2. What is a Target System?

### 1. The Missing Targets

Many phenomena in the natural world are complex, so, rather than studying them directly scientists construct models. These models are idealized representations of the systems in the world in which the phenomena take place (Giere, 2004), (McMullin, 1985),(Weisberg, 2007a; 2007b). They are studied with various goals in mind including generating predictions about how systems will change in particular circumstances and explaining why a particular set of circumstances came about in a system. In many of these cases the model is aimed at giving knowledge about a real-world system. Yet natural systems often cannot be represented by a model in their entirety, so scientists pick out some aspects of a system and represent only those in the model. The standard term used by philosophers of science to refer to what a model represents is ‘target system’ (Frigg, 2009; Giere, 2004; Godfrey-Smith, 2008; Knuuttila, 2005; Suarez, 2003; Weisberg, 2013).

There is a substantial literature in philosophy of science on ‘model—world relations’, which investigates the connection between models and natural systems. This literature focuses on the nature of scientific representation and on characterizing the relationship between models and the world, such as isomorphism (Van Fraassen, 1980), partial isomorphism (Costa & French, 2003), models as fictions (Frigg, 2009) and similarity (Giere, 1988; Weisberg, 2013). My purpose here is not to take a side in this debate but to address an issue which has been relatively neglected. Despite their differences, proponents of these views generally agree that models represent target systems, not natural systems in their entirety. Thus, accounts of model—world relations are actually accounts of model—target relations. However, there are few extended discussions of the notion of a target. In other words, there is a gap in the characterization of model—world relations, namely the relation between the natural system and the target system. The principal aim of this chapter is to fill this gap by explaining the nature of target systems and show how they are specified.

## 2. What is a Target System?

Philosophers of science agree that a target system is what a model represents, yet this simple definition does not give us much information about the nature of targets. Many authors who make explicit mention of target systems do not provide specifics about their nature. For example, Frigg (2010) states that a target system is a “particular part or aspect of the world that we are interested in” (2010, 252), while Giere states that “Scientists use models to represent aspects of the world for various purposes” (Giere 2004, 747).

Authors who give more specific definitions of target systems, usually characterize them as having a disjunctive set of essential properties. Suarez (2003), for example, distinguishes between a *source*, or vehicle, of a representation and its *target*, which is the object of the representation. He elaborates on the nature of sources and targets stating that both “may be concrete physical objects, systems, models, diagrams, images or equations” (Suarez 2003, 226). Another example comes from Weisberg (2013) who makes a distinction between target-directed modeling, generalized modeling, hypothetical modeling, and targetless modeling. In the case of target-directed modeling a target is a “single real system” (2013, 91) which is an abstraction over a phenomenon in the world. In the case of generalized modeling, the target is an abstract generalization over many phenomena. In hypothetical modeling, the target is an imaginary or hypothetical system while in targetless modeling, the target is completely absent.

This implies that targets are understood as a disparate group of entities, as they can be concrete or immaterial, objects or systems, aspects of the world or models, images or equations. Unfortunately, this pluralistic conception of the nature of targets obscures their importance in scientific practice and detracts from descriptions of model-world relations. If one of the aims of explaining model-world relations is to understand how and what models tell us about phenomena in the world, then we must explain not only how models represent targets, but how targets relate to the world. Yet the relation between targets and the world cannot easily be characterized if targets are sometimes understood as models and other times as parts of the world. The relationship between a model and the world is very different to the relationship between a part of a system in the world and the entire system.

In what follows, I will argue that target systems are *parts of the world*. In other words, target systems are real parts of natural systems which have the same ontological status as the natural systems. Moreover, even though there are various different types of models, such as mathematical, concrete and

computational models (Weisberg 2013), whenever a model represents a real-world phenomenon it only represents one kind of target: a real part of the world.

My account of target systems begins by considering how target systems are specified. I will break down target specification into four conceptually distinct stages: (a) identifying the *phenomenon of interest*, (b) locating the *domain of study*, (c) *partitioning* the domain and (d) *abstracting* to reveal the relevant parts and properties of the system. I should note that the four elements of target system specification do not always occur in this order, nor is target specification a linear process. Each of these elements can be (and often is) extensively revised as new information comes to light.

The first element of the specification process is the identification of the *phenomenon of interest*. Phenomena are identified within the framework of disciplines or sub-disciplines. For example, population growth, competition, predation and invasion are phenomena studied in the discipline of (population) ecology. In evolutionary biology the phenomenon of interest could be the mutation rate in a population or the frequency of a particular allele in a population. Other disciplines will have different phenomena of interest.

Phenomena in the natural world usually occur in specific spatiotemporal locations. Identifying the location in which the phenomenon of interest occurs is also part of target system specification. This I call the location of the *domain of study*. The dimensions and character of a domain depend on the particular discipline of the scientist. In the case of ecology, the domain of the real-world phenomenon is often geographically determined and can be an ecosystem, an island, an area defined by a particular microclimate and so on. Other disciplines have domains of study that look very different. In climate science the domain could be the Pacific Ocean, a particular country or even the entire planet, whereas in anthropology domains are usually particular groups of people.

Once the phenomenon and the domain are identified, a scientist partitions the domain. Put simply, *partitioning* is the division of the domain of study into parts.<sup>1</sup> That is, the scientist is deciding what sorts of things count as objects or units for the purpose at hand. Partitioning also involves identifying properties that correspond to the units. Properties will vary depending on the sort of thing that each unit is, but in most

---


<sup>1</sup> I will be using the terms 'parts' and 'properties' throughout this chapter. By parts, I mean identifiable units of natural systems. I use this term because I want to emphasize that these units have the same ontological status as the larger system which they find themselves in. The natural system is the whole and the units are some parts of it. The parts have not been altered in any way. I use the term 'properties' simply to refer to features of these units. A detailed account of the nature of properties is beyond the scope of this dissertation. Moreover, the nature of this analysis is such that it can remain silent on the fundamental nature of properties. All that is needed is a way to refer to features that are predicated or instantiated, universally or not, by the units of the analysis.



cases each unit will have a very long list of properties. The type of partition is determined by the question a scientist is investigating and the disciplinary framework in which the study is taking place. For example, in population genetics scientists model the change in frequency of alleles in a population. Hence alleles are the main units of study in population genetics. Similarly, in organic chemistry the main units of analysis are organic molecules, in physics they may be particles, in psychology they are individuals, in anthropology they are groups of individuals and so on.

Partitions can also cut across levels of specificity. For example, a plant ecologist studying competition might group an individual plant and the mycorrhizal fungi in its roots as one unit, and the nitrogen which the fungi help the plant to absorb as another unit. Figure 2.1 shows some ways in which a natural system can be partitioned. Partitions can vary in terms of fineness of grain, but they can also be partitioned at more than one grain by having partitions of different sizes, or even partitions that are hierarchical or nested.

**Figure 2.1. Partitions**



My notion of partitioning is similar to the formalized notion of partitions in set theory. A partition of set  $S$  is a collection of nonempty, mutually disjoint subsets of  $S$  (Lucas, 1990). The union of all subsets of  $S$  is the set  $S$ . Subsets of a set are also sometimes called parts (ibid). For example, the set  $\{ 1, 2, 3 \}$  has five possible partitions:  $\{ \{1\}, \{2\}, \{3\} \}$ ,  $\{ \{1, 2\}, \{3\} \}$ ,  $\{ \{1, 3\}, \{2\} \}$ ,  $\{ \{1\}, \{2, 3\} \}$  and  $\{ \{1, 2, 3\} \}$ . The power set of a set is defined as the set of all its subsets and is denoted  $P(S)$ . In the example,  $P(S) = \{ \{ \}, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\} \}$ . Partitions are generated from combining non-overlapping subsets given by the power set.

Similarly, in the case of target systems, every domain can be partitioned in a number of different ways given by each domain's power set. Partitioning amounts to choosing one of the available partitions

from this power set. Also, while a partition does not have to have parts which are all the same size, all the parts of the set are present when the set is partitioned. For example, a domain {Vancouver Island} might be partitioned into a subset of {marmots}, {grass}, {air}, {water}, {nutrients}, {other animals}, {other plants} and {everything else}. These subsets are all different types of things which have different sizes yet all the parts of the domain are present somewhere in this partition. We could also partition the same domain into {marmots}, {carbon molecules} and {everything else}, even though some of the carbon molecules are located within the marmots. This is not a problem, as long as we keep in mind that in this partition, the subset of marmots does not include their carbon molecules.

The next element in target system specification is to eliminate the subsets which are not relevant for our purposes. I call this a process of *abstraction*.<sup>2</sup> Abstraction has come to mean a number of distinct things in philosophy of science, including concept formation (Rosen, 2012), the de-concretization of a system (Cartwright, 1994), the simplification of a system (Frigg & Hartmann, 2009; Weisberg, 2007a) and the omission of parts from a system (Jones, 2005). Although the first three notions of abstraction are commonly used, I will restrict the notion of abstraction to the omission of irrelevant parts from a partitioned system.<sup>3</sup>

In the example of the Vancouver Island Marmots, the ecologist would adopt a partition such as the one mentioned above (marmots, grass, air, water, nutrients, other animals, other plants, everything else). She would omit the subsets {everything else}, {air} and {water} as irrelevant. The remaining subsets {marmots}, {grass} and {nutrients} are the target system.

An important feature of my view of abstraction is that the omission of parts does not distort the system. The omitted parts are not physically removed from the target, nor is the target itself altered in any way. For example, the scientist does not kill off all other organisms in the system, nor does she eliminate the air or water from the system. Instead, she simply focuses on some parts and properties that are relevant for studying the phenomenon of interest (Figure 2.2, grey squares) and ignores the rest (Figure 2.2, black


---

<sup>2</sup> Weisberg (2013) also argues that most target systems are generated by abstraction, yet his notion of abstraction is more general than the one presented in this dissertation.

<sup>3</sup> A full defense of this notion of abstraction is beyond the scope of this chapter but appears in chapter 4. Nonetheless, it should be noted that omission forms the core of most notions of abstraction. For Cartwright, who follows the Aristotelian view of abstraction (Frigg & Hartmann 2009), de-concretization and generalization rely on the process of omission. Similarly, the Lockean notion of abstraction as concept formation relies on the identification of common properties of a group of particulars and the *omission* of the properties that are not shared.

background). In other words, abstraction does not change the ontological status of any elements in the set, nor does it change the relations between any of the elements.

**Figure 2.2 Abstractions**

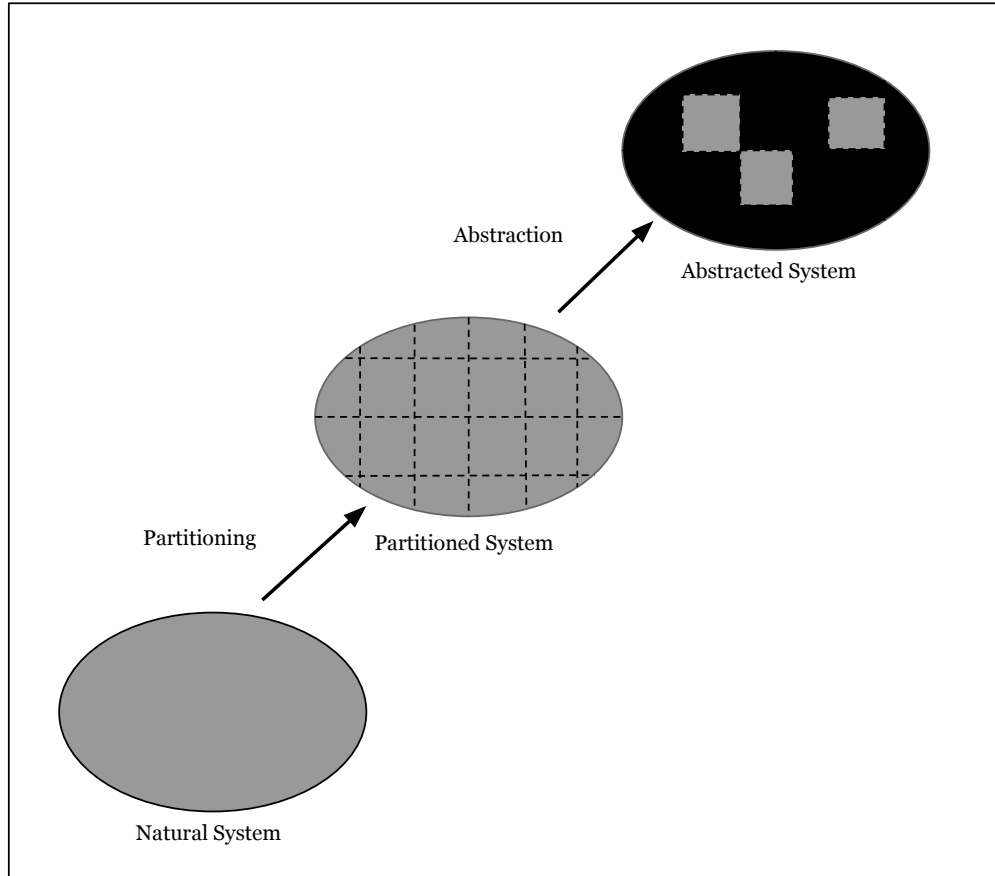


Just like partitions, abstractions can differ. Omitting many features results in more abstract target systems while omitting fewer features results in less abstract systems (see figure 2.3). However, abstracting indiscriminately can have unanticipated consequences. If we abstract too much we are in danger of missing important causal factors which can lead to a misinterpretation or distortion of the system we are studying. If, on the other hand, we do not omit enough parts or properties from the system, we can end up with a target system that is too complex to experiment on or model. The key lies in being able to identify the relevant aspects of the system for the particular model and omitting the rest.

To recap, target system specification (Figure 2.3 ) involves the location of a natural system situated in the world and its partition depending on the phenomenon we are interested in. We abstract the partitioned system focusing on some parts of the system and ignoring others. The partitioned, abstracted system in the world is the target system. In the next section I will present the advantages of my view of target systems. I will then consider two alternative views of targets and show that they do not share the

advantages of my view. Finally I will present some residual criticisms of my view and show how these criticisms can be overcome.

**Figure 2.3 Target System specification**



#### 4. Questions of Ontology

With my analysis of target system specification in place, it is now time to turn to the question of the nature of target systems. I will provide further arguments that target systems are parts of the world and stand in a part-whole relationship to natural systems. Therefore, target systems are concrete and real parts of the world.

The first advantage of my view is that it is the most straightforward way to understand target systems. On my view, target system specification does not distort the system in the world nor does it

involve any representation of the system. Partitioning picks out a subset from a list of *already available* subsets, while abstraction involves focusing on some aspects of the subset and not others. The parts themselves are neither changed nor removed from the real-world system, hence the target system is a real part of a real system.

The second advantage of my view is that the process of specifying target systems does not generate a new type of entity. As targets are parts of the world, they are the same kind of thing as the world itself. This means that target systems are not ‘intermediaries’ between models and the world, so we do not have to give a *metaphysical* account of how the target relates to the world. In other words, my view of targets does not complicate the picture of model world relations by adding a new type of entity. The only different ‘kind’ of thing with respect to the world is the model.

I will illustrate these points with the help of an analogy. Imagine a student studying for an exam using notecards. On each notecard she will write down important information in her own words, which she will then memorize. The notecards are the analogue of a model, as they are the student’s representation of the information in the textbook, while the relevant information in the textbook is the equivalent of the target system. However, in order to generate the notecards she must first determine which parts of the textbook contain the information relevant for the exam. The textbook is analogous to the domain of study, while the exam is the analogue of the phenomenon of interest. The student ‘partitions’ the domain into units, which are sentences. Her task is to determine which sentences in the textbook are important enough to represent on her notecards. In order to do this, she reads through the textbook and ‘omits’ all the sentences that are not relevant to her purposes. She is left with the analogue of the target system, that is the set of sentences which are relevant for her exam.

No part of this process involves changing the fundamental nature of the textbook. Partitioning the textbook into sentences does not change the textbook itself, nor does it alter the information within the textbook. It simply produces the power set of the sentences which could be relevant for the exam. The omission of irrelevant sentences reveals one subset of the power set, which *is* relevant for the exam. The act of omitting some parts of the text does not change the text. The student does not cut out the omitted sentences, nor does she change the order of the remaining sentences. She simply identifies which subset is the relevant set for her exam. Moreover, the student has not yet represented the textbook in any way. The sentences are still part of the textbook and hence part of the world. The representation occurs when the

student writes out her notecards, as she represents the relevant information from the textbook *in her own words*.

The third advantage of my view of target systems, is that it takes us a step closer to a full account of model-world relations. On my view, models are idealized representations of targets and targets are parts of the world. By analyzing a natural system, a target and its model, we can determine what the model tells us about a particular phenomenon in the world. As stated above, philosophers of science started using the notion of a target system in the context of model-world relations to denote the system that a model represented (Frigg, 2009; Giere, 2004; Godfrey-Smith, 2008; Knuuttila, 2005; Suarez, 2003; Weisberg, 2013). Most of the literature on model-world relations has focused on the relationship between a model and its target, in terms of isomorphism, partial isomorphism, similarity and so on. Even if my account is not fully accepted, I hope to show that a comprehensive analysis of target systems is important, especially given their place in model-world relations.

In order to provide some context for my view and its advantages, I will present two alternative views of target systems where targets are not real parts of the world. I should note that these views are not explicitly stated in the literature *as* views of target systems. Still, they are viable options of how target systems could be understood if they were not parts of the world.

#### 4.1 Targets as Models

The first possibility is if the target system is thought of as a type of model, which is itself represented by the mathematical model. An example of a system which is represented by a model but which is itself a representation of a system in the world, is Suppes's notion of *models of data* (Suppes, 1969). A model of data is an intermediate step between the *model of the experiment* and the *experimental design*. The experiment yields a set of raw data which the scientist can observe. She then proceeds to make sense of this data by eliminating any errors (outliers) and presents it in a comprehensible way, for example by fitting a curve to the set of data points (Frigg & Hartmann, 2009). As models of data involve distorting the raw data by presenting it in a 'neat' way, they constitute idealized models.

Models of data are supposed to circumvent a criticism of the isomorphism view of model-world relations. The criticism is that real world phenomena do not have the kind of structures which can be isomorphic to mathematical structures in a model (Knuuttila 2005). For example, the logistic model, a

typical model in ecology (which I mentioned in chapter 1), is used to measure how the growth rate of a population  $N$  is limited by the density of the population itself (see equation 1.1 p. xx). The important component for this discussion is  $r$ , the intrinsic rate of increase. This is the growth rate of the population independent of resource limitations, competition, or predation. In other words, it is the average number of offspring that an individual has, when the population is at low density. The question is how can  $r$  be isomorphic to a structure in the real world system?

The short answer is that it cannot, because there is nothing in a real-world system which is structurally similar to  $r$ . A growth rate is not a property of an actual population, but a statistical variable derived from data collected from parts of the system in the world. As stated previously, organisms are born, have offspring and die. Ecologists collect data on the number of births and deaths of the population within a specified timeframe. Then they construct a life table, a tabular summary of the birth rates, fecundity and death rates of a population, divided into age groups (Ricklefs & Miller 2000). The values from the life table are then used to calculate the intrinsic growth rate of the population. However, the actual birth and death *events* of the population do not have the same structure (or even a similar structure) to the mathematical notion of a growth rate.

The model of the data helps to circumvent the criticism because it is an ‘empirical model’ which has the necessary structures and can be isomorphic to the mathematical model. However, as it stands, the isomorphism account does not give a full picture of model-world relations. It may be true that the model of the data is somewhat ‘closer’ to the real-world phenomenon, yet it is still an idealized and manipulated model.

Suppes’s account includes two additional steps which are meant to show how the model of data relates to the phenomenon in the world. Working downwards from the model of data, we arrive at the stage that deals with the problems associated with ‘experimental design’. For example, it involves a formalized way of randomly assigning subjects into groups. This is the preliminary step for the experiment which yields the raw data. The issues that are ironed out at this level are usually determined by the lowest level, that is the determination of ‘*ceteris paribus*’ conditions, such as “control of loud noises, bad odors, wrong times of day or season”(Suppes, 1969).

In fact, I think that these other steps are similar to my account of target system specification. For instance, in the determination of *ceteris paribus* conditions scientists are deciding which of the many factors

which give rise to a particular phenomenon are actually relevant for the study. This looks very similar to my notion of abstraction, as both processes are aimed at identifying the relevant causes of a phenomenon which should be included in the study. In addition, partitioning is already implicit in the determination of the *ceteris paribus* conditions and the experimental setup, even though Suppes does not identify it as such, as the domain must be partitioned in units such as individuals or groups, whose properties will then be tested in the experiment. Thus, my view of target system specification is compatible with Suppes's account of model hierarchy, as long as the lower aspects of the hierarchy are seen as the analogues of target systems instead of the models of data.

Of course, it is still possible to designate the model of data as the target system. However, this seems rather arbitrary. Suppes's account identifies a hierarchy of three kinds of models: linear models, mathematical models and models of data. Given that target systems are defined as what models represent, we would be equally justified in calling the mathematical model a target system (as the linear model represents the mathematical model). This is problematic because it is at odds with what most philosophers of science have in mind when they are thinking of target systems.

An additional issue with this option is that it involves many intermediaries of different kinds. Models of data are empirical models, hence a different type of model to mathematical models. If target systems were to be thought of as models of data we would have three distinct entities, real world systems, target systems and mathematical models. In contrast, on my view, target systems are just parts of the world, hence there are only two entities involved in scientific modeling, models and the world. My view of target systems provides a much simpler and straightforward way to categorize the entities involved in scientific modeling. It also provides a coherent way to distinguish between target systems and models, which can be applied across different studies or disciplines, as target systems are partitioned and abstracted yet models are idealized.

## 4.2 Targets as Fictions

A second way in which targets could be ontologically different to natural systems is if they are thought of as fictions or imaginary systems that are described by idealized models. Scientists sometimes construct models of systems that do not exist. In fact, they sometimes construct models of systems that



could not possibly exist (Weisberg 2013).<sup>4</sup> For example, the exponential growth model (equation 3.1) represents how a population would grow if it were not checked by density effects. There is no population on Earth which is not subject to density effects, as there is no environment which can support an exponentially growing population indefinitely. The main difference between the exponential growth model and the logistic growth model is that the latter takes resource scarcity into account and therefore can be said to represent actual populations in actual environments. Consequently, its target systems are aspects of the world. Then, the question to ask is: what are the targets of the exponential growth model?

One option is to say that the target systems of the exponential growth model are hypothetical systems of populations in environments with unlimited resources, whose populations are infinitely large. Of course, it is impossible to have an environment without scarcity or an infinite population in the actual world. This means that these hypothetical targets cannot, strictly speaking, be thought of as aspects of the world. If this is the case, then we need to have an account for how the imaginary system relates to the actual system. In more general terms, models are supposed to tell us about real world phenomena, but if they represent imaginary systems, then there is an extra step to account for, namely how the imaginary systems relate to real-world phenomena.

### The Exponential Growth Equation

$$\frac{dN}{dt} = rN \quad 1.4$$

One way to account for this is through the notion *idealization* (McMullin, 1985; Weisberg, 2007a).

The exponential growth model is idealized because it represents an environment without resource

---

<sup>4</sup> Another group of models which can be claimed to have hypothetical targets are *generalized models* (Weisberg 2013). These are hypothetical targets of generalized phenomena. For example, a general model of sexual reproduction “isn’t supposed to be about kangaroo sex or fungi sex, but about sex itself” (Weisberg 2013 116). According to Weisberg, the target of a general model of sex will also be general, without the particulars associated with any particular population. However, just because the model is general, does not mean that it must apply to a single generalized target. In fact, a model can be considered general when it applies to many targets, hence the model of sex is about kangaroo sex *and* fungi sex *and* tasmanian devil sex etc. In fact, Weisberg agrees that so called generalized targets are actually constructed out of actual targets in the world, hence his view is compatible with my own. While I would resist referring to the generalized system as a ‘target system’, the important point for my argument is that ultimately these general models represent real-world target systems.

limitations and populations as infinite. One could argue, therefore, that the model contains the same idealizations as the target it represents, and it is the *target* which is idealized with respect to the system in the world. While this is a possible solution, it makes the target system redundant. In the standard account of modeling practice, idealization occurs in the construction of the model. If the target system is also idealized, then it will be identical to the model. This would add an extra step to the process of modeling without conferring any benefits. We would still have an incomplete account of model-world relations because we would not know how the idealized model and target related to the non-idealized phenomenon in the real world.

On the other hand, if we restrict idealization to the model, and maintain that the target system is part of the world, the inclusion of the target adds something important. Specifying a target system allows us to give a full account of model-world relations, given that we now have an account of how the target relates to the world. On this view, the targets of the exponential growth model are very similar to the targets of the logistic growth model. For example, if the exponential model were to be applied to the population of marmots outlined in the introduction, the target would be comprised of the marmots and some of their properties. The difference between exponential growth target and the logistic growth target of the marmot population, is that the latter includes the maximum number of marmots supported by the ecosystem. Hence, both models are idealized, yet both targets are merely abstract when compared to the entire ecosystem. Unlike the view of targets as imaginary systems, my view of targets provides a clear distinction between model and target, but also reflects the importance of target systems in scientific practice.

### 4.3 Targets are Real

Even though the alternative views of target systems create problems for the viability and usefulness of target systems, I anticipate some resistance to the claim that target systems are real. There are three criticisms which I will address in this section. The first is the view that we should adopt a pluralistic notion of targets, the second is the view that nature dictates how a system can be partitioned and the third is the view that abstraction necessarily involves distortion.

One possible criticism of my view of target systems is that it is too strict. While someone might agree that *some* targets are parts of the world, they might also want to allow for the existence of other kinds of targets. I have mentioned examples of generalized models, models of hypothetical phenomena, even

models which do not represent anything in the world. Why should we not refer to whatever it is that a model represents as a target system? The simple answer is that it is simpler, more consistent and coherent to use the notion of a target system in reference to a single kind of entity. By restricting the notion of target systems to the partitioned abstracted systems in the world, we have a clean distinction between models and the parts of the world that they represent. The main aim of this chapter is to argue that when models represent phenomena *in the world*, those target systems are real. Moreover, as I have shown, in many of the other cases, such as highly idealized general models, closer examination reveals that the models do represent real-world target systems. Having said that, however, my view is compatible with a more pluralistic conception of targets. I am not, in this paper, referring to cases where scientists intentionally construct a model of a *truly* nonexistent system, though my view does not preclude these special cases.<sup>5</sup>

The other criticisms are stronger, as they deny the possibility that *any* target systems are real parts of the world. The second criticism stems from the view that scientists must either carve nature ‘at its joints’, or admit that the partition is not real. This is a problem for my view because I stipulate that each real-world phenomenon can be partitioned in a number of different ways, while no particular partition is in principle privileged or better. Yet if this criticism is right, only some partitions and by extension, only some target systems, can be thought of as parts of real-world phenomena. Alternative partitions might be possible, yet they are imposed on the phenomenon by the scientist. If this is the case, then only a small minority of target systems can truly be thought of as parts of real world systems.

I have two responses to this criticism. The first is based on the view that scientists do not carve nature ‘at its joints’, but that there are countless ways of taxonomizing the world into kinds, all of which are equally legitimate. Following Dupré, one can argue that the structure of the world is very complex and can be categorized in many cross-cutting ways (Dupré, 1995). What determines a classification is not the world itself, but the interest a scientist is pursuing (Bird & Tobin, 2008). Dupré argues that within biological taxonomy, there are many ways of classifying species, depending on the model of biological systematics which a scientist uses, and that every one of these classifications is objective and real.

---

<sup>5</sup> I am referring to these cases as special, because I think that most scientists intend that most of their models tell us about systems in the world. Even cases which might seem too idealized to apply to real world-systems are, in fact, intended to elucidate real world systems. For example, Weisberg (2013) uses the example of 3-sex biology to show that some models have no real-world targets. However, as the scientists constructed the model in order to learn about 2-sex biology (ibid), the model can be understood as a *very* idealized model of real 2-sex populations.

The same can be said of target systems. Scientists from different disciplines carve up the same domain in a multitude of different ways, depending on the particular phenomenon they are studying. For example, a population ecologist carves up an ecosystem into organism-sized units, but an ecosystem ecologist might think of the entire ecosystem as a unit (Odenbaugh, 2007). A geneticist would carve up the same domain into smaller, allele-sized units, while a climate scientist examining the microclimate of the domain would partition the domain in terms of factors that affected precipitation, wind etc. Are any of these partitions illegitimate? Is there a partition which should take priority over the others? More importantly, is any of these partitions more real than the others?

Dupré's pluralism stems, in part, from his anti-reductionism (Dupré, 1995). Reductionistic approaches to science do prioritize some partitions over others, as they favor fine-grained partitions over coarser grained partitions (Rosenberg, 2008; Sterelny & Griffiths, 2012). Dupré, along with many others, has criticized reductionistic accounts for ignoring the kinds which occur at higher levels of organization, due to emergent properties and the complexity of higher level systems (Brigandt, 2013a; Dupré, 2009; Kitcher, 1984; Mitchell, 2003). Different phenomena occur at different levels of organization, hence scientists must partition their domains in different ways, in order to identify and study these phenomena. Moreover, for Dupré, the issue is not merely epistemic but also ontological. It is not simply that scientists find it easier to identify phenomena at different levels of organization, but that these phenomena only exist in some levels of organization and not others. In other words, there are many different natural kinds, which occur at various levels of organization (Dupré, 1995; Kitcher, 1984), and they can only be studied if scientists partition their domains accordingly.

Although I accept Dupré's radical pluralism, not everyone does. Yet there is another, less radical response to the criticism. It also starts from the observation that the world is very complex and does not give us sufficient reason to privilege a particular type of classification over another. The difference is that on this view, does have joints, yet we are not in a position to discover them. Our epistemic limitations as human scientists force us to take a pragmatic attitude towards classifications and partitions (Wimsatt, 1974). This means that we have no way of actually distinguishing between partitions, so, as long as they are useful they are legitimate.

Does the second response concede the main point of the criticism? If we have no reasons other than usefulness to distinguish between partitions, then they should all be thought of as equally *unreal*. If it

is the case that our basic perception of the world does not track the actual structure of the world, then all target systems are distorted representations of the world. This may be a legitimate criticism, yet it is not a criticism of my account of target systems. If this criticism holds, then it also applies to pre-partitioned systems. If our basic perception distorts a system, by definition, then our perceptions of target systems will also be distorted, yet *no more* than the pre-partitioned systems. Partitioned target systems are as real as pre-partitioned natural systems. More importantly, target systems are still less distorted than models, as they do not contain idealizations.<sup>6</sup>

It is also possible to criticize the view that target systems are real, by focusing on the notion of abstraction. There are two ways in which my notion of abstraction can be criticized. The first is that abstract objects are not concrete, hence target systems cannot be real. This criticism comes from a common notion of abstraction in philosophy of science which is that an object is abstract when it is not concrete e.g. (Cartwright, 1994), the implication being that abstract objects are not physically instantiated. However, this is not the only legitimate notion of abstraction. Earlier in this section and in the introduction, I argued that at least for the purposes of target system specification, abstraction should be understood as omission of irrelevant parts and properties. Omission does not imply or necessitate de-concretization. Therefore, if the criticism is that my notion of abstraction (omission) departs from a commonly-held notion of abstraction (de-concretization), then the issue is simply semantic. Those who do not wish to view this process as abstraction can simply refer to it as *omission*.<sup>7</sup> The point, for my purposes, is that omitting irrelevant parts and properties from a system does not alter its ontological status.

Nonetheless, there is another criticism which gets to the heart of the matter. This criticism stems from an understanding of abstraction as a kind of idealization, and states that there is no such thing as an omission which does not distort the system (Frigg & Hartmann, 2009; Humphreys, 1995; Weisberg, 2007a). The first version of the criticism was articulated by Humphreys (1995) in direct response to Cartwright (1994). According to Cartwright, when we abstract, “we strip away, -in our imagination- all that is irrelevant to the concerns of the moment to focus on some single property or set of properties ‘as if they were separate’” (ibid 197). An example of a concrete object is a triangle drawn on a blackboard. We can abstract a number of properties from the triangle, such as its color and what it is made of, until we get a

---

<sup>6</sup> I present a full account of the metaphysics of partitioning and responses to these criticisms in chapter 3.

<sup>7</sup> I should note once again that there is significant continuity between my view of abstraction and other views, as most views of abstraction incorporate and rely on omission. For a full argument of my view, and its relation to other views see chapter 4.

triangular shape. This, according to Cartwright, gives an abstract (de-concretized) idea of a triangle. According to Cartwright, “even when the chalk and the colour and all the other accidental features are subtracted, the shape that is left is never a real triangle. But let us pretend that it is ...” (Cartwright in Humphreys 158). Humphreys points out, however, that as soon as we start pretending that an abstraction is different than it seems, we have blurred the line between abstraction and idealization.

This is a strong criticism of Cartwright’s account, yet the problem centers on the extra step of ‘pretending’. It is true that if we pretend that the triangular shape is a true triangle and not merely triangular, then we are distorting the triangle. In fact, this kind of distortion is a true case of idealization, much like the idealization which occurs when an biologist ‘pretends’ that a population is infinitely large in her model. However, this is not a problem for my view of abstraction as it does not involve these types of distortions. On my view, it is *never* possible to get a true triangle from omitting its properties, only an imperfect triangular shape. Similarly, it is not possible simply by omitting properties from a real population to get a population which is infinitely large.

The second version of the criticism starts from the idea that systems in the world are dynamic, so each factor can be inextricably linked to other factors in the system with intricate feedback loops. If this is the case, then omitting a factor will result in a series of changes which count as distortions of the system. Hence abstraction must be a type of idealization (Frigg & Hartmann, 2009; Weisberg, 2007a). For example, when plants accumulate microbes such as mycorrhizal fungi and nitrogen fixers near their roots, this results in a beneficial effect on the plant which is called positive feedback (Callaway, Thelen, Rodriguez, & Holben, 2004). However, when they accumulate pathogenic microbes, they create increasingly hostile conditions for themselves in a process called negative feedback.

It is thought that while native species accumulate pathogens which cause negative feedback reactions in the soil, species that are introduced into a new habitat are free from their native underground enemies and exhibit neutral or positive feedback with the same microbial species (Klironomos, 2002). In some cases, an invader might have both positive and negative feedback going on at the same time, so that the effects of feedback are cancelled out. This means that even though these are important factors in plant competition (especially between native and invasive species) they are often overlooked. The omission of the feedback effects gives a distorted picture of the natural system, which can result in problematic explanations or predictions of competition between plant populations.

This is a legitimate concern, one which practicing scientists often have to grapple with. Yet it is an epistemic issue rather than an ontological issue. In ontological terms, target systems cannot be distortions, because they are real. In the example of plant-soil feedback, there are countless potential target systems which scientists might use, some of which include all feedback, some of which include only some kinds of feedback and some of which include no feedback. All of these target systems are simply parts of a system and are neither distorted nor unreal. Of course, if we were to *represent* a plant target system without feedback, we would be distorting the system. Yet this does not mean that the target itself will be distorted. It would still be a collection of real and concrete parts of the world.

It might seem that this response is skirting the issue. Surely not all target systems are ‘up for grabs’ in any scientific inquiry. I agree with this statement, because I think that what makes the choice of target system successful or not has nothing to do with ontology. Instead it depends on whether the specified target system is *apt* with respect to the particular study. In fact, part of the difficulty in finding the right target for a particular study lies in the fact that there are countless targets to choose from, all of which have the same ontological status.

To recap the argument so far, the two alternative views of target systems are problematic as they introduce unnecessary complexity to model-world relations and do not provide a clear way to distinguish targets from models. Instead, on my view, target systems are just parts of the world. The elements of target system specification, partitioning and abstraction, do not alter the ontological status of the system, hence target systems of natural phenomena are always real. In the next section I will sketch out the preliminaries of a theory for target system evaluation, based on epistemic issues. I will also give an explanation of cases where it seems that scientists have specified targets that are unrealistic or altogether unreal.

## 5. Problems of Specification

Throughout this chapter, I have only considered the relatively simple cases where the specification of an appropriate target system is successful. Yet part of the reason why formulating an explicit account of target systems is important, is that it can be very difficult to specify appropriate targets for a particular scientific purpose. For example, scientists sometimes erroneously classify relevant factors as irrelevant, or partition the system in a way which cannot be easily modeled. These target systems, while still real, are not successful given the scientists’ purposes. A full account of target system evaluation is

beyond the scope of this chapter, yet I will give the basic elements of such an account in the remainder of this section.

I begin my account of target system evaluation by considering three kinds of errors that theorists try to avoid. The first kind of error is when a target system is specified yet the abstractions and/or partitions are not optimal. The second is when a target system is specified, yet there is some mistake in the identification of its parts and/or properties. The third is when a theorist completely fails to specify a target system. I hope that the examination of these cases will provide additional support to my view that target systems are concrete real parts of the world, as it will show that many cases where targets seem to be immaterial are actually cases of mistakes in target system specification.

## 5.1 Inapt Target Systems

The way to think about the relative success of target systems is to determine whether or not they are apt. A subject or object is apt when it is suitable or appropriate *given a particular purpose or set of circumstances*. An apt description of a situation is suitable given the circumstances and gets the relevant information across to the audience. Whether a particular piece of information is relevant or not depends on the context of the situation. In the case of target systems, the context is the phenomenon we are trying to explain and the characteristics of the model or experiment that will explain it. The relevant information is the set of parts and properties that target selects from the actual system which will go into the model or experiment.

The relationship between target systems and our descriptions of them is similar to the relationship between things in the world and the terms we use to refer to them in natural language. For example, we might ask someone to bring “the glass on the table”. By doing so we are picking out a particular object in the world and saying something about its location. We leave out a lot of information, because we don’t consider it relevant. However, there might well be more than one glasses on the table, in which case our failure to specify which glass can lead to confusion. It would have been more apt to include some more information about the particular glass we were referring to, for example “the glass at the far end of the table”. Yet our failure to include this relevant information does not mean that we are failing to refer to the particular glass. We are simply using a description of it which is not particularly useful, given this context.



The same is true of target systems. Consider again the case of plant-soil feedback. Earlier target systems were rather inapt, as they did not include soil microbes. As it turns out, soil microbes *are* relevant for explaining plant competition, hence later targets, such as those specified in the work of Klironomos (2002) are more apt. It is possible to view these target systems as better than previous ones, but only in the sense that they help scientists explain a particular set of phenomena. They are neither better overall, nor more realistic *in themselves*. Target systems cannot be judged in terms of their aptness in isolation of their context, as they are just parts of the world. Whether or not a target system is apt depends on its relationship to the model and the phenomenon of interest. Of course, a description or representation of the target system which includes soil microbes and feedback is more realistic than another which does not include them. Still, the descriptions are descriptions of things in the world, which are parts of the world and hence real.

## 5.2 False Beliefs about Target Systems

Mistakes of relevance are not the only kind of errors that can occur in the specification of target systems. A different kind of mistake occurs when we manage to specify a target system, yet this specification induces or relies on false beliefs about the target. In the context of plant-soil feedback, this would occur if a scientist identified positive and negative feedback between the plants and soil biota, yet mistook mycorrhizal fungi for bacteria.<sup>8</sup> In this case the scientists would be specifying an actual target system, with particular plants and microbial biota, yet they would also have false beliefs about those biota (that they were fungi instead of bacteria).

Cases like the above suggest that we should maintain a conceptual distinction between target systems and our beliefs about them. Target systems are real parts of the world, but we can have mistaken beliefs about them. Going back to the analogy with language and reference, cases like these are similar to Keith Donnellan's water-in-the-martini-glass example (Donnellan, 1966). If someone looks at an "interesting-looking person holding a martini glass" and asks "Who is the man holding a martini?", they have succeeded in asking a question about the actual person in the room, even if the martini glass contains only water. In the case of target systems, the scientists are succeeding in specifying a target and asking questions about it, yet they have some false beliefs about it. These false beliefs might turn out to be

---

<sup>8</sup> This hypothetical situation might seem absurd, but it is in fact quite easy to mistake one for the other if the partitioning of the target system is coarse-grained enough, so that what is measured is photosynthetic output or total biomass.

problematic and require the re-specification of the target system, yet they do not change the nature of the target itself. The importance of the problems depend on what the target system will be used for and the manner in which it will be used.

### 5.3 Failure to specify a Target System

A more drastic type of problem occurs when scientists fail to specify a target system altogether. This might be the way to interpret the invocation of phlogiston to explain the processes of oxidation and reduction in the 18<sup>th</sup> century. The idea was that metals were composed of calx and phlogiston and would become dephlogisticated when burned (Weisberg, Needham, & Hendry, 2011). We now know that there is no thing with the properties of phlogiston. The question is whether scientists in the 18<sup>th</sup> century were specifying a target. If they were specifying a target, then this is a problem for my view, because they would be specifying something imaginary.

There are two ways in which we can interpret cases like these. The first is to show that they are, in fact cases like the ones in 5.2, where scientists are succeeding in specifying a target system, yet have false beliefs about it. Indeed, one way to interpret this 18<sup>th</sup> century science is to identify phlogiston as hydrogen, dephlogisticated air as oxygen and phlogiston-saturated air as nitrogen (Weisberg et al., 2011). The important point is that phlogiston was not part of the 18<sup>th</sup> century target system. Instead the scientists were specifying a target system including hydrogen, but mistakenly attributed properties to it that it did not have. In other words, the scientists succeeded in specifying a target but had false beliefs about the target.

Yet even if we think that 18<sup>th</sup> century scientists were simply failing to specify any actual target system, this does not jeopardize the nature of target systems in general. Failing to specify a target system is not the same as specifying a target system that does not exist. Sometimes scientists make big mistakes and are not talking about any real part of the world. Yet, as they continue to experiment and refine their views, they often do end up specifying an actual target, albeit with a number of mistakes and false beliefs. All this shows is that the natural world is complex and it sometimes takes centuries to even begin understanding what it is like and how it works.

## 6. Conclusion

Target systems are real parts of systems in the world. They have a part-whole relationship to natural phenomena, and are represented by models. Target systems are specified by the partitioning of a domain of study into units and the omission of those units that are not considered relevant for studying the phenomenon of interest. They are an important aspect of science yet their role and importance has hitherto been overlooked.

The next step is to extend the theory of target system evaluation. On my view, target systems are evaluated in terms of aptness with respect to a particular model. The issue is to determine the criteria for aptness in terms of partitioning and abstraction. Before I turn to the theory of aptness I will examine the metaphysical implications of partitioning and abstraction in more detail in order to provide comprehensive answers to the criticisms raised in this chapter.

### 3. Target System Metaphysics

#### Part 1: Partitioning

##### 1. Introduction

The marmots of Vancouver Island are living, breathing organisms. Does this mean that they are real? A common sense answer to this question is that of course they are real, they are as real as you or me. But what if a scientist is studying these marmots? Do interactions with scientists change their ontological status? Again, according to common sense, they continue to be real after interactions with scientists. Yet what if the scientist constructs a model of these marmots? Now common sense dictates that the ontological status of the marmots has changed. The model does not contain living, breathing marmots, but idealized representations of marmots. Thus, at some stage between the unobserved marmots on Vancouver Island and the construction of the model the marmots have ceased to be real. The issue for my account of target systems is to determine the stage at which marmots stop being real.

On my view, target systems are real, as target system specification does not alter the ontological status of a system in the world. Yet, in the previous chapter, I considered some criticisms of my view which argued that partitioning does change the ontological status of the system. If this is the case, then target systems cannot be considered real systems. In this chapter I will give a more extensive account of these criticisms and their motivations and show how they can be overcome. I will argue that the change in ontological status does not happen at the stage of partitioning, but in the construction of the model. I will start by giving an overview of the views which give rise to criticisms of my notion of partitioning by providing a discussion of the debate on natural kinds.

The goal of this chapter is not to argue for one particular view of natural kinds. Instead, I hope to show that one of the attributes of my view is that it is compatible with a number of views of natural kinds. In the following discussion, I will identify the characteristics which make views of natural kinds compatible with my notion of partitioning. I will then examine two moderate views, one realist and one

conventionalist and show how they are both compatible with my notion of partitioning. The key to this compatibility is that partitioned systems need to be as real, but only as real as pre-partitioned systems.

## 2. Natural Kinds

Any discussion of the metaphysics of partitions incorporates some aspects of the debate about natural kinds, as views about the ontological status of partitions stem, directly or indirectly, from particular conceptions of natural kinds. Kinds are classifications of the world. That is, they are groupings of individuals or particulars. Moreover, kinds tell us about the nature of individuals, as belonging to a kind defines what an individual is. If a kind is natural, then it exists as a classification in nature, independently of humans. Partitions are also classifications of a domain, as they are sets of individuals within a domain. Yet different views of natural kinds affect what can be counted as an individual for a partition, and in how many ways a domain can be partitioned.

A second reason for including an overview of the debate on natural kinds is that it is already well-established in philosophy. Moreover, it provides a framework for categorizing the various views. The dimensions along which views of natural kinds vary, have counterparts in views of partitions. My discussion of partitioning in this chapter will draw extensively on the framework for natural kinds provided by Bird and Tobin (2012).

There are two general debates about natural kinds. The first is metaphysical and concerns the nature of natural kinds, while the second is semantic, and concerns the meaning of natural kind *terms* and how they refer (Bird & Tobin, 2008). While the two debates are related, they are conceptually distinct, as metaphysical views of natural kinds do not determine the semantic content of natural kind terms and vice versa (ibid). Therefore, in my analysis, I will focus exclusively on the metaphysical debate.

A straightforward way to categorize metaphysical accounts of natural kinds is through two questions: (1) Are our classifications of the world determined by the world itself or are they merely conventional? (2) How many kinds does the world give us? <sup>1</sup>The two questions are straightforwardly metaphysical, as they refer to the nature of natural kinds. There is also an additional question, which is not obviously metaphysical but methodological, as it is concerned with how natural kinds are discovered: (3) What are the standards for determining whether a particular classification is a natural kind? Together, these

---

<sup>1</sup> These questions are based on the discussion in Bird and Tobin (2012) and on conversations with Michael Weisberg.

three questions reflect three dimensions on along which views about natural kinds can be placed. (1) reflects the dimension of *realism* versus *conventionalism*, (2) reflects the dimension of *monism* versus *pluralism* and (3) reflects the dimension of *strict rules* for determining natural kinds versus *no rules*.

## 2.1 Realism versus conventionalism

The first question is the most fundamental, and amounts to whether or not there are natural kinds. Generally, those who answer the question in the affirmative are *realists*, whereas those who answer it in the negative are *conventionalists* (or *constructivists*) (Bird & Tobin, 2008). These two categories are not as distinct as they might seem at first glance. Both realism and constructivism have stronger and weaker versions, and the weaker versions of each are often motivated by similar epistemic considerations. In fact, I think that the best way to understand this classification is as a continuum with *strong realism* and *strong conventionalism* as the two extremes.

Strong realism is the view that natural kinds are real entities which are ontologically distinct. On this view, each natural kind is an entirely different type of entity. Thus, for example, classifying one atom as hydrogen and another as helium is not just a natural way of distinguishing between lumps of matter, but it reflects a fundamental ontological division in nature. This point is important because it constitutes the basis of the distinction between strong and weak realism.

*Naturalism* is a weaker version of realism, which does not require the same ontological commitment (ibid). A naturalist maintains that there are natural ways of classifying the world, yet she does not necessarily think that natural kinds are different entities. For example, in the case of hydrogen and helium, a naturalist would agree that they are distinct natural kinds, but would also think that there is ontological continuity between them. In this sense, naturalism is compatible with *ontological reductionism*, a view which holds that objects are reducible to their smallest constituent parts (usually sub-atomic particles).

The realist dimension of naturalism comes from the view's association with scientific realism (ibid). Scientific realists believe that the world which scientists investigate is mind-independent and that scientific theories give us knowledge about the world (Godfrey-Smith, 2009). The idea is that successful scientific theories latch on to real entities and discover relationships between them, hence we can say that

we gain knowledge about the world. In fact, some scientific realists tie realism directly to the existence of mind-independent natural kinds (e.g. (Psillos, 1999)).

On the other side of the continuum lies *conventionalism*. In general terms, this is the view that the kinds picked out by science are conventions, and should not be privileged over other classifications. For example, while conventionalists might agree that angiosperms and gymnosperms are kinds which are important for science, they will also argue that vegetables and fruits are kinds, even though they do not match biological taxonomy. The latter classification is important, for instance, for dietitians and chefs.

As with the case of realism, there are weaker and stronger versions of conventionalism. Weaker conventionalists, such as *pragmatists*, do not deny the existence of natural kinds, yet argue that our *actual* classifications are not likely to be natural kinds, hence they should be viewed as conventions. Moreover, even if a particular classification happens to reflect a natural kind, we are most likely not in the position of perceiving it as such. It is interesting to note that the weak conventionalist's ontological stance towards natural kinds is actually rooted in epistemic considerations. I will examine these epistemic considerations and their relationship with ontological considerations in sections 2.3 and 2.4.

Stronger versions of conventionalism deny that there are any natural divisions in the world (Bird & Tobin, 2008). Hence, it is not possible to discover natural kinds, only to construct classifications which give us artificial kinds. The most extreme form of conventionalism is *ontological relativism*, the view that all classifications of the world are dependent on a particular conceptual scheme (Goodman, 1978). On this view, there are no natural kinds, only conventions. Each conceptual scheme has its own way of constructing reality, and science is just one of many. In fact, science can be further divided into disciplines or conceptual schemes, each with its own natural kinds. For example, on one reading of Kuhn, each scientific paradigm has its own method of constructing reality and its corresponding natural kinds (Kuhn, 1996).

## 2.2 Monism versus pluralism

The second dimension refers to the *number* of natural kinds that there are in the world. This dimension is also metaphysical as each view reflects a particular conception of how the external world is structured. In other words, the question is: how many joints does nature have? This dimension is more useful for classifying realist views, as it presupposes that there are natural kinds in the first place. At the

same time, however, it is also possible to rank conventionalist views in terms of the number of conventional kinds which are acceptable.

On one side of the dimension lies *monism*, the view that there is only one kind in the world. In order to be a monist one would have to believe that the world only has one true substance, while all other classifications are conventional. Moving along the dimension towards *pluralism*, each view accepts a larger number of kinds as natural kinds. A weak pluralist view is *ontological reductionism*, the view that the world is ultimately composed of a small parts (such as subatomic particles, strings, etc). Moreover, these are the only natural kinds, while all higher levels of organization are conventions.

The most pluralistic view is *promiscuous realism*. The term was coined by John Dupré, who argues the world is extremely complex and natural kinds do not form a neat hierarchy. Thus, the world can be categorized in many different ways, which can cut across traditional realist conceptions of natural kinds. Dupré's pluralism is motivated by an anti-reductionist view of the world. For example, he rejects the idea that sciences at different levels of organization, such as chemistry and biology are operating at a different level of the overall natural kind hierarchy. Instead, he thinks that the natural kinds of biology are not reducible to the natural kinds in chemistry. Moreover, he also thinks that scientists working within the same discipline can have different but equally legitimate categorizations. For instance, a population ecologist and an ecosystem ecologist can study the exact same ecosystem, but carve it up in very different ways. The point is that both ways are equally legitimate.

With the exception of promiscuous realism, realist views tend to cluster towards the lower end, or middle of the dimension. This is because realists think that natural kind classifications are fewer than conventional classifications. The idea is that conventional classifications are quite easy to create, whereas it is much more difficult to discover truly natural classifications. I think that this stems, at least in part, from the motivation for realism, namely that discovering natural kinds and basing our scientific and philosophical investigations on natural kinds helps to gain a deep understanding of phenomena in the world. This motivation highlights the importance of the epistemic dimension in the discussion of natural kinds, which I will turn to next.



### 2.3 The Methodological Dimension

The reason for including a methodological dimension in an otherwise metaphysical discussion of natural kinds, is that in many cases methodological or epistemic considerations provide the motivation for ontological commitments about natural kinds. As stated above, realist positions are motivated by the drive to understand which classifications are natural and which are conventional. This is often tied in with scientific realism, the view that the phenomena studied by science reflect natural divisions in the world. Hence, successful scientific divisions of the world are natural kinds and should be privileged over other divisions. This also means that the success of a scientific theory can be determined by whether it captures natural divisions in the world. Therefore, the epistemological objective for realists is simply to discover these natural divisions.

At the same time, commitments to conventionalism are also epistemically motivated. Conventionalists, especially pragmatists, argue that due to epistemic limitations, we are not in a position to discern which divisions are natural and which are not. Therefore, even if there are natural kinds, our scientific theories might not capture them. According to the pragmatists, this should not diminish the importance or value of scientific theories. However, it does mean that there is an additional methodological question, namely how to judge when a scientific theory is successful. There are many accounts of how to judge scientific theories in pragmatist terms and a discussion of these is beyond the scope of this chapter. However, a similar epistemological concern can be raised for the divisions that scientific theories employ. In other words, how should we judge when a division is successful? I will give a full account of my view in chapter 5; for now, I will give a general overview of the methodological dimension for natural kinds.

The methodological dimension to the natural kinds debate concerns the standards, or rules, by which classifications are judged. For realist views, it can be phrased as the question: (3) What are the standards for determining whether a particular classification is a natural kind? A traditional view is that natural kinds can be found by discovering the *essential properties* of a kind (cite). Once these are identified, then each particular can be classified based on these properties. Essential properties are those which make a particular object what it is. Therefore, it follows that if a particular has the requisite properties it is automatically part of the natural kind classification. On strict essentialist views, essential properties constitute the necessary and sufficient conditions for belong to a kind, hence each member of a kind must share all the essential properties of its kind and not share essential properties with other kinds. For example,

hydrogen and helium can be considered different natural kinds, because each kind has its own, unique atomic number. Therefore any element with atomic number 1 is classified as hydrogen, while any element with atomic number 2 is classified as helium.

Essentialism has been widely criticized. Critics argue that essential qualities are often difficult or impossible to locate. One of the most important set of criticisms of essentialism comes from biology, as critics of species essentialism argue that evolution precludes species essences (Hull, 1965). The processes of mutation, recombination and random drift can cause new traits to appear and others to disappear, thus members of a species will not necessarily share their traits. Moreover, members of a species often share a large number of traits with members of other species. A second argument is that the boundaries between kinds are often vague, hence it is impossible to distinguish between natural kinds. In the case of species, speciation occurs gradually, and there is no principled way of drawing boundaries between different species (ibid).

Realists must accept that we often make mistakes in our classifications. Some classifications are better than others, even if they seem successful at first. Even in the case of science, it is often very hard to determine whether or not a kind is discovered or if it is merely a convention. Many conventionalists think that this problem is in fact insurmountable. They argue that we simply do not have the tools to determine which classifications reflect natural kinds. However, this does not necessarily imply that all classifications are equally viable. In other words, being a conventionalist does not mean that anything goes in the selection of classifications. On the contrary, it is possible to have principled and strict ways of deciding whether a classification is appropriate, without appealing to the classification's ontological status.

The main criterion for judging classifications in conventionalist terms is usefulness. Intuitively, if a classification does not help with the overall goal of a project, then it is not appropriate for the project in question. For example, a biologist studying the growth of a particular population should employ a classification which includes organisms as units, as they are necessary for the experiments or models that she has at her disposal. A classification which has molecules as the basic units will not be useful in this instance.

The criterion of usefulness is not merely intuitive. In fact, most conventionalists devote a significant part of their accounts to setting out rules for judging the usefulness of classifications. This is usually achieved by determining what counts as acceptable and useful within a particular conceptual

scheme. Within each scheme, the criteria can be very specific and very strict. For example, according to Kuhn, different paradigms might have completely different ways of classifying phenomena in the world. They can also differ in the very concepts and methodologies they use to study phenomena and in the standards that they use to evaluate their theories (Kuhn, 1996). Moreover, these difference can make the contents of each paradigm unintelligible to the members of the other paradigm. Nonetheless, within each paradigm, there are strict rules for determining which classifications are appropriate.

The upshot of this section is that the methodological dimension to the natural kinds debate is important for both realist and conventionalist views. It often influences the ontological commitments towards kinds, yet it also shows that there are considerable differences within each camp in terms of how kinds are discovered or evaluated. In the next section I will provide a summary classification of different views of natural kinds, using all three dimensions.

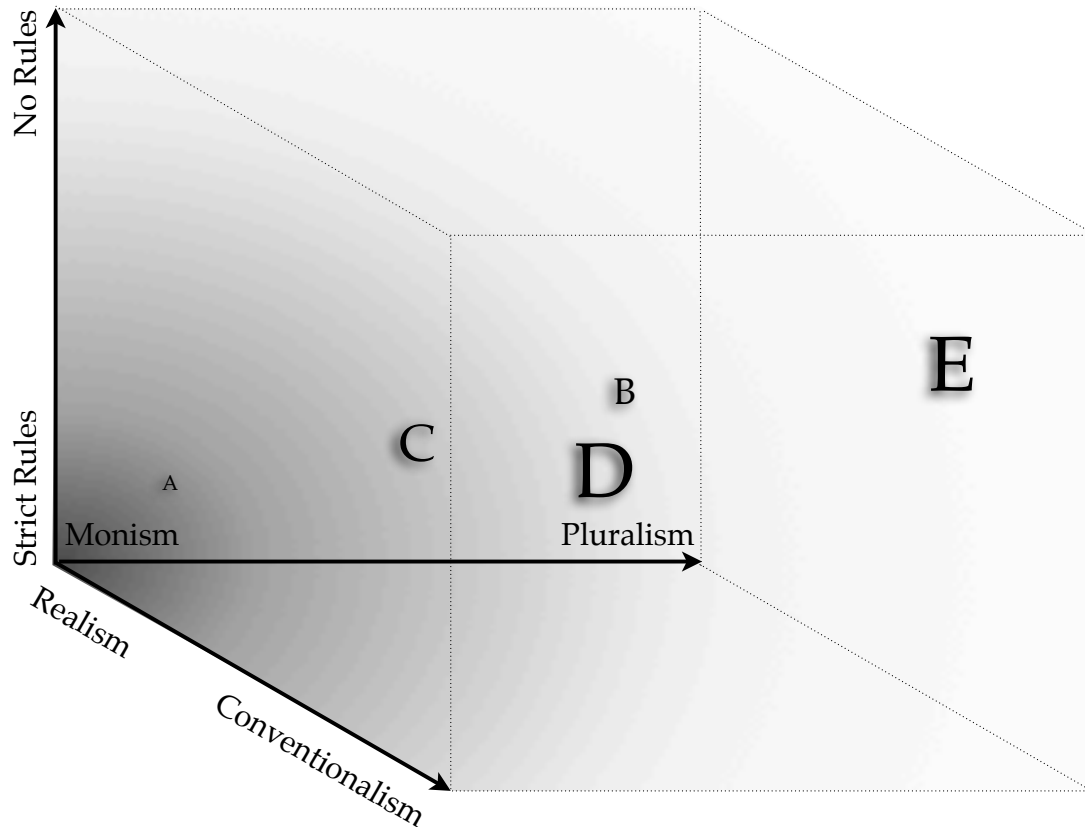
#### 2.4 Summary

To sum up the discussion of natural kinds, putting the three dimensions together provides a method for understanding each view and its relation to other views. Figure 1 shows a spatial representation of the debate in three dimensions, with each axis corresponding to one dimension. The figure also contains a selection of five views ranging from natural kinds realism (A) to ontological relativism (E).

The main point of this representation is to provide an easy way of visualizing the comparison of the different views. Thus natural kinds realism (A) is situated close to the starting point of all three axes, as it is a realist view, it is minimally pluralistic, and has strict rules for determining natural kinds. Promiscuous realism (B) is much more pluralistic and slightly less realist. In addition it also has slightly more relaxed rules for determining natural kinds. This is because the pluralism allows for many more kinds to count as natural kinds. In contrast, pragmatism (C) is conventionalist, yet has stricter rules for what can count as a kind. This is because it is less pluralist, hence there are fewer kinds overall. The Kuhnian view of paradigms (D) is more conventionalist, yet also more pluralist. However, it also has stricter rules for determining kinds, because even though the number of kinds are determined by each paradigm, within each paradigm the rules are exceptionally strict. Finally, ontological relativism is very conventionalist and pluralist, as each conceptual scheme determines its own kinds, yet the rules governing this determination

are not as strict as they are on other views. This is because the view allows for conceptual schemes which contain arbitrary classifications of kinds.

**Figure 3.1. Spatial representation of natural kinds views**



**Figure 1. Spatial representation of natural kinds views in three dimensions.**  
 The x-axis is the monism-pluralism dimension, the y-axis is the strict rules-no rules dimension and the z-axis is the realism-conventionalism dimension. The size of the letters is meant to reflect the three dimensional perspective, hence larger letters represent more conventionalist views and smaller letters represent more realist views.  
 Legend - **A**: Natural Kinds Realism (Armstrong) **B**: Promiscuous Realism (Dupré)  
**C**: Pragmatism (Kitcher) **D**: Kuhnian Paradigms **E**: Ontological Relativism (Goodman)

This leads to the second reason for including this diagram in my discussion of natural kinds. I hope to have shown that the three dimensions are related yet not intrinsically so. That is, a particular position on any one dimension can influence a position on the other dimensions, yet there is some leeway. Thus, for example, realist views tend to be less pluralistic than conventionalist views, yet promiscuous

realism is an exception. Similarly, while some conventionalist views do not have strict rules for determining natural kinds, this is not a necessary implication of conventionalism. This point is very important for the following discussion of partitioning and its relation to natural kinds, as the subtle similarities and differences between views determine whether a particular view of partitioning is compatible with my view of target systems as partitioned parts of the world. I will give an account of which views are compatible with target system partitioning in the next section.

### 3. Natural Kinds and Partitioning

In the previous section I gave a brief overview of the various positions on natural kinds. Views differ along three dimensions: in terms of realism, in terms of the number of natural kinds that the world gives us and in terms of the appropriate methodology for acquiring natural kinds. I will now turn to the relationship between natural kinds and partitions of target systems.

Natural kinds and partitions are intimately connected, as views of natural kinds imply views of partitioning. For example, a realist about natural kinds will partition the world in terms of those natural kinds. Thus, someone who thinks that hydrogen and helium are two different natural kinds will employ a partition which separates the two kinds of entities. Often, this partition will be hierarchical, with each level in the partition corresponding to a different natural kind. However, it is also possible to be a realist and allow for more than one real partition. For example, Dupré, who is both a realist and a radical pluralist about natural kinds, allows for many different and crosscutting partitions to exist at the same time.<sup>2</sup>

Allowing for many partitions to coexist is more closely connected with conventionalist views of natural kinds. However, on these views, the partitions themselves do not need to reflect kinds that are given to us by the world, but conventional kinds. Hence, on many conventionalist views, partitions are not real, but also conventions. Often, these partitions will be determined by a particular conceptual scheme.

The methodological dimension of the natural kinds debate also has a corresponding dimension in partitioning. It is especially important for pluralist views of partitioning, as determining which partitions are appropriate is key. This is true both of conventionalist and realist pluralist views. In each case, a substantial part of any theory of partitions must give an account of how to evaluate partitions, and how to

---

<sup>2</sup> It is important to note that the single hierarchical partition view stems from moderately pluralist views about natural kinds. A monist realist about natural kinds, i.e. someone who thinks that there is only one natural kind, would argue that it is not possible to partition the world into real parts.

distinguish between appropriate and inappropriate partitions. I will provide a full account of target system evaluation in chapter 5. Methodological considerations are also important for less pluralist views of partitioning, yet in a different way. In minimally pluralist, realist views of partitions, the methodological question is to decide which of the partitions reflect natural kinds, and which are merely conventions. However, on these views, the number of partitions is already restricted to the ones that have already been determined to be real. Thus, in some sense, the methodological dimension is less interesting for these views.

It is worth, at this point, to take a step back and take stock of the situation. The debate on natural kinds shows that there is an important tension between realism and conventionalism for my view of target system partitioning. On my view, target systems are real parts of the world, yet there are also many different types of target systems, corresponding to different scientific disciplines or phenomena. Realist views of natural kinds imply that partitions are real, yet do not usually allow for a plurality of cross-cutting partitions. On the other hand, conventionalist views of natural kinds allow for cross-cutting partitions, but at the cost of reducing partitions to conventions.

This diagnosis implies that not all views of natural kinds are compatible with my view of partitioning for target systems. Views that lie on the two extremes of the realism-conventionalism dimension will not be compatible with my view of partitioning, for the reasons outlined in the previous paragraph. More specifically, strong realist views tend not to be pluralistic enough to allow for cross-cutting partitions of domains in the world. For example, a very strong ontological and epistemic reductionist, will not accept that higher-level organizations are legitimate partitions, hence most partitions which are regularly employed by scientists (e.g. biologists) will not count as real. This means that the target systems will also not count as real.

At the same time, strong conventionalist views, such as ontological relativism are also not compatible with my view of partitions. This is because, by definition, no partitions are considered real. Consequently, on this view, no target systems employed by scientists can be considered real. Perhaps more importantly, ontological relativism does not have sufficiently strict rules for governing categorizations. Extending this view to partitioning would mean that an extremely large number of partitions are possible even within a particular conceptual scheme. This is not a problem in itself, yet without rules for determining which partitions are legitimate for a particular purpose, we have no method for evaluating

target systems. Being able to evaluate target systems for a particular phenomenon of interest is crucial, as without it the importance of spending time and effort to specify target systems is greatly diminished.

Fortunately, the number of views that are truly incompatible with my view of partitioning is quite small. Most of the remaining views are moderate enough to be at least sufficiently compatible with my view of partitioning for target systems. An examination of the compatibility of each view with my own would be laborious but not particularly useful. Instead, in the remainder of the section, I will give a detailed examination of two moderate views, *promiscuous realism* and *pragmatism*. In each case, I will give a more detailed overview of the view itself, show how it is compatible with my view of partitioning and consider some criticisms which can be leveled against it. The motivation for this particular choice is twofold. The first reason is that these are the two views which fit most naturally with my view of partitioning. The second reason is that these two views are rather exceptional, as they combine elements of traditionally realist and traditionally conventionalist views. Thus, in considering the main criticisms leveled against these two views, I will be addressing the main criticisms which can be used against other views. Thus, I hope that the following discussion will also imply how other views of natural kinds can be made compatible with my view of partitioning.

### 3.1 Promiscuous Realism

The most obvious candidate for compatibility with target system partitioning is promiscuous realism. Promiscuous realism, according to its advocate John Dupré, is a realist view of natural kinds, combined with a “metaphysics of radical pluralism” (Dupré, 1995, p.18). In order for the combination of realism and pluralism to work, Dupré argues for a weaker understanding of realism and a weaker conception natural kinds.

Dupré’s argument starts from the observation that the classifications picked out by ordinary language terms are at odds with the natural classifications picked out by scientific practice. For example, a carpenter might group together pines, oaks, chestnuts and cherries because these are trees which provide the right type of wood for making furniture. This classification is entirely at odds with the classification of a botanist, as these are all examples of different species of tree. In fact, a botanist would be especially irked at the inclusion of pines in this group, as pines are gymnosperms, while the rest are angiosperms. Angiosperms are a much younger group of plants which only emerged at the end of the Jurassic period

(Schulze, Beck, & Müller-Hohenstein, 2005). To put this into perspective, dinosaurs and early mammals were already roaming the earth when the flowering plants started their evolutionary life.

Dupré asks why the botanist's classification should be privileged over the carpenter's classification. One motivation for doing so is the belief that the scientific classification reflects natural kinds. However, according to Dupré, there are two ways of understanding the term 'natural kinds'. The strong version is the view that particular properties determine membership in a natural kind. In other words, the strong version ties natural kinds with essentialism. Dupré argues that any sort of essentialism is untenable for a number of reasons, including that it is incompatible with an evolution. Instead, he prefers a weaker version of natural kinds, defined as classes of things which are suited to particular roles. Moreover, he argues that it is a *natural fact* that these classifications exist and are mind-independent. That is, the carpenter's classification of trees which produce wood that is suitable for making furniture is discovered rather than invented, because it is a natural fact that some trees produce the suitable kind of wood and others do not.

Dupré concedes that it is possible to have a somewhat more restricted view of natural kinds, by thinking that some kinds are more natural than others. On this view, the carpenter's classification would be less natural than the botanist's classification. However, he argues that what makes a natural kind 'more natural' comes down to essentialism. Moreover, even if we accept this, the 'true' natural kinds would be few and far between, hence most of scientific practice would be operating with 'less natural' kinds. If this is the case, then we have no a priori reason to privilege the botanist's classification (which does not rely on essences) over the carpenter's classification. The only way to distinguish between the two is to determine which is better *for the purpose at hand*.

Perhaps more importantly, Dupré argues that pluralism is as plausible within science as it is between science and the rest of language. He states that "science, construed simply as the set of knowledge-claiming practices that are accorded that title, is a mixed bag. The role of theory, evidence and institutional norms will vary greatly from one area of science to the next" (Dupré, 1995, p. 242). It is even more difficult to distinguish between more and less natural kinds within scientific practice, hence the only way to evaluate classifications is by taking into account their context and intended function. Thus, promiscuous realism is both realist (though weakly so) and pluralist, as it allows for a number of different, cross-cutting classification to coexist and enjoy the same objective status.



The combination of realism and pluralism is what makes promiscuous realism naturally compatible with my view of partitioning for target system specification. Partitioning is a way of classifying a domain, just like any other. In the previous chapter I argued that partitioning amounts to choosing one set from the power set of a domain. This can be understood in the language of promiscuous realism in the following way. The power set is the totality of subsets of a domain, hence a particular partition is picked from the power set. In other words, the power set gives all the possible classifications (or natural kinds) from which we choose a subset. This means that partitions of a domain are objective, natural and mind independent. In other words, partitioned systems are real, because they pick out objective natural classifications in the world.

Of course, this does not imply that anything goes in partitioning. First, the domain itself restricts partitioning, as it precludes classifications which include objects outside the domain. In addition, the power set itself precludes some classifications. For example, each object can only be part of one subset, otherwise the addition of subsets would produce additional objects. More importantly, evaluating a partition in terms of its context and function actually amounts to strict rules about which partitions are viable and which are merely logically possible (I will give a more extensive account of evaluating partitions in chapter 5).

The downside to the compatibility of promiscuous realism with partitioning, is that some of the arguments against the former also apply to the latter. Here, I will focus on two criticisms of promiscuous realism, (1) that it is too promiscuous and (2) that it cannot be understood as a true realist position.<sup>3</sup>

The first criticism can be associated with the third, methodological dimension of natural kinds. It is the view that promiscuous realism allows for too many natural kinds and hence too many partitions. The problem with promiscuity is that it does not allow for a clear-cut and universal way to distinguish between appropriate and inappropriate classifications. In other words, as soon as we allow many classifications to count as natural, then we lose an important method for evaluating classifications. The framework for this criticism is realism. It is the same criticism that realists use against conventionalist views of natural kinds. One of the supposed attractions of strong realism is that the classification between natural and conventional kinds gives us a clean way of determining which classifications are appropriate, at least in scientific contexts. Hence, discovering which classifications are natural and which are not is a unifying goal of all

---

<sup>3</sup> I should note that the most famous criticism of Dupré's book comes from Rob Wilson {Wilson:1996ju}. However, as this criticism focuses on Dupré's treatment of the species debate and not on the general metaphysical implications of promiscuous realism, it does not apply to partitioning and so it is not relevant for the purpose of this discussion.

scientific practice. However, pluralism about natural kinds makes the difference between natural and conventional kinds fruitless.

I do not think that this criticism is a serious threat to promiscuous realism. Dupré's argument begins from the observation that there already exist different and cross-cutting ways of categorizing nature. Unless we accept (which I do not) that entire scientific disciplines or sub-disciplines are merely artificial, then we must accept that more than one classification is possible. But if we accept that more than one classification is possible, then there is no reason to think that we should accept few rather than many classifications.

A second counterargument is that a strong realist understanding of natural kinds seems to be untenable. Dupré shows quite convincingly that essentialism is problematic, as it is extremely difficult, if not impossible to discover the essences of natural kinds. Yet, as I stated above, without essentialism, the strong notion of natural kinds becomes untenable. This only leaves the weak notion of natural kinds, which is compatible with pluralism. Again, if modest pluralism is adopted, then there is no reason to think that radical pluralism should not be.

In the end, Dupré concedes that it may turn out that there are fewer natural kinds than his view seems to suggest. Yet he cautions that this is an empirically discoverable fact which cannot be determined a priori. He therefore urges that we adopt a 'categorical empiricism' and genuinely investigate whether "the discovery of a natural kind adds little, if anything, to the discovery of whatever correlations may turn out to characterize it" (Dupré, 1995, p. 80). In other words, before attributing essences to natural kinds, we should empirically determine whether we have any good reason to do so. The upshot of this concession is that we have good empirical reasons to adopt at least a moderate pluralism but also that we do not have good empirical reasons to deny a more radical pluralism. Still, if a moderate pluralism is all that the critic is ready to accept, then this is sufficient for a change in the way we think about natural kinds.

A third counterargument is that it is simply not the case that anything goes in the classification of natural kinds. As I argued above, there are both general and domain-specific reasons for favoring one classification over another. The fact that much of this evaluation is context dependent and determined by the usefulness of a classification for a particular purpose, does not mean that the rules *within the context* are not strict. Within a particular context there can be a fact of the matter about which classification is the best. Moreover, Dupré's view does not assume incommensurability between contexts, in the Kuhnian sense.

Integration between scientific disciplines is compatible with promiscuous realism, as long as there is an overarching common purpose. Hence it is possible to translate between different scientific disciplines and to decide that a particular classification is more appropriate given the particular problem which needs to be solved.

Finally, given our epistemic limitations and the untenableness of essentialism, it seems that determining the usefulness of a classification is an indispensable part of any evaluation. In fact, this is exactly how a lot of scientific practice actually operates. Thus, Dupré can lay the burden of proof on the strong realist, and argue that until a tenable way of distinguishing between natural and conventional kinds is discovered, then the more pragmatic approach is the only one which is available. Given this situation, our theories about the structure of the world should reflect these epistemic limitations.

The second criticism is much stronger. It also has a realist motivation and is related in some ways to the first criticism. This criticism denies that such a pluralistic view can be considered a truly realist view. It is connected to the first criticism because it focuses on the incompatibility of realism with radical pluralism, but instead of disputing the pluralistic aspect of the view, it disputes Dupré's claim to realism. The locus of the problem lies in Dupré's definition of a natural kind as a class of things which have a particular role. The realist will argue, for instance, that the carpenter's classification is wholly conventional, and that there is nothing *in nature* which makes that classification objective or mind-independent. In other words, if there were no carpenters, the classification would not exist.

This can also be applied to cross-cutting classifications within science. Realists can argue that some scientific classifications are just better than others and that this is the only way to make sense of scientific progress. For example, even within the species debate, the fact that Linnaean taxonomy is generally thought to be outdated and has been replaced with a cladistic notion of taxonomy shows that the latter classification is simply better than the former. Moreover, there are objective grounds for this claim, for instance, that the latter is compatible with evolution while the former is not. In other words, Linnaean taxonomy does not reflect natural kinds, and thus cannot be seen as a real classification.

The worry is that on Dupré's view, Linnaean taxonomy is an objectively legitimate way of classifying the world which is as real as the cladistic taxonomy. If Dupré insists that the Linnaean taxonomy is on a par with cladistic taxonomy then, according to the realist, Dupré's view misconstrues a

conventional classification as a real one. If this is the case, then the view cannot be considered a realist view.

As with the previous criticism there are a number of responses available. The first is a partial concession to the criticism, namely the reiteration that Dupré's view is incompatible with strong realism. As stated above, any form of realism which invokes essentialism cannot be reconciled with promiscuous realism. The second counterargument rests on a distinction between objective classifications and objective grounds for favoring a classification.

According to Dupré, the motivation for this criticism is the same as that of the previous criticism, and is methodological. That is, the realist insists on differentiating between real and conventional categories in order to have a way of choosing one type of classification over another. However, Dupré thinks that there are other, better ways of distinguishing between classifications, hence there is no need to question the reality of some classifications. In other words, we have objective (though pragmatic) reasons to favor cladistic taxonomy over Linnaean taxonomy and the botanist's taxonomy over the carpenter's taxonomy, when we are studying particular scientific questions, so we do not need to privilege any taxonomy on ontological grounds.

The same criticism can also be understood in a conventionalist manner. That is, a conventionalist could argue that the pragmatic motivations which give rise to pluralism, should force us to reject realism. This criticism has been applied to another view of natural kinds, which I will discuss in the next section. Both the criticism and the response are applicable to promiscuous realism.

### 3.2 Pragmatism

Pragmatism towards natural kinds can generally be understood as the view that we should take a pragmatic approach toward classifications. While pragmatists do not deny that there may be natural kinds, our epistemic limitations are such that we are not in a position to discover them or to know when we have discovered them. Therefore, we should use criteria other than realism to evaluate our classifications of the world.

The view which I will be considering is Philip Kitcher's pragmatist understanding of science. Kitcher's view is interesting because it started off as a more traditionally realist view (Kitcher, 1993) yet progressed towards a more pragmatist view over time (Kitcher, 2003; 2013). Strictly speaking, Kitcher's

view is a realist view, as he self-identifies as a realist (Kitcher, 1993), and others accept the identification (Stanford, 1995). However, there are two reasons why I think that Kitcher's view counts as a pragmatist view.

The most important reason is that the motivation for Kitcher's view is thoroughly pragmatist. In later writing, Kitcher identifies himself as a pragmatist in general terms, because his view of science shares the important pragmatist themes "in the emphasis on local studies, on introducing considerations of value, in treating Science (and other institutions) in relation to their goals." (Kitcher, 2013) (p.229). This pragmatist approach to science in general, informs his view of natural kinds. He states that "invocations of "genuine properties", "divisions in nature", even of "objective similarity" are hardly pellucid. Because I distrust these notions, I am suspicious of the enterprises in which they are employed" (Kitcher, 2003) p. 44. Instead, he advocates a pragmatist view which "allows a place for human values and human interests in the constitution of the goals of the sciences" (ibid).

The second reason is that Kitcher thinks that our classifications of the world are not wholly mind-independent. He argues that thinking of classifications as more natural than others, is tantamount to thinking that nature somehow has it's own language which picks out particular objects and sets boundaries around them. In other words, we discover natural kinds when our language (everyday, scientific etc.) picks out the same objects as "nature's own language" (ibid, 46). However, this matching presupposes an already anthropocentric perception of nature. The very phenomena or objects that we care about have something to do with the kind of creatures that we are (ibid).

For example, different elements in the periodic table are thought to constitute natural kinds, because of certain microphysical structures. That is, the atomic number of each element is supposed to have some fundamental set of properties which make it similar to other tokens of the element and different to tokens of different elements. Moreover, these properties have important causal implications as they affect the way in which each element behaves. However, Kitcher argues that "the fact that these microstructures play a systematic unifying role, depends on the *prior* identification of a class of manifest properties". We only notice these similarities between elements and focus on certain aspects of the world because they intrigue us. This interest lies in pragmatic reasons, such as "our sensory and cognitive capacities ... interests that people have developed, either naturally or as a result of the accidents of human history" (ibid 50).

A more traditional conventionalist would applaud these conclusions, however, they would also argue that Kitcher cannot be called a realist.<sup>4</sup> After all, Kitcher goes a step further and states that “different ways of dividing nature into objects will yield different representations of reality” (ibid 47). Moreover, he states that he is skeptical of the idea that there is an overarching aim for scientific inquiry, which is based on discovering natural kinds (ibid 59). However, this is entirely compatible with a conventionalist view of natural kinds and an instrumentalist view of science.

Nonetheless, there are two important differences between Kitcher’s view and full-blown conventionalism. The first is that, unlike the true conventionalist, Kitcher believes that there is sufficient continuity between different classifications to allow us to evaluate them using criteria that are not internal to each context. According to Kitcher, science can progress towards a better understanding of the world. Thus, for example we can say that *for the purposes of biology*, a classification which is compatible with evolution is better than the Linnaean taxonomy. Of course, the true conventionalist will argue that the notion of scientific progress and the discovery of truth is a futile endeavor, which should be abandoned (Stanford, 1995).

The second difference, is that Kitcher argues for a distinction between the world itself and our representations of it. The point is that different representations of reality can coexist but they do not change reality itself. That is, the fact that a chef might think of a tomato as a vegetable whereas a botanist thinks of it as a fruit does not in any way change the ontological status of the particular tomato. While the conventionalist will argue that this point ought to push us away from realism, Kitcher thinks that this is what ultimately grounds us in reality.

For Kitcher, our representations of the world are necessarily conventional, as they reflect our goals and purposes. This includes our scientific representations. However, to deny that the representations actually represent the real world implies that the world itself does not exist. Of course, the (moderate) conventionalist would agree with this point, as she would not want to align herself with extreme idealism. Yet if we agree that our representations of the world are representations of real objects, then according to Kitcher, we are ultimately trying to understand the real world. Thus, science aims to understand how the world really is, despite our epistemic limitations and our pragmatic ways of evaluating our categorizations.

---

<sup>4</sup> In fact, this criticism has been put forth against Kitcher’s view of species as natural kinds, even before his turn towards pragmatism, see {Stanford:1995tl}.

This discussion shows that the debate reaches an impasse. This should not be surprising, as these are the arguments which form the two sides of the debate on scientific realism and antirealism, which is far from being solved. This is, in part, because in the moderate sides of the two camps, the same epistemic observations and motivations lead to different metaphysical conclusions. However, this is not a problem for my view of partitioning, as it just shows that both moderate realist and conventionalist views are compatible with my view of target systems. I will show why this is the case in the next section.

### 3.3. Compatibility

The discussion of target system metaphysics in terms of partitioning, was meant to show that moderate views on both sides of the debate are actually quite similar, as they stem from the same epistemic or pragmatic motivations. Whether or not there are natural kinds, they are hard to locate, hence we should adopt criteria other than realism to evaluate our classifications.

In the case of target systems, saying that target systems are *real* is compatible with both realist and conventionalist views of natural kinds. This claim is based on a simple distinction. What matters for target systems is that the partitioned system is *as real* as the pre-partitioned system. In other words, partitioning does not change the ontological status of the domain. This should not be a problem for the realist *in principle*, because she would accept that at least some partitions are real. Hence, it is not the action of partitioning itself that changes the ontological status of the system. The problem, for the realist, is that only some of these partitions should be considered real. The ones that do not reflect natural kinds should be considered conventional.

The response to this worry is similar to that provided by Dupré'. I could concede that the term 'target system' be restricted to those partitions which pick out natural kinds. However, like Dupré', I think that we do not have good reasons to think that we are in a position to determine which partitions reflect natural kinds and which do not. Hence we should adopt a pluralistic approach to partitioning. I agree with the realist that this should not imply a methodological or ontological relativism, because there are strict rules for evaluating partitions, as I will show in chapter 5.

However, if I concede that partitions are constructed based on pragmatic motivations, both the realist and the constructionist might argue that partitioning does alter the ontological status of the system. However, I think that if we go down this route, we cannot restrict conventionalism to partitioning. If it is

the case that our scientific partitions are conventional, then our non-scientific classifications are also conventional. The argument for conventionalism includes our basic perceptions of the world. This means that any time we perceive a domain in the world, we are doing so under the influence of conventions. As scientists are humans, even their basic perceptions of a domain will be influenced by conventions. Yet this means that the pre-partitioned systems are also conventional, hence partitioning does not alter the ontological status of the system.

I will illustrate this point with an example. The biologists who studied the population growth of the Vancouver Island Marmots (see Introduction), decided that Vancouver Island was the domain in which the phenomenon of interest (population growth) manifested. We can easily imagine that another group of population geneticists from the same academic institution, were also present on the island at the same time, studying a particular trait of the marmot population. Their partitions of the domain will be very different, as each group of scientists is studying a different phenomenon. For instance, the population ecologists will partition the domain into marmot-sized objects whereas the population geneticists will partition the domain into marmot alleles.

There is no reason to think that one of the two partitions is more natural or more real than the other. I am assuming at least a modestly pluralistic view, where evolutionary biology and ecology are counted as sciences and have the same status in terms of their classifications. I do not think that this is particularly controversial, as it would take a strong version of reductionism to deny true scientific status to population ecology. Having established that both marmots and their alleles exist and are real, I will now ask what the pre-partitioned domains look like to each of the groups of scientists.

I think that it is safe to say that the pre-partitioned domains of the two groups of scientist will be very similar, if not identical. Their interactions with the pre-partitioned system is based on sensory perception. That is, their eyes, ears and language are sufficiently similar so that they perceive Vancouver Island in the same way. It is only when they start thinking of their domain, through the lens of their discipline that decide on different partitions. That is, it is the particular scientific discipline and the phenomenon which each group is studying which affects how they partition the system, yet in both cases the pre-partitioned system will be categorized in the same way. The question now is, what is the ontological status of the pre-partitioned system?



According to the conventionalist, the pre-partitioned system should also be a convention, as all categorizations are conventional. However, on this view the partitioning itself will not make the system any less real, hence we can say that the two target systems are *as real* as the pre-partitioned system. On the other hand, the realist will think that the pre-partitioned domain is real, yet this means that the partitioned system will also be real. In fact, most realist positions would argue that scientific partitioning is better at latching on to natural kinds than our everyday sensory perception. The point of doing science is to find the tools which help us locate natural kinds. Yet this means that the scientific partitions are *more real* than the pre-partitioned domains.

To summarize, my view of target system partitioning is compatible with both moderately realist and moderately conventionalist positions. This is because, the pre-partitioned systems and the partitioned systems have the same ontological status, hence target systems are as real as their pre-partitioned domains. At the same time, my view is not compatible with more extreme versions of realism and conventionalism. Realist views that are not sufficiently pluralistic will not allow target systems of accepted scientific disciplines, and therefore are not compatible. On the other hand, strong conventionalist views which do not have strict rules for determining classifications will also not be compatible with my view.

#### 4. Conclusion

In this chapter I examined the metaphysical questions concerning the partitioning of target systems. I hope to have shown that partitioning itself is a metaphysically neutral operation, which does not change the ontological status of a system. This should quell worries of both realists and conventionalists regarding the ontological status of target systems, as they are (only) as real as the pre-partitioned system.

This discussion does not exhaust the metaphysical and epistemic questions regarding target systems. As I have shown, it is important to be able to evaluate our target systems, in order to avoid the problems associated with ontological or methodological relativism. I will provide the framework for a theory of target system evaluation in chapter 5. Before I can turn to target system evaluation, however, I must first examine a different set of worries, this time pertaining to abstraction. In the next chapter, I will examine the notion of abstraction, starting with a brief overview of some common notions of abstraction in philosophy of science, considering some important criticisms and finally providing an account of abstraction which avoids them.

## 4. Target System Metaphysics

### Part 2: Abstraction

#### 1. Introduction

The west pediment of the Parthenon is a physical object that exists in space and time, but it is also triangular. We say that the west pediment is concrete, but that triangles are abstract. What accounts for this difference? The term ‘abstraction’ is used in many academic disciplines and everyday parlance, and is generally used to denote objects and systems that are not concrete or physical. People usually think of mathematical objects such as numbers and triangles as abstract, because they are not physically instantiated. Examples of abstract objects in academic disciplines are concepts and ideas which are not tangible (e.g., fairness, evil, superego). This intuitive notion of abstraction is also the received view in philosophy of science, which I call the *material view* (Cartwright 1994). Proponents of the material view think that the process by which we get abstract objects is one of omission. For example, we can start off with two roses, omit properties such as color, smell, photosynthetic capacity, chemical composition and so on, until we arrive at the number two, an abstract idea. Historically, philosophers writing on abstraction (e.g. Aristotle and Locke) have held versions of the material view and explained how we arrive at abstract objects through a process of omission (Rosen 2009, Cartwright 1994).

The problem with this view of abstraction is that it suffers from an important criticism, namely that it is indistinguishable from idealization. The criticism is based on the view that the process of abstraction misrepresents the system under consideration, as it introduces distortions. As distortions are the hallmark of idealization, abstraction is a kind of idealization. I call this the *distortion view* of abstraction. If proponents of the distortion view are right in their understanding of abstraction, then as target systems are abstract they are also distortions of natural systems. This means that they cannot be considered real. Yet there is an alternative view of abstraction, the *omission view* which focuses on the process of abstraction by omission and does not make claims about the ontological status of an abstract system (Jones 2005). A

version of the omission view can overcome both criticisms of the material view of abstraction, and provide the conceptual space for a view of target systems as abstract, concrete and real.

In this chapter, I will present the three existing views of abstraction in philosophy of science. I will examine the criticisms leveled by proponents of the distortion view against the other two views and show that a revised version of the omission view can overcome them. This revised version highlights the distinction between three aspects of abstracting, omission, de-concretization and generalization, and argues that omission is sufficient for abstraction. Moreover, the process of omission is not itself distorting, hence abstract systems can be concrete and real. I will argue that target systems are examples of systems which are concrete particulars, though incomplete with respect to the full natural system. I will then consider some cases which at first glance seem difficult for my revised view of abstraction as omission, and I will show that they can actually be used to support my view.

## 2. Abstraction in Philosophy of Science

### 2.1 The Material view of Abstraction

The main proponent of the material view of abstraction is Nancy Cartwright. Her view is presented in a chapter titled 'Abstract and Concrete' (Cartwright 1994). She distinguishes between two uses of the word 'abstract' (Cartwright 1994 197) and thinks that philosophers are prone to conflating the two. This results in the more general conflation of abstraction with idealization. The first use of the terms comes from physics. It was used by Duhem and Kelvin as a synonym for 'symbolic' and described systems in mathematical physics that could not represent the world accurately. It arose from the need to give simple representations of complex natural phenomena. A concrete particular phenomenon may be described by a set of mathematical formulae. However, taken literally, these formulae do not represent that particular phenomenon but an imaginary one. For example, the sun is often represented as a geometrically perfect sphere, even though its surface is highly irregular and has huge protuberances (Duhem in Cartwright 194). The best (simplest/most easily comprehensible) way of describing phenomena such as the motion of the sun from a particular point on the earth, is to construct geometrical formulae which describe spheres, not gaseous giants with highly irregular surfaces.

Cartwright does not think that this process is best thought of as abstraction. It is closer to what we think of as idealization. When we idealize, we start off with a concrete object and “mentally rearrange some of its inconvenient features -some of its specific properties- before we try to write down a law for it” (ibid 187). Her paradigm example is that of a frictionless plane. We do not just delete factors of a plane that cause friction, but we replace them with others that make our calculations easier. In the example of the sun, we do not smooth out the sun’s surface, but replace the actual sun entirely with a gigantic sphere, complete with all its geometrical properties. We then construct formulae that represent the sphere, not the actual sun.

True abstraction, for Cartwright, corresponds to the second use of the term. It is an Aristotelian notion which involves subtracting properties of an object or system instead of changing them. As with idealization, we start with a concrete particular but then “we strip away, -in our imagination- all that is irrelevant to the concerns of the moment to focus on some single property or set of properties ‘as if they were separate’” (ibid 197). The object or system in the world has all its properties. All we need do when we abstract is to identify the properties, disentangle them (separate them from each other) and dismiss those that are considered unimportant or irrelevant for the particular model or experiment. An example of a concrete object is a triangle drawn on a slate. We can abstract by stripping away from the object various properties, such as the color of the triangle, the chalk that it is drawn from “and other properties incidental to being a triangle” (ibid 213). The now abstract system consists of the remaining properties, the geometrical properties of triangles in general.

Both the processes of abstraction and idealization have their opposites. The opposite of idealization is de-idealization, and it is achieved when the features that were rearranged in our representation are systematically put back as they originally stood, as we learn more about the object we are studying. The opposite of abstraction is concretization. This process involves adding back properties that were previously omitted. Though de-idealization and concretization may seem similar, they have some important differences. De-idealization is itself an important goal for modeling. We are supposed to de-idealize when we can, because models that contain fewer misrepresentations are better representations of real systems. On the other hand, we don’t need to concretize when we model. A concrete model is not particularly useful, because it contains a lot of irrelevant information. Concretization is useful in other parts of science, for example, when applying the results of a model to a particular system. In other words, while

de-idealization is ultimately a goal of models themselves, concretization is not part of modeling. Another important difference is that concretization is easier to achieve than de-idealization. This is because adding real properties to a set of other real properties is easier than converting idealized properties to real properties.

An implication of Cartwright's account of abstraction is that we can rank objects in terms of their level of abstractness. A concrete object is thought to be closer to 'substance' in the Aristotelian sense than an abstract object (ibid). This is because the concrete object still has all the inessential (accidental) properties, such as color etc. However, merely counting how many properties have been omitted is not always a good method for ranking in terms of abstractness. This is because it is often very difficult to individuate properties. For example, are 'being-white', 'not-being-green' and 'not-being-red' all inessential properties of the chalk triangle? Color may be relatively easy to distinguish as a type that can be instantiated with a variety of tokens, but this is not always possible, especially when we are dealing with systems of great complexity. Nonetheless, Cartwright, following Aristotle, believes that focusing on the essential properties of objects can help with at least a partial ordering of abstract objects: "A is a more abstract object than B if the essential properties, those in the description of A, are a proper subset of the essential properties of B" (ibid 214). For example, a right triangle is more concrete than a triangle because 'having-a-right-angle' is an essential property of a right triangle which is abstracted in the case of a triangle. On the other hand, all the properties of triangles also apply to right triangles (ibid).

This ability to partially order abstract objects based on their essential properties is very important for the distinction between abstraction and idealization. Essential properties found in abstract objects are found in the real world, i.e. they are part of the world. A concrete object has them and retains them after the process of abstraction is complete. This is the main difference between abstract and idealized objects or systems: all the properties of an abstract system exist in the real world whereas at least some of the properties of an idealized system have no counterpart in the real world. This means that abstraction should not be viewed as involving distortions. A description of an abstract system will contain only truths. It may not contain all the truths, but it will not contain any falsehoods.

## 2.2 The Omission view of Abstraction

The main proponent of the omission view of abstraction is Martin Thomson-Jones. His view is similar to Cartwright's in that he believes that while idealization is the result of misrepresenting features of a system, abstraction results from merely omitting features (Jones 2005). However, his account of abstraction and idealization is not identical to Cartwright's, as he focuses explicitly on the process of omission as the hallmark of abstraction. Here, abstraction and idealization are distinct because idealization requires the assertion of a falsehood, while abstraction involves the omission of a truth (ibid). Thus, while both idealization and abstraction can result in the distortion of a system, the distortion is very different in each case. When we abstract, we do not describe the system in its entirety, so we are not telling the whole truth. However, when we idealize, we add properties to the system that it does not normally possess. Therefore, our description of an idealized system contains falsehoods. Thomson-Jones also highlights some issues that are of interest especially in the case of abstraction modeling. These points are intended to show precise ways in which abstraction and idealization differ.

First, Thomson-Jones considers a criticism that has been leveled against his (and Cartwright's) account. The distinction between abstraction and idealization construed as omission versus misrepresentation is problematic because in some cases omission of factors results in distortion or misrepresentation of the system. For example, we can model a food web of an aquatic ecosystem without including primary producers, yet this would result in the misrepresentation of the energy flows through the system. There are countless other omissions of this sort that can result in misrepresentation. However, Thomson-Jones attempts to avoid this problem by restricting abstraction to precisely those omissions that do not result in misrepresentation (ibid). A 'mere omission' does not misrepresent a particular feature of a system because it retains 'complete silence' with respect to whether the system contains the feature (ibid). So if an omission results in a misrepresentation, then it is not the type of omission that is part of abstraction.

It is important to note that this setup presupposes that we have enough knowledge of the system in question to know which features we are misrepresenting or omitting. There may be times where we omit a feature of a model thinking that it is irrelevant, only to find out later that it resulted in a misrepresentation of the system. In cases like these the problem is our epistemic limitation. Thomson-Jones's account is

construed as a general theoretical framework and applies to cases where such limitations have been overcome.

An interesting consequence that arises from this first point, is that idealization and abstraction are mutually exclusive. That is, a particular feature of a model can be either abstract or idealized, but not both. If the primary producers from the previous example are absent from the model of the food web, then they cannot be misrepresented as being in another trophic level. Conversely, if their population is represented as infinitely large (a common misrepresentation in biology) then they cannot be entirely absent from the system.

Second, Thomson-Jones shows that successful idealizations are thought to *approximate the truth*, whereas this notion is meaningless for abstractions. Idealizations are supposed to be useful fictions so they need to be distinguishable from non useful fictions (mistakes). The condition as Thomson-Jones presents it, is quite weak. Either the idealization itself can approximate the truth, or it can be part of a model which ‘captures the approximate truth overall’ or it can be part of a model that gets the basic ontology of the modeled system right (ibid 186). Moreover, he stresses that this is supposed to constitute neither a necessary nor a sufficient condition for a misrepresentation to count as an idealization. It is simply a feature of idealizations in general, that they depart less from the truth than outright mistakes.

The third point refers to *simplicity*. Idealizations are often introduced in a model, to make the model simpler (ibid 187). This is because a simple model is more computationally tractable. Thus, a successful idealization will simplify a model. The case of abstraction is slightly different. It makes no sense to say that a successful abstraction simplifies a system, because all abstractions simplify a system by virtue of their being omissions (ibid 188). Therefore, we can use simplicity as a criterion for identifying successful idealizations, but it is not useful for distinguishing between abstractions.

The fourth point of divergence between the omission view and the material view concerns *relevance*. Idealizations are thought to misrepresent relevant features of a system. They are relevant in terms of predicting and explaining the particular behavior of the system for which the model was constructed (ibid 187). The idea is that accidental or irrelevant features of the system will not need to be misrepresented, thus idealizations will focus on features that are important to the functioning of the system. At first glance it seems that abstractions by definition omit features that are irrelevant, which would give us a neat way to distinguish between idealizations and abstractions. However, Thomson-Jones thinks that the

situation is more complex. Though irrelevant features of models are omitted as a result of abstraction, some relevant features can also be omitted. This can happen when we wish to screen off a feature, even when we believe it is relevant. For example, we may know that nitrogen uptake is necessary for the growth of a plant, yet we might want to isolate other important factors of plant growth. We can then construct a model that abstracts away from nitrogen uptake and focuses on differences in other causal factors.

### 2.3 The Distortion view of Abstraction

There are two versions of the distortion view of abstraction, a stronger and a weaker version. The first version is articulated explicitly by Paul Humphreys in his review of Cartwright's chapter 'Abstract and Concrete' (Humphreys 1995), though it also implied in many accounts of idealization (McMullin 1985). The view is presented in the form of a criticism of Cartwright's view, which states that abstraction and idealization are not kept apart as easily as Cartwright asserts. In fact, Humphreys goes so far as to assert -in direct contrast to Cartwright- that abstractions are useless unless idealization is already possible (Humphreys 1995 159).

Humphreys takes issue with the idea that all the properties of an abstract object or system exist in the real world. Even though he agrees with Cartwright that 'symbolic' or idealized representations contain falsifications and that 'they give us concepts that do not describe reality', he thinks that this also applies to the case of abstraction (ibid 158). He points out an important passage that Cartwright uses to illustrate the difference between abstraction and idealization. The example, outlined in the previous section refers to a triangle: 'Consider a triangle, a real triangle, drawn on a blackboard. Even when the chalk and the colour and all the other accidental features are subtracted, the shape that is left is never a real triangle. But let us pretend that it is ...' (Cartwright in Humphreys 158). As I showed in chapter 2, for Humphreys, as soon as we start pretending that an abstraction is different than it seems, we have blurred the line between abstraction and idealization.

Cartwright wants to pretend that the abstracted triangle is real because that means that all the properties it has are real and are still there after the process of abstraction. If this were an idealization, the triangle would involve distortions and would have properties that were not there in the world before the process of idealization started. The abstract triangle is not entirely real, but it is not unrealistic in the sense of idealized representations. The problem is that we start off with a shape that is not a true triangle because



we have drawn it with chalk and it is not perfectly triangular. When we abstract away its accidental properties, the triangle that remains is therefore also not perfectly triangular. According to Cartwright, this is a different sort of issue than the unrealistic nature of idealizations. In the case of abstraction, we start off with all the distortions and imperfections that end up in the abstract system. No additional misrepresentations are introduced in the process of abstraction. The distortions that are present, are themselves part of the world and arise because of our limited capabilities as humans, i.e. we cannot draw perfect triangles. Still, these triangles are close enough to reality that the imperfections do not distort the geometrical properties of triangles when we abstract. Therefore, Cartwright thinks that it is legitimate to pretend that the abstracted triangle is, for all intents and purposes, real.

Humphreys thinks that this move undermines Cartwright's position. He grants to Cartwright that idealizations may not always be useful, because they fail to give us realistic representations. Even so, he thinks that making these sorts of pretenses adds an element of idealization to abstractions. As mentioned above, Cartwright distinguishes between cases such as representing the sun as a sphere and conceptualizing triangles in the abstract. The former is considered to be an idealization, because we substitute properties of the sun with the geometrical properties of spheres, whereas the latter is an abstraction, because we simply omit irrelevant properties of the particular triangle and focus only on its geometrical properties. Yet Humphreys points out that we cannot legitimately focus on the triangle's geometrical properties because an imperfect concrete triangle will remain imperfect after we abstract. If we have to pretend that the concrete triangle has geometrical properties, then we will also have to pretend that the abstract triangle has them. This is because in the case of true abstraction all the properties of the abstract object exist in the real world. If we are forced to pretend, we must continue pretending throughout the process of abstraction, and consequently there seems to be no substantial difference between the 'imperfect real triangle' and the 'imperfect real sphere' (ibid).

Humphreys takes his criticism a step further. He asserts that whenever a system is represented mathematically, the process of abstraction will contain elements of idealization. Though Humphreys does not elaborate, the problem here is that mathematical representations of real systems have properties that the systems themselves do not. Consider again the logistic growth model described in the introduction.<sup>5</sup> It measures how the growth rate of a population  $N$  is limited by the density of the population itself. An

---

<sup>5</sup> For a full description and use of the logistic growth model see chapter 2, pp. (x-x).

interesting aspect of this equation is that it is differential. That is, it is a set of functions whose derivative can be found for every point in their domain. The derivative, in turn, is a secondary property which shows how the dependent output of the equation changes with respect to its input. In the case of logistic growth, the rate of change the population  $N$  with respect to time  $t$ , is given by a linear equation. The problem is that it does not make sense to say that real growing populations have the same properties. For example, the marmots of Vancouver Island have a number of biological properties, but differentiability is not one of them. Only the mathematical representations of the populations have these special mathematical qualities.

If this is the case, then every time we represent a system mathematically, we introduce properties in the representation that are not there in the original system. Yet the introduction of new properties is a mark of idealization, not abstraction. This does not mean that abstraction is impossible. We can still omit properties of the system in our representation. However, any such abstractions can only go hand in hand with idealization. Without the mathematical representation and the idealization it introduces, we would not be able to model the system in the first place. Therefore, it is only because mathematical representation makes idealization possible that abstraction can be useful (ibid 159).

The weaker version of the distortion view can be seen as a middle ground between the material view and the strong version of the distortion view. Proponents of this view think that abstraction is distinct from some kinds of idealization, but not from all kinds of idealization. They think that idealization itself is not a unified concept, and abstraction falls under one of these kinds of idealization. There are two accounts of the weak version of the distortion view, Frigg & Hartmann (2009) and Weisberg (2007). According to Frigg and Hartmann, idealization in general terms is 'a deliberate simplification of something complicated with the objective of making it more tractable' (Frigg & Hartmann, 2009). Further, they think that it can be divided into two basic kinds, Galilean and Aristotelian. Galilean idealization was first described in detail by Ernan McMullin, who outlined Galileo's use of distortions in his explanations, the problems they caused and potential solutions (McMullin, 1985). The important aspect of Galilean idealization is that the distortions are supposed to be temporary, at least in principle. We are forced to distort our representation of the system for greater ease of explanation or computational tractability, yet our aim is to eventually de-idealize our model as we learn more about the phenomenon we are studying. Aristotelian idealization, on the other hand, involves stripping away irrelevant properties from a representation of a system (ibid).

Moreover, the two kinds of idealization are thought to coexist in models. For example, a mechanical model of the planetary system only takes into account some properties, omitting a number of others and hence engaging in Aristotelian idealization while distorting the remaining properties engaging thus in Galilean idealization (ibid). The result is a model that describes planets as ideal spheres with a ‘rotation-symmetric mass distribution’ (ibid). This distinction mirrors the one presented in Cartwright and Thomson-Jones between abstraction and idealization. The difference here, is that abstraction is thought to be a kind of idealization.

A more nuanced version of this view is presented by Weisberg in his paper “Three kinds of Idealization” (Weisberg, 2007a). As the title suggests he identifies a third type of idealization, in addition to the two identified by Frigg and Hartmann. Galilean idealization features in Weisberg’s account, as the first kind of idealization. As in the case of Frigg and Hartmann, it is presented following McMullin’s account. That is, one starts with an idea of what a non-idealized representation of a system would look like, and then introduces distortions that help simplify the representation and make it computationally tractable. The model that is created includes these distortions (Weisberg 2007). An important aspect of this kind of idealization is that it allows for the systematic removal of distortions, as the system is better understood.

Another type of idealization is called ‘Multiple Models Idealization’. This occurs when theorists create a number of related but incompatible models that describe a particular (usually extremely complex) phenomenon (ibid). Each model will have a different set of idealizations and simplifications, and will highlight different aspects of the phenomenon. Taken together, the collection of models is meant to give a better understanding of the phenomenon. A good example Multiple Models Idealization is found in climate science. The United States Weather Service use three different models of global circulation patterns, each with different assumptions, to predict weather patterns (ibid).

The third kind of idealization is ‘Minimalist idealization’. This is ‘the practice of constructing and studying theoretical models that include only the core causal factors which give rise to a phenomenon’ (ibid). A representation of a system that has omitted all factors except the core causal factors is called a ‘minimalist model’ (ibid). Cartwright’s account of abstraction is seen as an example of this kind of idealization. This is because abstraction results in isolating the relevant causal factors which affect the behavior of the system and omitting the remaining irrelevant factors. The model based on abstraction will therefore be a minimalist model.

On Weisberg's account, there is an important difference between Galilean idealization on the one hand and Minimalist and Multiple Models idealization on the other. Only Galilean idealization has de-idealization built in to the process. That is, this kind of idealization is justified by the possibility of de-idealizing, when more knowledge about the phenomenon is available. On the other hand, the justification for Minimalist idealization is that the omission of irrelevant factors allows theorists to find the underlying mechanisms that give rise to the phenomenon. The idea is that the omission of irrelevant details reveals underlying truths about real phenomena, so that there is no need for de-idealization in the future. The minimalist model is supposed to give precise and useful representations of phenomena. Similarly, there is no de-idealization built in to the Multiple Models type of idealization. Here, each of the models has different idealizations built in, and they are supposed to complement each other, or cancel each other out. Taking the models collectively counteracts any problems that arise from a particular assumption or idealization.

As in the case of Frigg and Hartmann, this view can be seen as a middle ground between two extremes. On the one hand, abstraction is a process in its own right, which has different motivations and justifications from other kinds of idealization. It can coexist with Galilean idealization, but it is far from being the same process. The various kinds of idealization are separable and each has its place in science. Whenever minimalist models are useful, abstraction is the key. On the other hand, Weisberg views the process of abstraction as involving distortions, as omitting factors in a representation distorts the phenomenon which is represented. In this sense, abstraction it is no different to other kinds of idealization.

Even though these two views are less extreme than Humphreys's in their dismissal of abstraction as a process distinct from idealization, they nonetheless present a challenge which the material view and omission views cannot overcome. As long as abstraction distorts a system in any way, then cannot be seen as an altogether different type of process. If this is the case, then it seems that proponents of the material and omission views do not have the requisite support for their case for abstraction. For my purposes, this criticism is even more devastating. On my view of target systems, they are abstract but concrete and real entities. If the process of abstraction necessarily introduces distortions, then it is not possible to claim that target systems are at once abstract and real.

In summary, here are two distinct criticisms which arise from the distortion view of abstraction. The first, associated with Humphreys, is that abstraction and idealization are indistinct. The second,

associated with Frigg & Hartmann and Weisberg, is that abstraction and idealization are not identical, yet both are distorting processes. As they stand, the material and omission views of abstraction cannot overcome these criticisms. However, I will argue that a revised version of the omission view can overcome them.

### 3. The Revised Omission view of Abstraction

All the accounts of abstraction presented so far have an important point in common. Even though they may give the process different names or disagree on some implications or details, they all define abstraction as the representation of a real-world system or phenomenon with all the irrelevant factors or parts omitted. The process of omission is the key to defining and understanding abstraction. I propose a revised view of abstraction which is restricted to omission. In order to present the revised omission view, I will first introduce a distinction between omission, de-concretization and generalization. I will argue that pulling these three aspects apart is important for overcoming the criticism of the material view of abstraction. I will then analyze the main criticism of the material view of abstraction, stemming from the weak version of the distortion view, which is mainly aimed at the standard omission view.

The main criticism of the material view is the one outlined in chapter 2 concerning the triangle. In Cartwright's example, a physical triangle drawn on a board becomes abstract when we remove, in our minds, all the properties which tie it to the world. In the end we are left with an abstract and not physical idea of a triangle. Continuing on the theme of triangles, the west pediment of the Parthenon is a physical object which exists in space and time. When we talk of the west pediment in isolation from the rest of the Parthenon, we are abstracting. That is we start off with the entire Parthenon and focus only on some aspect of it, the west pediment. On Cartwright's view, we can also abstract further, omitting the colors, the frieze etc., and we are left with an abstract idea of a triangle.

The problem is that, as Cartwright herself notes, we could never arrive at a true triangle from omitting properties of a real physical triangle. The reason is that no physical triangle is a true triangle, but a mere approximation of a true triangle. Therefore, in order to get a true triangle, with all the important geometrical properties, we must 'smooth it out' or, to use Cartwright's phrase, we must pretend that it is a real triangle. According to Humphreys, as soon as we pretend or 'smooth out' we are misrepresenting and idealizing. A closer look at this last step, reveals that two processes are happening, omission and de-

concretization. The omission of the properties itself, does not make the triangle immaterial. This is precisely why we have to add the ‘smoothing out’ step. However this also means that simple omission of properties does not suffice for getting an immaterial, true triangle from an actual object. The final step, or steps, involves de-concretizing and idealizing from the physical triangular object to the immaterial true triangle.

We must concede Humphreys’s criticism. He is right in saying that the last step is an idealizing step. However, as argued in chapter 2, it is also possible to bite the bullet and restrict abstraction to omission without de-concretization. In the example, when we abstract we are left with the actual physical west pediment, without taking into consideration any other properties other than its triangular shape. Importantly, this shape merely approximates a true triangle and is itself not a true triangle.

De-concretization is a central aspect of the material view of abstraction, as an object is abstract precisely when it is de-concretized. However, it is also responsible for blurring the distinction between abstraction and idealization, which is the main point that proponents of the material view set out to prove. Therefore the first important revision for a successful omission view of abstraction, is to distinguish it from de-concretization. De-concretization is often tied together with omission in views of abstraction, because abstraction is thought to be a mental process which involves representing the system with parts and properties omitted. For example, for Cartwright abstraction, like idealization, happens “in our minds” (Cartwright, 1994). Thomson-Jones also implies that abstraction involves representing a system, even though he does not explicitly tie it in with de-concretization (Jones 2005).

While I agree that we sometimes construct a mental picture of the process of abstraction, this does not mean that an abstract system is itself a representation. That is, the fact that we sometimes find it helpful to construct a mental representation of the abstracting process does not mean that the process itself is happening in our minds. In addition we might represent an abstract system by referring to it with language or a model. Yet, this does not mean that the system that we are referring to is itself a representation. For example, when we talk about the west pediment of the Parthenon in isolation from the rest of the Parthenon, we are referring to an abstract system and hence representing it. Yet the system itself is not a representation, it is the actual physical system. In other words, when we refer to the west pediment of the Parthenon, we are not representing a *representation* of the west pediment of the Parthenon; we are representing the west pediment *itself*.

The same is not true of idealization. The mental representation which is connected to idealization results in models, paintings, poems etc., systems which are themselves representations of actual systems. In fact, in the case of modeling, the important representation is not the internal representation in each scientist's mind, but the external mathematical, computational or concrete representation of the system. The internal representation is relevant only in the sense that it provides the original intention of the scientist, for the model to count as a representation of the system. This is important for getting the possibility of representation off the ground, yet once that happens, the nature of the representation, its analysis and evaluation does not deal with the internal representation but the external one. If the internal representation has a limited role to play in the process of idealization, I do not see why it should have a more extensive role to play in the process of abstraction.

Another concept which is often tied together with abstraction is generalization, which also helps to blur the lines between abstraction and idealization. Abstraction is thought to be a generalizing process in the material view because an abstracted system is thought to be a general idea which refers to many particularizations of that idea. For example, the notion of a triangle is a general one, which can refer to many particular physical objects that are triangular, such as the west pediment of the Parthenon and the chalk triangle on Cartwright's blackboard. Yet general notions or ideas unlike particular instantiations of them are not grounded in the physical world. Hence, if we assume that omission and generalization are one and the same, then abstract systems cannot not be concrete and real.

There is a relationship between omission and generality, in the sense that in order to get from a concrete particular to a general notion we must first omit the properties which make the particular unique and use the remaining properties to construct the general idea. Hence omission is a prerequisite of generalization when going from the concrete to the general. However, this does not mean that the process of omission implies generality. It is possible to omit properties and not proceed to the next step of generalization. For example, when we talk about the west pediment in isolation from the rest of the Parthenon, we are abstracting away the metopes, the friezes, the doric columns etc. Yet we would not say that the west pediment is general. Even if we abstracted more and focused only on the triangular shape of the pediment without the statues of Athena, Poseidon and co., it would still be a concrete particular, an imperfect instantiation of the notion of a triangle.

To sum up the argument so far, I have analyzed the first criticism which comes from the strong distortion view and is mainly directed against the material view of abstraction. I have shown that the criticism can be overcome by separating the notion of omission from the notions of de-concretization and generality. The revised view of omission explicitly restricts abstraction to omission. I will now turn to the second criticism, which comes from the weak distortion view and is directed mainly against the standard omission view of abstraction.

The crux of the weak distortion view of abstraction is that omission constitutes a distortion of a system, even though this distortion might be different in kind to the distortion which results from idealization. In other words, all omission is by definition a distortion of the system. Thomson-Jones attempts to avoid this problem by restricting abstraction to precisely those omissions that do not result in misrepresentation (Jones 2005). As stated above, a 'mere omission' does not misrepresent a particular feature of a system because it retains 'complete silence' with respect to whether the system contains the feature. So if an omission results in a misrepresentation, then it is not the type of omission that is part of abstraction. To emphasize the distinction further, Thomson-Jones states that a property that has been omitted from a system cannot be misrepresented, hence with respect to a particular property, abstraction and idealization are mutually exclusive.

The problem is that the criticism presented here is much stronger, as it denies the possibility of 'mere omission' altogether. The standard omission view cannot overcome this criticism, because it assumes that abstraction and idealization both occur at the level of the model which is a representation of the system. As stated above, if we allow abstraction to be part of the representational process, then we cannot distinguish it from idealization. This is because abstraction is distinct from idealization because abstracted systems are not representations of systems, hence abstract systems are concrete physical and real systems. Real systems are not distorted, they simply exist. By restricting abstraction to the specification of the target system and idealization to the model which represents the target, we can have omissions which are not misrepresentations.

Another issue which is important is the question of relevance. The standard view of omission allows for relevant factors to be omitted from a system. Thomson-Jones states that even though many omissions are of irrelevant factors, it is still possible to omit relevant factors in two cases: (a) when we want to screen off factors that are relevant and (b) when we mistakenly omit a factor even though it is



relevant. The problem with this statement is that it leads to inconsistency within the view. Thomson-Jones argues that ‘mere omissions’ remain silent on whether the system has the property or does not. Yet, remaining silent means omitting the property from the model. But why would we omit a property from a model if it really is relevant? This would result in a model with distortions that are not the intentional kinds of distortions which occur from idealization.<sup>6</sup>

I think that the underlying cause of the problem is the particular use of the term ‘relevant’. Thomson-Jones is stating that factors which are *generally* relevant to the functioning of a system can be omitted. I previously illustrated this point with an example from ecology, where a scientist can screen off nitrogen from a model of plant growth, even though she knows that plants cannot survive without nitrogen uptake. I agree that it is possible to screen off factors that are generally relevant, yet only when they are not relevant for the particular problem we are trying to solve. Nitrogen can be omitted from a study if it is not a limiting factor, and occurs in sufficient abundance so as not to have an effect on the plants in the study. For example, the scientist in question might be trying to determine how different plants respond to predation from a particular predator, and the amount of nitrogen will not affect regeneration. Yet in doing so, the scientist is implicitly asserting that nitrogen is *irrelevant* for the particular experiment, even though it is generally relevant for plant growth.

The point is that the relevance of a particular factor is determined (at least in part) by the intentions of the scientist, given the specifics of a phenomenon and the method used to study it. Whether a factor is generally relevant or of limited importance, the point is to determine if it is relevant *for* a particular model or experiment. Therefore, for the omission view of abstraction to be consistent, it must be the case that only irrelevant factors are omitted.

To sum up, the revised omission view of abstraction departs from the material view by distinguishing omission from generalization and de-concretization. It also departs from the standard omission view by isolating abstraction from representation, and maintaining that only irrelevant factors can be omitted in an abstracted system. In doing so, the revised omission view provides a coherent way to distinguish between the process of abstraction and idealization and their respective products. It also results in a view which is internally consistent. Finally it creates the conceptual space for understanding target

---

<sup>6</sup> It is important to note that I am not claiming that the omission of relevant factors does distort the system itself, because it is not a representation of the system (see previous paragraph). However, I am claiming the resulting model will contain additional distortions which are not intended.

systems as parts of the world, which are both abstract and real. In the next section I will consider three cases which can be seen as counterarguments to the revised omission view of abstraction, and show that the view has the necessary tools to overcome them.

#### 4. Difficult Cases

My exposition of the revised omission view of abstraction dealt with the general theoretical framework concerning abstraction and aimed to provide a conceptual argument for distinguishing between the notions of abstraction and idealization. However, there are three cases that appear in philosophy of science literature, which could be seen to undermine the revised omission view. The first is the case of errors in abstraction brought up by Thomson-Jones, which occurs when features are mistakenly thought to be irrelevant. The second is an example of a case where an idealization appears to be an omission, and thus threatens to weaken the distinction between abstraction and idealization. The third is the notion of “generalized targets”, which threatens the distinction between abstraction and generalization.

##### 4.1 Mistakes

According to Thomson-Jones, the second reason why we should allow for the possibility of omitting relevant factors from a system, is in order to allow for mistakes. Scientists often do not realize that a particular factor is important for the explanation of a phenomenon and re-design their models and experiments to accommodate these realizations. For example, in the case of plant ecology, there was a recent realization that abiotic factors and the soil microbes through which plants absorb them are a lot more important than previously thought (Klironomos 2002). In fact, these microbes can even be used to predict and explain some plant invasions. It is thought that while native species accumulate pathogens which cause negative feedback reactions in the soil, species that are introduced into a new habitat are free from their native underground enemies and exhibit neutral or positive feedback with the same microbial species (ibid). In some cases, an invader might have both positive and negative feedback going on at the same time, so that the the effects of feedback are cancelled out. This means that even though they are important factors in plant competition, especially between native and invasive species, they are often overlooked.

Thomson-Jones wants to allow for the possibility of omitting relevant factors, in order to make sense of cases like these in scientific practice. However, the distinction to be made here is between the possibility of making inapt omissions and the goal of scientists while abstracting. Cases like these show that one cannot deny the possibility of making inapt omissions, yet one must also recognize this as a mistake. When scientists realize that their omissions are inapt, given the phenomenon being studied, they re-specify their targets in order to include all relevant information. Hence it is still possible to say that the goal of abstraction is to omit only irrelevant factors, yet this is often difficult to achieve.

The question is whether these cases of inapt omissions count as distortions of a system. However, as I argued in the previous chapter, they do not. The problem is an epistemic one, not an ontological one. Omitting aspects from a system does not change the ontological status of the system, because omission is, like Thomson-Jones states, remaining silent on whether the system has the property or not. This implies that a scientist thinks that the omitted factor is not relevant, so it does not have to be included in the model. It does not mean that the scientist physically removes the factor from the system or distorts the system in any way.

The fact that this is an epistemic rather than an ontological problem does not imply that it is not important. In fact, I think that part of the reason why the study of target systems is important, is because it can help us understand when factors are relevant and when they are not. The way to deal with this problem is to have a method for evaluating target systems, which is the subject of the next chapter. There is no need to think of the problem as an ontological one in order to make it worth noting.

## 4.2 Friction

There are many models in science which contain idealizations which, at first glance, seem to be omissions. A standard example is a model of harmonic oscillation without friction. If we say that the model represents a target system, then one way to understand the difference between the model and an actual pendulum is that the friction has been omitted. If this is the case, then the revised omission view of abstraction faces a problem. Omitting friction by all accounts counts as a distortion of the system, as a target system without friction cannot possibly be real.

However, I think that this problem rests on a misunderstanding. The harmonic oscillation model potentially represents many target systems. Let us assume that in this particular case it represents Ron

Giere's grandfather clock. What is the target system of the model in this case? It is not the entire clock, but only the pendulum mechanism. All the other parts of the system (the clock) have been omitted. Aspects of the pendulum, such as its color and chemical composition are also omitted. What about friction? Friction could be omitted, in principle, if the scientists thought it to be irrelevant. This would result in a very inapt target system, like the case mentioned above. However, omitting friction as irrelevant does not mean that the scientist is asserting that there is a system in the world without friction. She is simply asserting that friction is not important. In other words she is saying that the friction of the system could have many different values, yet this is simply not important given the goal of the model. Whether this turns out to be true or not is an epistemic issue, like the one mentioned above.

In fact, making any sort of change to the amount of friction in the system is an idealization, and happens at the level of the model. In other words we can represent the target system with the full amount of friction, a diminished amount of friction or zero friction in the model. Yet this does not alter the state of the actual pendulum in the particular clock, which always continues to have friction. To illustrate this point more clearly, I will consider another idealization which is used very often in biology. Many biological models misrepresent populations of organisms as infinite, as it simplifies the mathematics involved in the model's computations. We would not say that this counts as an omission, as all real populations are less than infinite. Yet the change that occurs in these models is the same kind of change that occurs in the harmonic oscillator model. The world is a certain way: a particular population has 2 or 85 or 1003 individuals, just as particular pendula are subjected to certain amounts of friction. Increasing or decreasing these values even in the extreme, happens at the level of the model and are misrepresentations of the world. They are distortions, but distortions resulting from idealization.

#### 4.3 Generalized Targets

Scientists construct generalized models when they want to learn about classes of phenomena in general, instead of particular instances of phenomena. These models are often very simple, in the sense that they aim to capture only the basic dynamics of a system and include very little detail. Many models in population ecology fall into the category of general models, as they describe general phenomena such as population growth, competition and predation. Other examples can be found in population genetics. An important account of generalized models is found in Michael Weisberg's book *Simulation and Similarity*

(Weisberg 2013). Weisberg uses a model of sexual reproduction as an example of a generalized model. According to Weisberg, a “generalized model of sexual reproduction isn’t supposed to be about kangaroo sex or fungi sex, but about sex itself” (Weisberg 2013 116). The point of studying sex in this general manner is not to discover facts about sexual reproduction in particular populations, but to understand larger issues such as the relative merits of sexual reproduction when compared to asexual reproduction. Weisberg points out that the emergence of sex is a strange phenomenon and there is no real consensus on why it evolved. Therefore, our best bet for determining why sexual reproduction confers some evolutionary advantage, is constructing and analyzing a generalized model of sexual reproduction (Roughgarden 1997, in Weisberg 2013).

According to Weisberg, generalized models are used to understand *abstract* targets. The term ‘abstract’ here means that the targets are themselves general and do not have details specific to particular instantiations of the phenomenon. This is important for Weisberg’s account, as otherwise it would be difficult to connect a general model to a non-general target. As Weisberg points out, there is no such thing as ‘sex in general’ in the world, but only a collection of instances of sex in particular populations “kangaroo sex, Tasmanian devil sex, human sex, but not sex in general” (Weisberg 2013). He continues: “sex in general is an abstraction over these more specific kinds of sex” (ibid).

But this raises an important question, namely how does an abstract and general model represent particular instantiations of a phenomenon? Weisberg gets around this problem by stating that the target of a general model is also general. In other words, according to Weisberg, a generalized model of sex will have a *generalized target*. This is a target of generalized phenomena, not particular instances of them and is “by its very nature, more abstract than any specific target” (ibid).

The problem, for my view, is that generalized targets are not real. The number of general models in science is substantial, so if their targets really are general targets, then the view that targets are real parts of the world can at best be limited to a small subset of targets. However, a closer look at the generation of generalized targets reveals that they rely on local (and hence particular and real) targets.

On Weisberg’s view, there are two ways in which generalized targets can be generated. The first occurs when the generalized target lies at the intersection of a number of specific targets. That is, the generalized target has the features common to all targets that the model can represent (Weisberg, 2013). In the case of sexual reproduction, the generalized target has the set of properties which are shared by all

sexually reproducing species. Once the relevant features shared by all specific targets are identified, the scientist can abstract further by leaving out all the specifics of each case and focusing only on those generalized properties. The generalized target is therefore more abstract than the set of specific targets, and more importantly, it is exactly as abstract as the model. This is the simple case of generalized modeling, which works as long as: “1. The relevant set of specific targets actually share the relevant features, such that an intersection of their sets of features is an informative generalized target, and 2. A model can be constructed at the appropriate level of abstraction so that just those features can be modeled” (Weisberg, 2013, 117).

It is obvious here that these generalized targets rely on local targets for their generation. The generalized target is nothing more than a generalization over a number of real local targets. But this means that there is no reason to claim that *the* target of a general model is a generalized target. It is one thing to state that a general model gives us information about a phenomenon in general, and quite another to state that it *represents* that phenomenon in general. Instead, what the general model represents, are *many local* targets.<sup>7</sup> Therefore, even though the general model of sexual reproduction might tell us about sex ‘in general’ it *is about* kangaroo sex *and* Tasmanian devil sex *and* human sex, and so on.

Yet, for Weisberg, there are also more complex cases where the second criterion is not satisfied, as many models are “less abstract” than their generalized target. Models that are less abstract than their targets contain feature types that go beyond what is shared between the specific targets. Examples of these types of models are individual-based computational models. If a scientist were to construct an individual-based model of sexual reproduction, she would not represent the distribution of genotypic fitness in a population (a very abstract property) but specific genotypes tied to specific organisms which are represented as agents. In order to do this, the scientist would have to include concrete properties such as the life-cycle, spatial distributions and fitness of individual organisms.

Weisberg correctly points out that in cases such as these, the model does not obviously connect to the generalized, abstract target. Instead, the scientist must place special constraints on the construal of the model and its intended scope. Even if the model is capable of generating concrete detail, the scientist must restrict its scope to the more general results which are compatible with the generalized target. Thus, if a

---

<sup>7</sup> The actual number of targets that the model represents is determined by the world itself, that is the model can potentially represent all populations throughout space and time which reproduce sexually. The model actually represents only those populations in space and time that scientists have applied it to.

model has features that cannot be part of the generalized, abstracted target, such as being able to generate specific predictions about a particular population, the theorist “must regard these features as remaining outside the model’s scope” (Weisberg, 2013, 117).

This complex case is problematic for Weisberg’s account, but it is not really a problem for my view of target systems. This is because the problem is in fact *generated* by the introduction of a generalized target. The only reason the mismatch occurs, is because the model is being compared to a generalized target in the first place. An individual-based computation model of sexual reproduction need not be understood as a model of sex in general, but as a model of kangaroo sex and tasmanian devil sex and human sex etc., just like the population genetic model. The models might have different types and levels of idealization and abstraction but they represent the same type of target, local targets.<sup>8</sup>

I understand Weisberg’s point that it is useful to have a term to refer to a collection of targets which a model can represent. Yet there is no need for the collection to take on a different ontological status from other targets. There is a much simpler way to have a name for a collection of targets without undermining the coherence of the account of targets as real systems in the world. The solution is what I call a ‘family of targets’, i.e. a group of local targets which shares some common features. Applying this to Weisberg’s account, it is the group of targets which shares common features before the extra step of abstracting to a generalized target. This group of similar targets has important properties, namely that it focuses the scientists’ attention to the features that are shared by a number of systems in the world. The general model is general *because* it applies to these local targets which share a number of properties. Moreover, identifying another target to which the general model might apply is straightforward, at least in principle. All that the scientists need to do is determine whether the particular instantiation of a phenomenon shares the relevant features and can be incorporated into the family of models.

## 5. Conclusion

Abstraction is an important process in science, especially in modeling. However, its role has been overlooked by many and those who have defended it have been criticized. Some of these criticisms are truly valuable, because they have shown that existing accounts of abstraction need to be revised and

---

<sup>8</sup> It is important to note that the targets will not be identical, as the two types of models will have different abstractions because different properties will be relevant in each case. Nonetheless, the important point is that the targets will all be local and quite similar to each other.

refined. The key to understanding the notion of abstraction is omission, and this forms the core of what it means to abstract in science. The important and difficult part of the process is therefore the decision of which properties are relevant and which are not. The other notions which are often associated with abstraction, de-concretization and generalization can be divorced from the notion of abstraction. Doing so will help to overcome the charge that abstraction is nothing more than a kind of idealization. For the purpose of target systems, abstracting by omitting irrelevant parts and properties is an extremely important and difficult endeavor. In the next chapter, I will give an account of how this endeavor can be constrained and evaluated.



## 5. Target System Evaluation

### A Theory of Aptness

#### 1. Introduction

The North American lakes have been invaded by an exotic plant. Like many other ecosystems in the world, an alien species has become firmly established and is spreading through the ecosystem at the expense of the native flora. In this case, the culprit is *Typha angustifolia*, a species of aquatic cattail, native to Europe. In order to explain this particular phenomenon, scientists must determine *why* it is the case that *T. angustifolia* has succeeded in invading the Indian Creek ecosystem in the North American lakes.

There are a number of ways in which scientists can go about studying this invasion. For example, they can focus on the invader (*T. angustifolia*) and look for traits that help it outcompete native species. Alternatively, they can focus on the community of Indian Creek and determine whether there is something about these communities which makes them susceptible to invasion. These two alternatives are quite different, as they involve distinct conceptual schemes for understanding invasions. In other words, each alternative affects the way scientists conceptualize the phenomenon of invasion, and the methodology for studying invasions.

Translating this into the language of target systems, the scientists must choose one of the available target systems in order to model or experiment this particular invasion. Given this choice, there are two sets of questions we can ask. The first, is why the scientists chose this particular target system, and by extension what governs target system choice, in general. The second is how we can know that the scientists made the right choice of target system, and, more generally, how we can evaluate target systems. A brief examination of the first set of questions will pave the road to answering the second set of questions.

In this case, the choice pertains to partitioning and abstraction, as the domain of study (Indian Creek) and the phenomenon of interest (invasion by *T. angustifolia*) are already set. Yet there are at many partitions to choose from, and numerous ways to abstract in each partition. Instead of specifying the various

target systems that the scientists could choose from, I will focus on three, which differ in important and illuminating ways. Target System 1 (TS1) is partitioned at the level of communities and its properties are community-level properties such as biodiversity and disturbance. Target System 2 (TS2) is partitioned at the level of individuals, thus organisms (of natives and invaders) are its main units. It is also highly abstract, as it does not include any other parts of the ecosystem. In addition, it includes only the traditional properties of plants, associated with competition (such as growth rate, seed size, photosynthetic rate, root/shoot ratio biomass). Target System 3 (TS3) has the same partition as TS2, yet it differs in abstraction, as it contains an additional property of *T. angustifolia*, *allelopathy*, which is not contained in TS2. Allelopathy is the exudation of toxins from the root of a plant which inhibit the growth of other plants.

In their study, published as “Allelopathy as a mechanism for the invasion of *Typha angustifolia*”, Jarchow and Cook chose TS3 over the other two targets (Jarchow & Cook, 2009). The fact that they included the property of allelopathy in an otherwise traditional target system of population ecology gives us a lot of information about their choice. The particular partition shows us that they are part of what I call *trait-level* invasion research, a sub-field in invasion biology which takes its concepts and methods from population ecology.<sup>1</sup> Trait-level researchers believe that the best way to study invasions is to identify the traits which make an invader a good competitor. In contrast, if they had chosen TS1 they would be part of *community-level* invasion research, a sub-field which comes from conservation ecology and uses very different methods for studying invasions. These partitions are larger as their main units are entire communities, with aggregate or emergent properties such as biodiversity, biotic resistance and disturbance. In terms of abstraction, it is significant that they included an additional property not usually included in standard competition studies. This means that they thought that allelopathy was important enough to be included explicitly in the target system.

In order to understand *why* TS3 was chosen over the other two targets, we must look at the method used to study the invasion. Jarchow and Cook set up an experiment, consisting of four different treatments, some of which contained activated carbon. Activated carbon can counter the effect of the allelopathic chemicals in the soil/water, because it absorbs hydrophobic molecules (such as the toxins exuded by allelopathic plants) but does not absorb hydrophilic molecules (which include most of the nutrients in the

---

<sup>1</sup> The existence of these two distinct sub-fields in invasion biology is well documented, as are their differences in conceptualizing and studying invasions {Richardson:2011un}. For a full examination of these two fields and their history, see chapter 6.

soil such as Nitrogen and Phosphorus). Thus, if allelopathy was the cause of the competitive advantage and hence the invasion of *T. angustifolia*, then we would expect to see its competitive advantage reduced, when grown with the activated carbon.

This experimental setup requires a partition at the level of individuals, and cannot work with a partition at the level of communities. The results of the experiment are measured in the growth rate and size of individual plants after they have competed with each other. This type of measurement cannot happen (or at least will be extremely difficult to accomplish) with a partition which measures biodiversity and disturbance. Thus, it seems that the particular partition goes hand in hand with the experimental setup and a particular sub-field in invasion research. In fact, it seems that belonging to a particular sub-field determines the experimental set up, or that the choice of set-up brings with it the concepts and norms of the subfield. Either way, both of these affect the choice of partition as not all partitions work equally well with each experimental setup. In other words, some partitions are more *useful* than others given a particular experiment or model.

While usefulness of a partition may explain the choice between TS3 and TS1, it does not explain the choice of TS3 over TS2, as both targets have the same partition. Nonetheless, as stated above, they differ in terms of abstraction as TS3 includes allelopathy but TS2 does not. Why then did Jarchow and Cook choose a level of abstraction which broke with tradition. The answer is that they must have thought that TS2 and the experiment associated with it, would not provide them with a good enough explanation of the invasion. Instead they hypothesized that including allelopathy would give them a better explanation of the invasion. In other words, they hypothesized that allelopathy was *relevant* for the explaining the invasion.

This discussion is meant to show that the criterion for a successful partition (usefulness) is distinct from the criterion for a successful set of abstractions (relevance). Thus, a complete theory of target system evaluation should take into account both of these criteria. It should also determine whether there is some relationship between the two or if they are entirely independent.

In this chapter I will give a more general account of target system evaluation, based on usefulness and relevance. I will argue that they are not entirely independent criteria, and will therefore refer to the theory as a whole as a theory of *aptness*. In order to gain a deeper understanding of the notions of usefulness and relevance, I will look an account of salience and an account of relevance from the literature

on scientific explanation. I will then show similar accounts can be constructed for the evaluation of target systems.

## 2. Aptness

The point of departure for a theory of target system evaluation, is that target systems cannot be evaluated in themselves. That is, target systems are real parts of the world, hence they cannot be considered better or worse *per se*. In the example given above TS1, TS2 and TS3 are all just parts of the world and no one is inherently better than the others. However, as I showed in the introduction, as soon as the scientists' goal is taken into account, TS3 can be judged as the best of all three options. In other words, any account of target system evaluation must take into account the aims of the scientists and the context of their scientific practice.

This notion of target system evaluation is entirely coherent with my view on the nature of target systems. Most importantly, a pluralistic notion of partitioning means that there will be more than one possible partition of each domain, and those partitions will not be distinguishable through ontological criteria (see chapter 3). In addition, it will possible to abstract in many ways from each partition, again without changing the ontological status of the target system (see chapter 4). Therefore, it is not possible to distinguish between targets on ontological grounds. Instead, the criteria for target system evaluation will be epistemic and in some cases, pragmatic. I hope to show that Despite its pragmatism, my theory of target system evaluation is sufficiently strict to be a useful tool for scientific practice.

I should note that this theory of target system evaluation is only directed towards target systems which are unsuccessful because they are inapt, yet still real target systems. To clarify, I showed in chapter 2 that there are three ways in which target system specification can fail: cases where target systems were inapt, cases which involved false beliefs about targets and cases where target systems were not specified at all. The theory of target system evaluation is restricted to the first case, as this is only case where target system actually exist (even though they are not optimal). In the other two cases, scientists are not actually succeeding in specifying a target system, hence the theory cannot be of use. In those cases, scientists must first realize their failure to specify a target and specify an actual target before it can be evaluated in terms of aptness.

The notion of ‘aptness’ is itself apt for target system evaluation, as it implies context dependence even in our everyday use of the term. What we usually mean when we say that a subject or object is apt, is that it is suitable or appropriate given a particular purpose or set of circumstances. Imagine a situation with three friends: Sara, Aisling and Meera. Sara has just received a glowing set of teaching evaluations from her students. Aisling, her friend, might give an apt description of her emotional state to their mutual friend Meera by saying that Sara feels happy but also vindicated, given the large amount of time and effort that went into preparing for her course. The description is apt because it takes into account the context of Sara’s emotional state, identifies the important causal factors that give rise to Sara’s set of emotions and conveys all the important information to Meera. In contrast, we would think that Sara’s description would be inapt if it did not contain all the relevant information, or if it contained a wealth of irrelevant information. In addition we would imagine that Sara would not be able to communicate her point to Meera if she used a method of communication which Meera could not understand (such as ASL). In all of these alternative cases, the words (or signs) that Sara uttered would not be problematic in themselves, but only given the context of the conversation.

In the case of target systems, the situation is somewhat different, as targets are parts of the world and not representations of parts of the world like models or utterances. That is, a scientist specifies target systems in order to provide explanations of a phenomenon, yet does so by representing it through a model or experiment. Also, in the case of target systems, the context can be understood generally as the framework in which the phenomenon is conceptualized, but it also contains the specific model or experimental setup. Thus, the characterization of the context is often more precise than it is in everyday natural language.

I find it useful to think of target systems as containing causal factors which give rise to the phenomena we are trying to explain through science. This may seem strange, as a traditional way of understanding causal relations are between two (or more) *events*, the cause(s) and the effect. Yet, it is not controversial to think of a phenomenon we are trying to explain as an event, or the parts of a system as causal factors. In the standard example used to illustrate causation, the event is the a billiard ball moving (or falling into a pocket). The ball itself and the cue ball that hit it are part of the system and they are related to each other in a causal way. Yet, on my view, models and experiments are tied to target systems, and those target systems are are spatiotemporally strictly defined. For instance, we can think of the invasion of

T.angustifolia into Indian Creek as an event, and the parts of the system as factors which have some kind of causal relation to each other and which somehow bring about the event. At this point, I do not wish to say any more about the metaphysics of the causal factors or the causal relation itself, as I think it can remain open for now. The account of causal explanation I will use in the next section also has a liberal notion of the metaphysical aspect of causation.

With all these pieces in place, it is now possible to give a general definition of aptness for target systems:

- (1) A target system  $T$  is apt for a scientific purpose  $S$  just in case:
  - (i) its partition  $P$  is *useful* for  $S$  and
  - (ii) It contains the *relevant* causal factors of  $S$ .

At this stage, the definition is not particularly illuminating, as there needs to be an analysis of the terms ‘usefulness’ and ‘relevance’. In order to do this, I will make use of two accounts, both from the literature on scientific explanation. The first is Bas van Fraassen’s account of salience as part of the ‘pragmatics of explanation’ and the second is Michael Strevens’s account of relevance in his *kairitic* account of causal explanation.

I should note that scientific explanation is by no means the only place where relevance is discussed. In fact, relevance is key in philosophy of cognition and philosophy of language, as it enters into pragmatic aspects of communication. That is, communication involves much more than just the syntax and semantics of a speaker’s utterances, as it depends on the context in which the utterances are made and the expectations of the interlocutors. For example, Grice observed that a lot of conversation relies on information that has not been directly stated, but implied. He also thought that conversations should be viewed as cooperative enterprises, where speakers are trying to understand each other and make themselves understood (W. Davis, n.d.). In order to do this, speakers must try to obey maxims, which fall into certain categories. The category of interest is that of Relation, whose maxim is “be relevant”. In another account, Sperber and Wilson (1995) go beyond language and communication as think that relevance is key for understanding human cognition. The main idea is that “Human cognition is geared towards the

maximization of relevance (that is, to the achievement of as many contextual (cognitive) effects as possible for as little processing effort as possible)” (Korta & Perry, n.d.).

While these are important and comprehensive accounts of relevance, I think that the scientific explanation approach is better suited to understanding usefulness and relevance in the case of target systems. There are two reasons for this. The first is that the type of ‘context’ investigated in scientific explanation is the same as in the case of target systems, and it is affected by scientific disciplines, norms and practices. While these may have analogues in communication and cognition, it is simpler to use the theory where context is restricted to scientific practice.

The second reason is more complicated. As I will explain further on, most accounts of pragmatics or relevance rely on one (or a few) main principle(s) to explain both contextual factors and relevance. This is also true of Strevens and van Fraassen, yet they are interesting because the former thinks that relevance is key and influences the context whereas the latter thinks that context is key and influences relevance. This makes the accounts complimentary, provided that they can be combined under a single umbrella account. I think that this is exactly what is needed in my account of aptness, as I want to distinguish between usefulness of partitions (which can be illuminated by van Fraassen’s account of salience) and relevance of abstractions (which can be illuminated by Strevens’s kairetic account of causal explanation).

In the next section I will give an overview of these two accounts, situated in a more general (albeit brief) summary of some topics in scientific explanation.

### 3. Scientific Explanation

The topic of scientific explanation is arguably one of the most important debates in philosophy of science (Woodward, 2009). The debate has many dimensions, yet the basic issue is to determine the criteria which make an explanation of a scientific phenomenon successful. In the early 20th century, the received view in philosophy of science was that a successful scientific explanation of a phenomenon showed how the phenomenon in question was an instance of a more general pattern. Examples include Hempel and Oppenheim’s covering law model of scientific explanation (Hempel & Oppenheim, 1948) and Kitcher’s unification account (Kitcher, 1981; 1989). Moreover, Hempel and Oppenheim also thought that explanation and prediction are symmetric as they have the same logical structure. That is, the structure of a scientific explanation contains laws and antecedent conditions and the explanation follows logically from these.

While Hempel and Oppenheim's view of scientific explanation gained some popularity it then came under heavy attack. Much of the criticism came in the form of counterexamples, many of which have now become famous in their own right (Salmon, 1999). More specifically, some counterarguments brought to light that the covering law model did not account for causal factors in the explanation of an event. Thus, for example, we can derive the length of a flagpole's shadow, given the height of the pole, the angle of the sun and laws governing the rectilinear propagation of light. However, the covering law model also allows us to determine the length of the flagpole from the same laws, the angle of the sun and the length of the shadow. The problem with this explanation, however, is that while the flagpole's height causes the shadow and its length, the existence of the shadow cannot cause the flagpole, nor can it cause the flagpole to have a certain height.

Problems like these led to a number of causal accounts of scientific explanation, which focused explicitly on the factors which caused a particular phenomenon (Salmon, 2006). Views differ on many dimensions, including the nature of causation and the adequate method for determining the causal factors relevant to an explanation. The main issue for causal accounts of explanation must be able to determine which causal factors are actually relevant for a particular explanation. Each phenomenon is influenced by a great number of causal factors, yet not all of these factors have the same degree of influence. Most of these factors will have a negligible degree of influence, as they will be spatiotemporally distant (Strevens, 2008).

The problem is that an explanation which sites all of the causal influences or even just most of them will be extremely laborious and time consuming. For example, if we had to specify the gravitational force of every single planet in the solar system, the ambient temperature, the elasticity of the materials in an explanation of the length of the flagpole's shadow, scientific investigation would not get off the ground. More importantly, it seems intuitive that an a true explanation of an event must isolate the factors which *actually* caused the event, from the factors which merely made the event's taking place possible. In other words, not every causal factor that gives rise to an event is explanatory. The problem is that if irrelevant factors are left in the explanation, it can lose explanatory power as irrelevant factors might obscure the explanation of an event.

While there are numerous accounts of explanatory and causal relevance, I will focus on one of these accounts, which I think is both promising and easily applicable to target system evaluation. This is the *kairitic* account of causal explanation, which I will examine in the following section (2.1).



A second criticism of the covering law model is that it is too formalized, as it only accounts for the syntactic and semantic aspects of explanation, but not its pragmatic aspects (Salmon, 1999). According to Bas van Fraassen, it is impossible to construct a theory of scientific explanation or a method for evaluating explanations without these pragmatic factors. Instead determining what makes an explanation successful must take into account the context in which the explanation is put forward. This pragmatic account of explanation focusing in salience will be the key to explicating the term ‘usefulness’ in my account of aptness. I will explain the basics of van Fraassen’s account in the next section.

### 3.1 The Pragmatics of Explanation - van Fraassen

Bas van Fraassen’s discussion of scientific explanation begins from the view that scientific language is very similar to natural language, as it is characterized by syntactic, semantic and pragmatic factors. Therefore, a full account of scientific explanation must include a discussion of these pragmatic factors. According to van Fraassen, Hempel and Oppenheim’s covering law model was doomed to fail, because it focused exclusively on syntax and semantics (Van Fraassen, 1980).

According to van Fraassen, the problems started when Hempel and Oppenheim attempted to give an objective, context-independent theory of scientific explanation. Moreover, he does not think that the alternative causal accounts of explanation address this problem. For van Fraassen, giving an account of relevance is not sufficient for generating a successful theory of scientific explanation. This is because he thinks of relevance as a weaker notion, which does not capture the importance of pragmatic factors in scientific explanation. He argues that once the context of an explanation is determined in a sufficiently fine-grained manner, the relevant factors will be easily determinable. He states that “*the context, ... determines relevance* in a way that goes well beyond the statistical relevance about which our scientific theories give us information” (ibid p. 128).

I will explain this point with the use of a term coined by Peter Railton, the *ideal explanatory text* (IET) (Railton, 1981). This vast hypothetical document contains all the facts, laws, causal connections and hidden mechanisms of a particular phenomenon or event. What counts as a successful explanation of said event is an explanation that uses elements of the ideal explanatory text. Of course, no explanation can include everything in the IET, so selections must be made. The problem is that *everything* in the IET is

*relevant* for the phenomenon and can be used as a sufficiently good explanation. But if everything is relevant, then how are we justified in using explanations which do not reference the IET in its entirety?

According to van Fraassen, this is where pragmatics come in. While everything in the ideal explanatory text is relevant *for the explanation of the phenomenon*, it is not the case that everything in the IET is *salient* for the human beings providing the explanation or trying to understand it. For example, we give many different explanations of why a brick broke a window, many of which are generated from the IET. However, explaining the event in terms of molecular structure and properties, would be futile, if the audience consisted of young children. The salient aspects of the IET will be differ depending on the background knowledge and interests of the humans involved in the explanation.

However, picking out the salient factors from an IET is what is meant by ‘picking out the relevant causes of an explanation’. All the factors in the IET are causal, and all of them are contributing factors, however, the particular subset which is picked out as an explanation is determined by factors which are considered salient. He states that “the salient feature picket out as ‘the cause’ in that complex process, is salient to a given person because of his orientation, his interests, and various other peculiarities in the way he approaches or comes to know the problem - contextual factors” (van Fraassen, 1980) p.125.

The positive part of van Fraassen’s argument begins from the idea that scientific explanations should be understood as answers to *explanation-seeking why-questions* (ibid 126). That is, the explanation of a phenomenon  $P$  can be understood as an answer to the question Why  $P$ ? The point of involving questions is that working on clarifying the question both illuminates the phenomenon being explained and constrains the answer. In more formalized terms, a why-question  $Q$  is determined by three factors:

- (i) the topic  $P_k$
- (ii) the contrast class  $X = \{P_1, \dots P_k, \dots \}$  and
- (iii) the relevance relation  $R$

Yet these three factors are not equal partners in the relationship. In order to determine whether a proposition is *relevant* to a question  $Q$ , we must determine whether it is relevant *to a certain topic with respect to that contrast class*. I should note that there is not one kind of relevance relation, but many relations could be considered relevant, given the context. Thus, for example, in this case, the relation of ‘affecting the tides’ will not be relevant but the relation of ‘affecting plant growth’ will be relevant.

This process of clarification and determination of relevance is primarily a pragmatic endeavor. Consider again the example mentioned in the introduction, where *T. angustifolia* has invaded the North American lakes. If our aim is to explain the phenomenon of invasion, an obvious why-question is: Why did *T. angustifolia* invade the North American lakes? However, as van Fraassen points out, this question needs pragmatic clarification. By asking this question I could be interested in at least three different questions: (a) why it was *T. angustifolia* rather than another organism which invaded (b) why it *invaded* rather than failed to invade or (c) why it invaded the *North American lakes* rather than another community.

According to van Fraassen, these are more than merely different interpretations of the same question, they are actually different questions which have distinct *contrast classes*. A contrast class is simply a set of alternatives, hence the contrast class of (a) includes other species of cattail, other plants of a different genus, and could even include other objects or animals. The contrast class of (b) includes other interactions, such as competition, predation, etc. Identifying the contrast class is extremely important, as it determines what constitutes an explanation of an event or phenomenon. That is, once we have identified the contrast class of the question, we are in a much better position to know what it is that we are explaining and how we should go about explaining it.

In the case of the *T. angustifolia* invasion, the way in which the question is clarified is very illuminating. If, for instance a group of scientists chose (a), then they would be focusing on the invader. In asking why it was that this particular organism rather than another managed to invade, the scientists are asking what it is about *T. angustifolia* which makes it a good invader, that is, what traits does *T. angustifolia* which other plants do not have? Scientists who choose this clarification and contrast class are part of what I call *trait-level* invasion research. This is a sub-field within invasion biology which has its roots in population ecology, and uses many of the concepts, models and experiments developed by population ecologists. On the other hand, scientists who choose (c) are focusing on the community which is being invaded and asking what traits of *the North American lakes* community made it susceptible to invasion. These scientists are part of what I call *community-level* invasion research which has its roots in conservation ecology and studies invasions with its own set of concepts, models and experiments. In short, the way in which a question is clarified affects the concepts and methods used in answering that question, hence it is an extremely important aspect of the explanation itself.

If it is determined that  $Q$  arises in the right context (that is, if it is an appropriate question given the context), then it can be *answered directly*, in the form (\*)  $P_k$  in contrast to (the rest of)  $X$  because  $A$ , provided that the following conditions hold:

- (i)  $A$  is true.
- (ii)  $P_k$  is true.
- (iii) No member of  $X$  other than  $P_k$  is true.
- (iv)  $A$  bears  $R$  to  $\langle P_k, X \rangle$ .

According to van Fraassen, we are now in a position to evaluate the answer (\*), in probabilistic terms, depending on our prior probabilities and the relationship of  $A$  to the probabilities of other reasons.

The important upshot of this discussion is that the factors which are considered relevant, are only considered thus *because of* the context which frames the explanation. In other words, this is a one-sided view of the relationship between context and relevance where context determines relevance but not vice versa. In addition, the connection between the two is quite strong, as determining relevance often follows easily from the determination of the context.

Having given a summary of van Fraassen's account of salience, I will now show how the notions of 'why-questions' and 'contrast class' can help to explicate the notion of usefulness for partitions.

### 3.2 Useful Partitions

Usefulness is an inherently pragmatic criterion for evaluation, as something can only be useful given a certain purpose and context. However, this context dependence does not imply that there are no rules for determining when something is useful, or at least when it is more useful than something else. In the case of aptness, this means that we can use contextual factors to determine which of the available partitions will be useful for a particular scientific purpose.

Before I show how this can be done in actual scientific practice, I want to try and give a more precise definition of 'usefulness' for partitions. What makes a partition useful? As I argued in chapter 3, a partition is a way of grouping the contents of a domain of study into units. Thus, a partition is useful when it groups the contents of a domain in the 'right way'. But what does it mean to group the contents in the right way? Given that usefulness is a pragmatic criterion, the 'right way' of grouping results in units which can be used for a particular scientific purpose, where a scientific purpose is the study of a phenomenon of

interest. Yet the study of a phenomenon of interest is constrained by the tools and methods available to the scientist. Hence, the scientist must partition the domain in such a way as to individuate units which can be used in the study of the phenomenon. This means that a partition is *maximally useful* when it individuates *all* the units which can be used in the study of the phenomenon of interest.

There are two points to note. The first is that this individuation does not necessarily mean that the finest-grained partition is always the most useful. For a start, fineness of grain depends on factors such as the possibility of measurement and technology, which means that there is a lower (and a higher) limit to the grain. Moreover, many scientific partitions are multi-level, as they contain units of different sizes and at different levels of grain. The point of the criterion is to make sure that all units which can be used to study the phenomenon are categorized or individuated *as units*.

The second point to note is that this individuation is not meant to show which units will *actually* be used to study the phenomenon, as this is the job of abstraction. Abstraction is the process by which a scientist identifies the *relevant* units from the others. However, in order for it to be possible to identify which units are relevant and which are not, the scientist must have first categorized them as units. Thus, a partitioned domain will have a large number of units, yet in order to be maximally useful, all the units which could turn out to be relevant must be individuated as units.

Of course, it will probably never be possible to identify the maximally useful partition for the study of a phenomenon. However, we can still keep the idea of maximally useful partition as a goal to which we can strive for in our scientific endeavors. Thus, instead of trying to find the unique maximally useful partition for each scientific purpose, we can start at the other end and try to eliminate those partitions which we know will not be useful.

I should note that given that the starting point for any partition is the domain of interest, we cannot include any units which are not actually part of the system. This is not because they will not be useful, but because a partition is a categorization of what is actually there, and cannot include anything additional. Thus, for example, units such as the plants in the scientist's garden will not be part of the partition (even if they are individuals of the same species as those found in Indian Creek). In addition, units which are not located in the system at the time in which the study is taking place will also not be included. Thus, for example, the prehistoric species of plants which once occupied Indian Creek cannot be part of the domain.

With the domain is spatiotemporally defined, we can start eliminating partitions which are not useful. The first set of partitions which can be eliminated are the ones which do provide scientific classifications. This is determined by the fact that the phenomenon of interest is a scientific phenomenon. For example, a partition which classified the plants of Indian Creek into two units of ‘magical’ and ‘non-magical’ would not be useful for studying the invasion of Indian Creek by *T. angustifolia*, as they cannot be used in the study.

The second set of partitions which can be eliminated are the ones which will not be useful, given the specifics of the phenomenon of interest. Within science, there are a number of ways in which a particular phenomenon can be studied. A scientist might be looking for an explanation of a phenomenon, a prediction about the phenomenon, a way of manipulating the phenomenon etc. Each type of scientific endeavor can be formulated as a different kind of question. In our example, the phenomenon is framed in explanatory terms, hence it is formulated as the why-question “Why did *T. angustifolia* invade Indian Creek?” While van Fraassen only considers the contrast classes within the framework of why-questions, we can also think that the entire question has a contrast class. That is, if we were interested in prediction, then we could ask the question “Will *T. angustifolia* invade Indian Creek?” or even “When will *T. angustifolia* invade Indian Creek?”.<sup>2</sup>

The next step is to consider the contrast classes of the actual why-question so we can further narrow down the number of useful partitions. The easiest one is the contrast class of ‘invade’. The contrast class here includes other relations between units in the system. We can first eliminate all non-ecological relations, as we are within the framework of invasion biology, which is part of ecology. We can also eliminate other ecological relations such as predation, facilitation and intraspecific competition.

We are now left with two of the questions which came up in the previous section: (i) why it was *T. angustifolia* rather than another organism which invaded and (ii) why it invaded the *North American lakes* rather than another community. As I stated in the previous section, each of these two questions represents a different sub-field in invasion biology. The scientists in the paper from which the example was taken are answering the first question, which means that the partitions associated with the second question are also ruled out. This leaves the partitions associated with trait-level invasion research. All of these partitions

---

<sup>2</sup> I should note that there is another constraint implied here. The predictive questions about the invasion of Indian Creek would make for an unsuccessful study, as the domain has already been invaded. This means that the domain itself precludes the study of some phenomena in some ways.

count organisms as units but some of them also contain smaller units such as nitrogen, phosphorus and the chemical compounds which make up toxins. Of these, the scientists will choose one of these multilevel partitions. At this stage, the partitions will all be useful, and any difference between them will be minimal.

To sum up, a partition is useful when it individuates units in a domain such that the phenomenon can be studied. In addition, we now have an answer to the question posed in the introduction: how should scientists study the phenomenon of the *T. Angustifolia* invasion of Indian Creek? The answer is that they should choose a partition from trait-level invasion research. This kind of analysis helps when the context of a field or sub-field is known and when methods that are already well established are used to study the phenomenon in a different domain. Another way of putting this point is that if we are within a Kuhnian paradigm engaged in normal science, the context of the paradigm will determine the useful partitions for each puzzle.

The problem with this answer is that it is overly reliant on the norms of the field (or paradigm) and does not give us a way to determine whether the methods of one field are better than another. In this example, the field is already established, which makes the choice for Jarchow and Cook easy. However, what if community-level research is actually a better way to study invasions? Merely relying on context and usefulness of partitions will not help us in this case. Still, there is another criterion with which we can evaluate target systems, namely *relevance*. The account of relevance which I will use, is aimed at finding the factors which *make a difference* in the occurrence of an event. While the relevance of a particular factor is affected by the context (and in the case of target system, the chosen partition), identifying the difference makers can also affect the usefulness of a partition. These two criteria together will determine the aptness of a target system.

### 3.3 *The kairetic account of causal explanation* - Strevens

Michael Strevens's account of causal explanation focuses on the factors which *make a difference* to the explanation of an event. His motivation for developing the kairetic account was to give an account of explanation which combines aspects of classical causal accounts of explanation and the unification approach to explanation. A causal explanation identifies an event's causal history, while a unification explanation aims to explain an event using a theory that unifies many other phenomena. An important feature of unification accounts is that they can account for more general explanations, because they do not

aim for causal accuracy. In fact, unification accounts are often enhanced by omitting or changing details (Strevens, 2004). Causal accounts, as I showed above, face the problem of *explanatory relevance*.

For Strevens, the key for explaining an event is to pick out, from the complex causal network, those factors which *made a difference* to whether or not the event in question occurred (Strevens 2004). Strevens's germane example concerns the death of Rasputin. Rasputin was famously poisoned, shot (twice) and finally tied up and drowned. Strevens's account allows us to make sense of this complex situation, by helping us to identify *drowning* as the factor which ultimately caused Rasputin's death.

To this end, Strevens proposes a test called the *eliminative procedure*, with which we can determine the difference makers in the explanation of an event *E*. Before we can administer the procedure, we must first determine the event to be explained. Each event is individuated at a particular level of grain, which is determined by the explanatory request. That is, the way we ask the question about event *E*, determines how finely we specify that event, and how much causal detail will be present in the explanation. A fine-grained individuation of an event will elicit an explanation with more difference makers than a coarse-grained individuation. For example, asking why Rasputin died will elicit an explanation with fewer causal details than the explanation elicited by the question 'why did Rasputin die in exactly the way he did?' (ibid 159).

Once the event has been specified, we look at the causal network in which the event *E* is embedded, and pick out a part of that web that was sufficient to cause *E*. By 'part', Strevens means a set of actual initial conditions and laws that could, in itself, cause *E*. We then construct a representation of this set of actual causes, in terms of propositions. This is the *veridical deterministic causal model* of *E* (ibid 162). Strevens calls it a veridical model, because the conditions are actual causal influences on the event. In the model, relations of causal production are represented through relations of logical entailment. In other words, Strevens uses the notion of logical entailment to represent a situation in the real world, where a set of initial conditions produce an event in virtue of laws (ibid 163). The second step is to eliminate from the causal model any factors that are not necessary in the causal production of *E*. The factors that are not necessary are the ones that do not play a role in the entailment of *E*. What remains is the set of initial conditions and laws that *made a difference* to the occurrence of *E*. It is called the *explanatory kernel* of *E*.

I will illustrate how the eliminative procedure works through Strevens's own example of Rasputin's death (*E*). The event is specified through the question "What caused Rasputin's death?", and the



causal network includes the drowning, the shooting and the poisoning, but also a number of other factors which have some causal influence on  $E$ , such as “the length of Rasputin’s beard, the day’s pollen count, the gravitational influence of Mars” (ibid 158). The next step is to determine the veridical causal model for Rasputin’s death and eliminate all the factors which can be eliminated from it.

It is easy to see why drowning counts as a difference maker. ‘Being thrown in the river’ is part of the causal network that *entails*  $E$ , therefore is part of the veridical model. Also contained is a law which states that ‘people thrown in the river under certain conditions die’, and the conditions for which the law holds. Therefore, we cannot remove the being thrown into the river without invalidating the entailment. This means that being thrown into the river is a difference maker of  $E$ . Contrast this with explanation of Rasputin’s death by poisoning. To do so we would construct a causal model which contained the proposition that Rasputin was poisoned, together with a law which stated that a person poisoned under a set of conditions, will die. The problem is that in Rasputin’s case, these conditions did not hold. We now have two options. Either we bite the bullet and state the conditions were present or we eliminate the conditions from the model. However, if we choose the first option, we end up with a causal model which is not veridical, as it asserts something which not the case. On the other hand, if we eliminate the conditions from the causal model, then we are left with a model which does not entail  $E$ . This is because, as Rasputin showed, simply being given a certain poison does always result in death. Thus, the poisoning cannot be considered a difference-maker for Rasputin’s death.

The eliminative procedure also gives us a precise method for eliminating the tangential factors which are definitely not difference makers yet have some sort of weak causal influence, such as the gravitational force of Mars, the length of Rasputin’s beard etc. Strevens’s explanation focuses on the ‘gravitational force of Mars’, which is part of the veridical causal model, as it is present and has some influence over events on Earth. However, it can only be part of the veridical causal model if the ‘being thrown in the river’ is also part of the model (otherwise the model would cease to be veridical). Yet as long as the ‘being thrown in the river’ is part of the causal model, then the ‘gravitational force of Mars’ is superfluous. Removing it from the model will not affect the model’s entailment of  $E$ , as the actual cause of  $E$  (being thrown in the river) is sufficient to have caused it. Thus the ‘gravitational force of Mars’ and all similarly superfluous causal factors are not difference makers and can be eliminated.

Strevens also considers a trickier situation. What if the ‘gravitational force of Mars’ entailed another factor which *was* a difference maker (being thrown into the river)? It seems that in this case, it would be a difference maker, and we would have to keep it in the model. This would spell trouble for the kairetic account, yet Strevens has the means to overcome it. He argues that that it is not possible to construct a veridical causal model where a factor such as the ‘gravitational force of Mars’ entails a difference maker such as ‘being thrown into the river.’ This is because the notion of entailment he uses here *represents* an actual causal relation in the real world. Yet there is no physical system where the gravitational force of Mars causes someone to be thrown in a river. Thus, the model that included the gravitational force of Mars *as entailing* a difference maker, would not be veridical.

To sum up, the kairetic account of causal explanation, gives us a method for determining which factors from a complex causal network actually make a difference in an event’s being brought about. In the next section, I will show how some aspects of this method can be applied to target system evaluation, so that we can determine whether the factors in the target system are relevant and whether those omitted from the partition are irrelevant.

### 3.4 Relevant Abstractions

While there are many similarities between causal explanation and target system analysis, there are also some important differences, which mean that ‘applying’ an account of causal explanation to target systems is not entirely straightforward. The most important difference is that target systems are real parts of the world, whereas causal explanations are representations of actual factors and events in the world. Still, this difference does not need to be problematic, provided that the some adjustments are made for the different ontological status of targets and causal models.

The second issue is that, at least on my account of target system evaluation, the criterion of usefulness and the selection of a partition renders a substantial part of Strevens’s account redundant. On my view, there is no need to follow the first steps of the kairetic account, namely specifying the event, identifying the causal network and constructing the veridical causal model. In the case of target systems, specifying the domain of study and the phenomenon of interest are the equivalent of specifying the event. In addition, the role played by the identification of the causal network and the veridical causal model in the kairetic account, is played by the choice of partition in my account of target systems.

There is an important functional similarity between partitions and veridical causal models. A partition is not a model, but a particular categorization of the domain. Also, a veridical causal model is constructed so that it contains all the causal factors which influence the event, yet the partition contains everything in the domain categorized in a particular way. Thus, while the veridical causal model is the selection of a subset of the causal network, in the partition nothing which is part of the domain can be omitted. This means that a partition might contain factors which have little or no influence on the phenomenon of interest. Still, the important point of similarity is the ‘veridical’ nature of the causal model and the partition. The causal model is veridical by definition, as it contains all the causal factors which influence the phenomenon. The partition is veridical in the sense that it is real, and contains all the factors which are present in the domain.

In the case of the *T. angustifolia* invasion, the partition will contain everything that exists in the Indian Creek ecosystem, categorized in a particular way. This includes: all the organisms in the ecosystem identified as organisms, all the properties of these organisms (i.e. their traits, such as allelopathy, phenotypic plasticity, births, deaths etc), abiotic factors (such as nitrogen, phosphorus, water, radiation, toxins and some of their properties (such as temperature, salinity, velocity (of the wind))). We can now apply the eliminative procedure to this partition, in order to determine which of the parts and properties are difference makers for the phenomenon of interest.

An obvious candidate is *T. angustifolia* itself, and at least some of its properties (births, deaths, allelopathy). If we eliminate *T. angustifolia* these properties from the target system, then we will have a target systems whose factors do not entail the phenomenon. That is, we cannot explain an invasion (in the context of trait-level research) without the invader or the properties which help it outcompete the native plants.

The scientist must also choose how to represent the community in their experiment. One option would be to include a number of different species, yet a closer examination of the system shows that some species are more important than others. The plant which seems to have suffered the most is *Bolboschoenus fluviatilis*, probably because it was in direct competition from *T. angustifolia* and the direct recipient of the toxins. In other words, the invasion might not have occurred if a different plant was the main competitor of *T. angustifolia*. This means that the *B. fluviatilis* and some of its properties are difference makers in the invasion of Indian Creek. In contrast, the rest of the organisms in the system are not difference makers.

While it is true that other plants might have suffered from competition with *T. angustifolia*, yet they have not suffered as much, because they were not in direct competition with *T. angustifolia*. In other words, the invasion would still have occurred if the other plants were absent. The animals are not relevant because changes to their populations would occur as effects of the invasion, not causes of it. Animals could prey on particular plants or use native plants for cover, and hence suffer because their resources are reduced by the invasion. Nonetheless there would come about as a result of the invasion and hence not important for explanation of why the invasion occurred.

There are a number of other factors in the partition which are definitely not difference makers. These have the same status that the gravitational force of Mars has on Rasputin's death. Factors of the ecosystem that are constant over time, or between the two plant populations will not be important. For example, the amount of radiation, precipitation and temperature are not likely to be relevant. The close proximity of the two plant populations means that these factors will not affect the growth or reproductive rate of the plants. Similarly, with exception of a case where the toxins inhibit the uptake of a particular nutrient, nutrient availability will not be relevant, because nutrients tend to be dispersed uniformly in lakes.

There is a third category of factors, which are neither obviously relevant or irrelevant. These could be factors that are only relevant in some invasions but not others, or factors whose effects we do not clearly understand. For example, we might not know if the acidity or salinity of the water is important, because it may affect the spread of the allelopathic toxins or their uptake by the native plants.

To sum up the argument so far, applying the eliminative procedure to the partitioned domain helps us to identify the factors which are relevant for causing the phenomenon. Another way of putting this point is that it can help us determine the right level of abstraction given a particular partition and phenomenon of interest. By applying the eliminating procedure we can abstract as much as possible, thus specifying a target which is simpler and easier to experiment on or model. The eliminative procedure is a way to test our abstractions so that we can determine if we are abstracting too much or not enough.

#### 4. A General Theory of Aptness

Now that all the parts of my theory of aptness are in place, I can put them together and give a general definition of aptness for target systems. This definition is pragmatic to a certain extent, as the context of a scientific purpose provides the criterion for usefulness, which in turn places constraints on what is considered relevant, yet this does not mean that any target system is equally apt for given a certain context. Within the contextual framework, there is a fact of the matter about which factors are relevant and which are not, hence there are strict rules for evaluating target systems. A phenomenon is either caused by a combination of factors or it is not. The pragmatic aspect constrains how we individuate those factors, yet the world itself determines whether and how the phenomenon occurs. Therefore, provided that a partition is sufficiently useful, we have an objective measure for determining which factors are relevant and which can be omitted.

This point will be made clearer by the discussion of two issues. The first concerns the difference between a maximally apt target, and a sufficiently apt target. A target is maximally apt when its partition is maximally useful and its abstractions are maximally relevant. A maximally useful partition is one which individuates all the factors which could be relevant, and a maximally relevant set of abstractions omits all those factors which are not difference makers. The definition of a maximally apt target then becomes:

A target is maximally apt for the explanation of phenomenon  $P$  just in case:

- (i) it individuates all the causal factors which are potential difference-makers for  $P$
- (ii) all factors which are not difference-makers for  $P$  are omitted

The problem with a maximally apt target is that there are few, if any, cases where a scientist can specify a maximally apt target for her purposes, or indeed be in a position to know that she has specified such a target. These difficulties can arise in relation to both partitioning and abstraction.

In the case of partitioning the difficulties can be grouped into two categories. The first occurs in cases where the contextual framework within which the scientist is working rules out the partitions which are most useful for studying the phenomenon. For example, if it were the case that biodiversity was a difference maker for the invasion of Indian Creek, the scientists operating within the trait-level framework

would not be in a position to determine that biodiversity was a limiting factor, as their framework excludes the partitions which individuate communities as units, with properties such as biodiversity. In fact, as I will show in the next chapters, there is an important debate in invasion biology, concerning the appropriate framework for studying invasions.

The second category deals with more extreme cases, where there is simply no available framework which can provide useful partitions for the study of the phenomenon. For example, it could turn out to be the case that none of the available frameworks individuate all the difference-makers for invasion biology. This is a real enough problem, which has led some invasion biologists to ‘go back to the drawing board’ and try to identify new frameworks for invasion biology (see chapter 6). This means that in some cases the contextual framework itself is limiting scientists from specifying the maximally apt partitions.

A possible answer to these problems is to dispense with pragmatics and contextualization in the evaluation of target systems and focus only on relevance. In fact, in many of the accounts of causal explanation, context plays a minimal role. In the kairetic account, for example, context is minimally relevant, as it helps individuate the event and the causal factors but it does not itself preclude other individuations. For Strevens, the causal model should always be made as abstract as possible, without invalidating the fact that the event is causally produced. Strevens’s notion of abstraction is not identical to mine, as it is not the omission of irrelevant factors, but the fineness of grain with the factors are individuated. Thus, for Strevens, a more abstract individuation is a less precise, or vague individuation of a factor. For example, in the case of Rasputin, the veridical causal model could be made less abstract by individuating the causes more precisely. Thus, the shots could be individuated, as could the precise location of the entry wounds. The drowning could be more precisely specified in terms of the length of time Rasputin would have to be underwater for him to drown, and so on.

According to Strevens, these modifications are not necessary for the veridical causal model of explaining Rasputin’s death because the more abstract model leads to the desired result. However, he agrees that if the event needed to be specified in a more fine-grained manner, then the causal factors would also have to be individuated more precisely (Strevens, 2004, p. 159). Yet it seems to me that in this case, context is a lot more important than the account implies. The problem I am trying to solve for target system evaluation is the analogue of being able to determine how the event of Rasputin’s death should be specified in the first place. For Strevens, this is not an important problem because he assumes that the reasons for

choosing one individuation over another are not themselves important, and also that it is very easy to change how the event is individuated. Still, this is an extremely important problem for target system evaluation.

There are also difficulties which pertain to the identification of difference makers, and create problems for identifying the maximally relevant set of abstractions for a phenomenon. Even if the partition is sufficiently useful, a scientist might not succeed in determining the difference-makers for the phenomenon. In real-world systems, it is not always possible to know what the causal connection between objects is, hence we cannot use the notion of entailment, as it represents the actual causal connections in the system. That is, we might not know how a particular object affects other objects in a system so we might not know if the object or facts about that object really are relevant. For example, one could hypothesize that in a complex ecosystem, there might be additional dynamics that affect the course or strength of the invasion. The existence of animal populations might have dynamical effects which affect the course of the invasion, by magnifying or diminishing it. However, if there is no available data which suggests that such an effect exists, then the scientists might not identify in the first place.

This difference has important implications. The notion of entailment is very important in the kairetic account because it helps us eliminate factors which have some influence yet are not difference makers. In Strevens's account the gravitational force of Mars is part of the veridical causal model, yet it is eliminated because we can be sure that Rasputin still dies if we not include it in the account. We are not always able to determine whether eliminating a factor will affect the system unless we actually eliminate it in an experiment. This means that we might not be able to know if our target is apt until we actually conduct the model or experiment.

The same is true in the other case that Strevens describes, namely the situation where the gravitational influence of Mars entails a factor that is a difference maker. In his example, we are able to determine that the gravitational influence of Mars does not entail someone being thrown in a river. However, we might not be able to determine entailment when dealing with target systems. We can imagine a situation where allelopathy, which is a difference maker, is inhibited by another factor in the system, which is usually not considered important. For example, it is conceivable that water temperature will affect allelopathy if *T. angustifolia* cannot produce toxins below a certain temperature, or if the toxins bind to

molecules in the water above a certain temperature. If this is the case, then we will not be able to know that a particular target is apt.

One way to deal with these problems is empirically, though a process of trial and error. As we cannot manipulate a target system in the way we can manipulate a model, we have to actually conduct an experiment in order to test whether the target is apt. Scientists can go back and forth between hypothesizing what sort of target is apt, conducting an experiment which fails and then going back and re-determining the target. For example, it could turn out that allelopathy did not cause aquatic invasions, because the toxins became diluted in the water and did not inhibit the growth of the native plants.

This should not be particularly surprising. After all, if we were always in a position to identify the difference makers of each phenomenon then there would be no need to actually conduct an experiment or a model. One of the most important reasons to practice science is so that we can determine the difference makers of a particular phenomenon. When we do, this knowledge becomes part of a scientific discipline and can be used later on, when a similar phenomenon or domain is being investigated.

Consider again the factor of temperature. Jarchow and Cook decided that it was not relevant for their purpose and omitted it from the target system. While they could not know that temperature was not relevant, they could make an informed guess given the body of knowledge of ecology, with which they are familiar.

We know that the temperature of the water in a lake varies within a certain range throughout the year and across years. However, we also know that this variation did not have an effect on the structure of the ecosystem before the arrival of *T. angustifolia*. We also know that since the invasion, the temperature of the water has continued to vary within the specified range, but this has not accelerated or decelerated the rate of the invasion. We know this because there has not been any change in the sizes of the native and invasive populations that correlates with the changes in water temperature. Moreover, we did not need to conduct any additional empirical tests to determine this. All we had to do is use already existing data to check whether there was variation in this particular factor, and whether this variation was correlated with any change in the system. This sort of data is often readily available to scientists working within a particular framework, so it is a useful way to determine whether a factor can be omitted from a system.

In the end, however, this data might not always be available, hence scientists will sometimes make mistakes and specify inapt target systems. In many of these cases, the only way forward is to move back



and forth between the model (or experiment), the target and the world, until a more successful target-model combination can be found. Still, I think that there are actually more constraints within the theory of aptness, as I have set it up, which will help minimize these cases.

The key for explaining to this point is to determine the exact relationship between usefulness and relevance. As I have already shown, each of these criteria alone is not sufficient to determine a maximally apt target system. Simply looking for the most useful partition does not automatically give us a way of determining the factors which are actually relevant, while it is impossible to determine the difference-makers independently of the partition and by extension, the contextual framework which affects the choice of partition. This means that the two criteria are not independent of each other. However, this also means that the two criteria are distinct from each other, and cannot be collapsed into one. Identifying the maximally useful partition still leaves us with a number of causal factors which are not difference makers, which have to be omitted from the target system, while a maximally relevant target still can only be evaluated as such given a certain context. These issues are compounded by the fact that we are rarely, if ever, in a position to identify either maximally useful partitions or maximally relevant abstractions.

However, in many cases, these two constraints together will be able to provide us with a *sufficiently apt target system*. The definition of a sufficiently apt target system is:

A target system is sufficiently apt for phenomenon  $P$  just in case

- (i) the partition individuates enough causal factors which could be difference-makers for  $P$  and
- (ii) the abstractions succeed in identifying enough of the difference-makers which causally entail  $P$

I have already shown how, given a certain partition, a scientist can identify enough of the difference-makers in order to have a sufficiently relevant set of abstractions. The problem which remains is how to identify a sufficiently useful partition, without relying entirely on the contextual framework. This is a more general version of the worry mentioned above, namely given there are two well-established contextual frameworks for studying invasion (individual-level and community-level) which one of the two provides the most useful partitions?

I will argue that in some cases, relevance can constrain usefulness, as testing for difference-makers can help determine the usefulness of a partition. The individuation of difference-makers is built into definition of usefulness, and this is not accidental. This implies that a partition which does not individuate difference-makers will not be useful. Yet it also implies that the difference-makers themselves can also cut across the boundaries determined by the context. That is, a partition which individuates fewer than sufficient difference-makers will be less useful than a partition which individuates sufficient difference makers.

I should note that there is no upper bound for identifying difference makers for the usefulness of a partition. That is, the more difference-makers that a partition individuates, the more useful it is. It can be argued that this implies that a more finely-grained partition will automatically be more useful but this is not so. The term 'individuation' here guards against this possibility. In order for a factor to be a difference-maker, it must make a difference as a causal factor. That is, it cannot be jointly sufficient with other causal factors to be a difference-maker (though it can be jointly sufficient with other difference-makers to entail the phenomenon). This means that a partition which is too fine-grained will individuate factors which are not themselves difference-makers. This will make the omission of irrelevant factors much more difficult, even impossible. For example, if it were the case that the community-level traits were difference-makers for invasion, then the partitions of individual-level research would be too fine-grained to be useful. That is, if emergent properties of communities (such as biodiversity) alone, resulted in resistance to invasion, then scientists using individual-level partitions would not be able to identify the relevant factors of the invasion.

However, as it turns out, the opposite is the case. Empirical results show that biodiversity, as an emergent property of communities is not a difference-making factor for invasions. This is because there have been many documented cases of highly diverse communities being invaded, yet communities with low diversity resisting invasion (Levine & D'Antonio, 1999). This means that the partitions which individuate community-level properties are not useful for studying invasions, and therefore Jarchow and Cook were right to choose the individual-level framework. Of course, the knowledge that biodiversity is not a difference-making factor can only come about after empirical investigation. It is often the case that this kind of knowledge is not available. Still, the point is that when this knowledge is available, it can be used to evaluate the aptness of target systems. Moreover, this knowledge can be used to cut across contextual frameworks making the criteria for evaluating target systems not merely pragmatic.

At this point, we have an answer to the question posed in the introduction to the chapter. The best framework to study the invasion of Indian Creek by *T. Angustifolia* is the individual-level framework. This framework provides the most useful partition, given our knowledge of invasion so far, one which individuates the trait of allelopathy as a difference-making factor for the invasion. This is an important achievement, which can be used to further refine our knowledge of invasions.

## 5. Conclusion

The ultimate goal for providing a theory of target system evaluation is not in order to identify the maximally apt target for each scientific purpose. This would not be possible due to our epistemic limitations. However, in many cases, identifying a sufficiently apt target is a goal that is well within our reach, as the evaluation of the target systems for invasion of Indian Creek demonstrates. This method of evaluation relies on two criteria, usefulness and relevance, which are distinct conceptually, yet closely related.

There is also a broader goal for the theory of target system evaluation, which is to identify the precise location of a problem in order to determine the best way to overcome it. Target systems are part of a more general network which makes up scientific practice, as they are represented by models or manipulated by experiments. When a model or experiment is successful, it is not so important to complete a full analysis of the target system. However, if a model or experiment does not meet its goal then the scientist must be able to determine the reason for this failure. In some cases the problem will lie within the model or experiment itself. For example, if a model contains idealizations which distort the target system too much, then the model will have to be adjusted. However, in many cases the problem lies with the specification of the target system. As I have shown, the theory of aptness can further distinguish between problems of partitioning and abstraction. This is important because, as we have seen, different problems require different solutions. The theory of aptness, while not infallible, provides us with tools which can provide significant assistance in such endeavors.

In the next chapter, I will examine the field of invasion biology in more detail. In doing so, I hope to show that the issues raised in my theory of target system evaluation are questions which are currently being raised by practicing scientists. In chapter 6, I will show that some invasion biologists are reconsidering the usefulness of the two main existing frameworks (individual-level and community-level

research), as they are capable of providing explanations for particular invasions, yet are not capable of providing successful predictions of invasions. I will then (in chapter 7) show how a closer examination of the target systems of individual-level and community-level research shows the different ways in which each is limited, and provides the space for the creation of a new, integrative framework for invasion biology.

## 6. Case Study

### Can there Be a Unified Theory of Invasion Biology?

#### 1. Introduction

Invasion biology is a relatively young discipline which is extremely important, interesting and is currently in turmoil. Invasive species become established outside their native ranges and displace existing inhabitants (Jeschke *et al.*, 2012). In doing so, they can “threaten global biodiversity, introduce diseases, cause other ecological problems, or incur economic costs” (*ibid* p. 2). Because of this, invasion research has captured the attention of the media and the public and has entered the political sphere (Davis, 2011). However, the field is also in turmoil as there are deep-rooted disagreements about the definition of invasive species (Richardson, 2011), the threat they pose to ecosystems (Lodge, 1993), the best way to manage invasions (Hulme 2009), and the conceptual framework which dictates the methods used to study invasions (Gurevitch *et al.* 2011, Jeschke *et al.* 2012, Moles *et al.* 2011). I will be focusing on this last point of disagreement, as it brings up interesting questions but has not yet been examined from a philosophical perspective.

The main issue in the controversy is whether the field invasion biology can be provided with an overarching unified theory. In order to understand why this is important for invasion biologists, we must look more closely at the history of the field (Davis, 2011). The history reveals that the discipline is very heterogeneous both with respect to conceptualizing the problem of invasions and with respect to the methodology used to study it. It also reveals that invasion itself is very heterogeneous, as invasions are caused by numerous and diverse factors. Moreover, these factors are so diverse that providing a unified theory of invasion based on the causes of invasion is extremely unlikely. A unified theory of invasion is possible, yet it must reconceive the problem of invasion. One way to achieve this is to focus on the process

of invasion, and classify particular invasions in terms of stages and barriers that need to be overcome in order for the invasion to pass to the next stage (Blackburn et al., 2011).

However, there is an additional problem, which relates to the prediction of invasions. Achieving successful predictions is perhaps the most important goal of invasion research, given the ecological and economic implications of invasions (Moles et al., 2011). The problem is that so far, few studies have produced successful predictions, while the ones that have are local, and focus on the causes of specific invasions. These studies are limited to small groups of invaders, and cannot be generalized to a greater theory of invasion. This means that the goals of unification and prediction must be pulled apart in invasion biology, as studies can either focus on providing a unifying theory of invasions, based on non-causal classifications of invasions, or they can focus on identifying the causes of particular invasions in order to provide predictions.

The point of this analysis is not to show that one of these goals is unimportant. Instead, it is meant to provide the conceptual space for both goals, yet to highlight that each must use different conceptualizations of invasion and different methodologies. This distinction is important because there is a view within invasion biology, which argues that the way to achieve successful predictions is to strive for greater theoretical unification of invasion. I will argue that causal accounts of invasion can achieve integration but not unification, and that this integration is the key to successful predictions.

## 2. A brief History of Invasive Species Research

The recognition that species outside their native range can cause problems for indigenous species dates back to Darwin (Reichard & White, 2003), who realized that in each habitat species have friends and enemies which affect their population's growth. Moreover, he saw that if a species was introduced to a new habitat without enemies, it could take over at the expense of the native species (ibid). After Darwin, there was little interest in the problems associated with invasive species, until the publication of Charles Elton's book *The Ecology of Invasions by Animals and Plants* in 1958 (Davis 2006). Elton thought that invasive species were exceedingly problematic and his main aim was to provide a framework for the conservation of native communities (ibid). Of particular interest is Elton's focus on the relative susceptibility of various

communities to invasion, as it eventually gave rise to one of the important conceptual frameworks of invasion research, ecosystem-level research.

The next landmark in invasion research was the first Biological Sciences Symposium held in Asilomar, California in 1964. The aim of the meeting was to understand the evolutionary significance of introducing species to a new habitat (ibid). The proceedings of the symposium, published as *The Genetics of Colonizing Species* (Baker and Stebbins 1965, in Davis 2006), reveals that the main focus of the meeting was on the invading species themselves and the traits that made them invasive (Davis 2006). This focus gave rise to another notable conceptual framework of invasion research, individual-level research.

Research continued largely separately within each conceptual framework until 1982, when the Scientific Committee on the Problems of the Environment (SCOPE) attempted to establish a single overarching framework for the discipline (Davis 2011). The committee identified three questions for invasion research: (1) What traits make species invasive? (2) Which communities are invaded? (3) How can we use the answers to (1) and (2) in order to manage invasions? The first two questions are a direct reflection of what had become the two main strands of invasive species research. Individual-level researchers were those who followed the approach set out by the Asilomar conference, and aimed to identify the traits of invaders as the causes of invasion. A notable example of this was Baker's account of the traits that make up the 'Ideal Weed' (Baker 1974). On the other hand, ecosystem-level researchers followed the ecosystem-level focus that originated in Elton's book. They sought the causes of successful invasions in the characteristics of invaded communities, arguing that irrespective of the invaders' traits some communities can resist invasions whereas others cannot. The factors that are most commonly referred to as affecting the susceptibility of a community to invasion are biodiversity (Levine & D'Antonio, 1999), biotic resistance (Shea & Chesson, 2002) and disturbance (Burke & Grime 1996, D'Antonio *et al* 1999), all of which are thought to be emergent properties of ecosystems.

Since the SCOPE publication, invasion research has made progress on the first two questions. However, there has been little -if any- progress made in answering the third question set out by SCOPE. Invasions remain exceedingly hard to predict and therefore difficult to avoid. What is troubling is that successful predictions elude scientists despite the advances that have been made in answering the first two questions. There is a reason for this. The results from studies and experiments tend to be specific to a particular system, and do not easily generalize to other systems. For example, traits that are thought to

promote invasiveness are context-specific (Pyšek and Richardson 2007). That is, a trait may predispose a plant to be invasive, yet only in a particular environment. Achieving a deep understanding of an invasion by a particular organism in a particular community, does not provide sufficient insights for predicting new invasions by similar organisms or invasions in similar communities. In fact, some invasion biologists believe that invasive species research is a science that is simply non predictive, much like earthquake science (Williamson 1999). Because of this, they think that the best we can hope for as a goal for invasive biology, are better explanations for particular invasions.

Despite this pessimism, there have been efforts to generalize results, mainly within taxa or geographical location. Most of this type of research has focused on plants and traits that universally make plants successful invaders. Notable examples include a series of papers on pine species (Rejmánek & Richardson, 1996) which found that compared to non-invasive species, invasive species had shorter juvenile periods, shorter intervals between seed crops and smaller seed mass. Other relatively general results come from the study of invasions on island ecosystems, as there is generally a higher rate of invasion in small habitat fragments and islands (Lonsdale, 1999). Yet these lower-level generalizations are quite limited. Most of them rely on idiosyncratic characteristics of a genus or geographical feature (Moles *et al*, 2012), while in some cases the generalization does not even apply to all genera within a taxon (Valentine *et al*, 2007).

In addition to the continuation of the research within the two major frameworks endorsed by SCOPE, a number of other lines of research into biological invasions have emerged, most notably propagule pressure, resource fluctuation, the enemy release hypothesis, evolution of increased competitive ability of invaders and the impact of invasive species on ecosystems (Davis, Grime, & Thompson, 2000; Gurevitch *et al*., 2011; Moles *et al*., 2011). Scientists usually advocate for one of these lines of research, trying to show how it is better than others in explaining the most biological invasions. There is a tendency of a group of scientists to publish a succession of papers within one line of research, which has led to increased specialization and fragmentation of invasive species research. This has led some scientists to believe that due to the diversity of approaches “researchers may overlook that they are studying different pieces of the same invasibility puzzle and that they should combine their findings in order to get a complete picture” (Milbau *et al* 2009).



In light of this, there has been a recently renewed effort to integrate the various frameworks of invasive species research, in an attempt to bring together many of the approaches in a fruitful way (Blackburn *et al.*, 2011; Catford, Jansson, & Nilsson, 2009). One of the first attempts to integrate hypotheses for invasion and bring them under a unified general framework appeared in a landmark paper by Richardson & Pyšek (2006), whose aim was to “merge the concepts of species invasiveness and community invasibility”. They kept the main framework provided by SCOPE, which divided invasive species research into individual-level and ecosystem-level approaches to invasions, but incorporated the other hypotheses into this framework. Thus, the tens rule, residence time, taxonomic affiliation, enemy release, phenotypic plasticity and long-distance dispersal come under the umbrella of trait-level research, while community invasibility is joined with propagule pressure, biotic resistance and invasional meltdown in the ecosystem-level camp. Perhaps the most important advance, however, was that two hypotheses, the theory of seed plant invasiveness, developed by Marcel Rejmánek (Rejmánek, 1996; 2000; Rejmánek *et al.*, 2005a; 2005b) and the theory of fluctuating resources (Davis *et al* 2000), were seen as theories that integrated the two lines of research identified by SCOPE, as each of them focuses on aspects of both invaders and communities, attempting to match particular traits which increase invasiveness with environmental conditions that facilitate invasion *because* of the particular invasive traits. This type of integration which focuses on pairing up traits and ecosystem characteristics, has since been adopted by a number of invasion biologists (Catford *et al* 2009, Davies & Sheley 2007, Shea & Chesson 2002).

Still, this attempt at unification was still quite limited, as it was meant as a first sketch at unification, and as it was only meant to apply to plants. Since the publication of that paper, there have been many more attempts to provide invasion biology with a unified conceptual framework. According to Gurevitch *et al* (2011) “Invasion biology is clearly ripe for conceptual synthesis and integration, by subsuming these individual hypotheses in a broadly applicable conceptual framework grounded in basic principles of ecology and evolutionary biology”. A number of invasion biologists have voiced the worry that the lack of integration between frameworks can inhibit progress in invasion research. For example, Milbau and colleagues (2009) lament the scarcity of interdisciplinary collaboration, despite the “knowledge that combining studies encompassing different factors and scales usually largely improves our understanding of ecological processes”. Some researchers believe that within the last few years, progress has been made towards developing a general, synthetic conceptual framework for invasions (Gurevitch *et al*

2011). Nonetheless, many agree that there is a lot more work to be done if the integration is to be successful and yield fruitful results (Moles et al 2012).

### 3. Two Conceptual Frameworks

Before turning to the analysis of unification and integration in invasion biology, I will give a more detailed analysis of the two conceptual frameworks outlined by SCOPE. This is necessary for gaining a detailed understanding of the background concepts and methodologies that drive different types of invasion research. It also serves to highlight the extent of diversity between different approaches to invasion.

An interesting fact about invasions is that the actual fraction of plant species that spread and impact native populations is rather small (about 1%) (Weidenhamer & Callaway, 2010). In fact, most introductions of species to a community do not result in a successful invasion. Even successful invasions are usually not the first introduction of a particular species to a particular community (Sax and Brown 2000). Facts like these have motivated some scientists to look for the causes of invasion in the ability of a community to resist invasion. As most invasions are resisted by native communities, these scientists believe that understanding the factors that affect resistance are the key to understanding invasions.

I call this approach ecosystem-level research, as it gets its name and many of its concepts and methods from ecosystem ecology. The framework for ecosystem-level research was set by Elton, whose treatment of invasion focused on the effects it had on native communities (Davis 2006). Elton saw invaders as threats to native ecosystems (ibid). An invader destabilizes an ecosystem, resulting in the loss of native species. Because of the strong connections between members of an ecosystem, the loss of some species often results in the loss of more and more species. The result is often catastrophic. The job of the ecosystem ecologist or conservationist is to protect these native ecosystems and minimize any type of disturbance to them.

There are numerous versions of ecosystem ecology, which view ecosystems slightly differently. In the most extreme version, ecosystems are thought of as individual entities (Odenbaugh 2007). Here the differences between biotic and abiotic components of the system are broken down and the system is studied in terms of the flow of energy through it. For example, a food web is not conceptualized in terms of the organisms or populations that constitute it, but in terms of the transfer of energy from one level of the web

to another. As energy moves up the levels of the web, it is lost. In a less extreme version, communities and ecosystems are thought of as wholes, with emergent properties that their parts do not possess (de Laplante & Odenbaugh 2006). These researchers think that it is meaningful to speak of properties such as “stability” that often occur in a system, which no individual part of the system possesses. At the same time, they may attribute causes of balanced states or equilibria to individuals or populations in the system.

Ecosystem-level research incorporates many different methods for studying invasions. Biodiversity studies use models and experiments which attempt to find lower levels of invasion in species-rich environments (Kennedy et al., 2002). There are also related experiments that focus on mechanisms of biotic resistance to invasion. These often take the form of experiments on island ecosystems, as continents typically have higher biodiversity than islands, because they are geographically and evolutionarily isolated (Jeschke et al. 2012). More recently, niche-based ecological models such as GARP and MAXENT are used to identify similarities between ecosystems in different parts of the world and use that to see whether they might be susceptible to the same invader (Sobek-Swant *et al*, 2012).

The second conceptual framework is individual-level research, and views invaders as individual competitors. Researchers look for the traits that make an organism a successful invader. Traits can contribute to invasion success by helping an organism to disperse, by helping it to utilize resources and by conferring a competitive advantage so that it can become established and spread at the expense of native populations. As in the case of ecosystem-level research, individual-level research has links to general ecology, more specifically to population ecology. This sub-discipline studies the dynamics that emerge from interactions between organisms. One of the most important population interactions is the competition within and between populations for resources (Cooper, 1993). A successful invader can be thought of as a good competitor. For example, a plant which is allelopathic<sup>1</sup> can outcompete other plants and successfully invade a new environment. In population ecology and individual-level research, the main unit of analysis is the individual, as the focus is on characteristics which can potentially make an individual a successful invader.

The majority of individual-level research takes the form of experiments which aim to find the causes an invaders competitive advantage over native species. For example, many experiments on plants

---

<sup>1</sup> Allelopathy is the release of toxic compounds from a plant which results in the suppression of neighboring plant growth (Fitter 2003). It confers a competitive advantage to the invader by inhibiting the reproduction and growth of native populations.

identify traits that affect their growth and reproduction. Increased relative growth gives plants a competitive advantage because they are faster to reach the stage of optimal resource uptake. This means that they are capable of utilizing a greater fraction of the limited resources, which gives them an added growth boost but also means that the native species' growth is restricted. Faster reproduction rates increase competitive advantage simply because there are more invasive propagules than native ones. As a population, the invaders then utilize a greater fraction of resources.

Other experiments have focused on phenotypic plasticity and its relation to invasions. An organism exhibits phenotypic plasticity when it can express different phenotypes (i.e. when it can change its chemistry, physiology, development, morphology or behavior) depending on changes in the biotic and abiotic environment (Agrawal, 2001). Phenotypic plasticity has been thought to facilitate biological invasions in communities with more specialized native species (Hulme, 2007). Individual-level research also utilizes a variety of models. These include growth models which examine the connection between density dependent growth and invasion (Taylor & Hastings, 2005), and individual based models, which aim to determine the potential spread of an invader in a new environment (Higgins, Richardson, & Cowling, 1996).

In the next section I will address the notion of synthesis, used by invasion biologists and determine its connection to the notions of unification and integration, as they are used in philosophy of science.

#### 4. What is a Synthesis?

The discussion in the previous section makes it clear that invasion biologists are calling out for synthesis in their discipline, as they believe that this will help the discipline move forwards. From a philosophical standpoint, there are two important issues which need examination. The first is to understand what exactly is meant by synthesis or integration and the second is to determine how exactly synthesis can help improve the discipline.

There are three terms used in invasion biology to denote union within the discipline: 'synthesis', 'unification' and 'integration'. In many cases these three terms are used interchangeably, yet there are some distinctions which should be made, which may prove useful in understanding the motivation and aims of the scientists. In the philosophical literature, *unification* is understood as the provision of a single theoretical framework for a field or a discipline (Mitchell, 2003). It can be a 'regulative ideal' for theories

or explanations (Kitcher, 1999), as theories aim to unify by being able to explain a wide range of phenomena. Important unificatory accounts of scientific explanation include reductive accounts of science, where phenomena at one level are explained by mechanisms at a lower level (such as Nagel 1949), the ‘covering law model’ of scientific explanation, where phenomena were viewed as instances of general patterns explained by laws (Hempel and Oppenheim, 1948), and Kitcher’s unification account where a unifying theory “advances our understanding of nature by showing us how to derive descriptions of many phenomena, using the same pattern of derivation again and again, and in demonstrating this, it teaches us how to reduce the number of facts we have to accept as ultimate” (Kitcher, 1981).


In biology, the notion of *integration* started being used as an alternative to the traditional idea of reduction (Brigandt, 2013b). According to Plutynski (2013), who put together a list of various uses of integration in biology, integration is usually attributed to particular research programs which aim to solve a particular problem and usually involves some form of collaboration between fields or research programs. It often takes the form of merging together *results* of particular studies, and ““integrative explanations are frequently local, that is, they often involve establishment of epistemic and ontological connections between two specific areas of research” (p. 468). The main contrast to unification is that integration is usually attributed to particular research programs which aim to solve a particular problem, whereas unification is associated with large scale explanations, theories and models. This focus on problem-solving is what brings together researchers, methods and materials from different disciplines and sub-disciplines, and often results in the creation a ‘meta-methodology’ for solving the particular problem (Plutynski, 2013, O’Malley, 2013).

Despite the alleged lack of collaboration in invasion biology (Milbau *et al.* 2009), integration is already present. For example, scientists have attempted to integrate results from studies of invaders’ traits with results from studies of invasibility of communities (such as Richardson and Pyšek 2006). The idea driving the integration is that pairing up particular traits which are advantageous with environments that have certain characteristics will help us identify the types traits and communities that have a higher risk of invasion. The problem with this method is that it is considered to have been unsuccessful so far. The few studies and meta-analyses that have attempted this sort of integration have usually produced long lists of traits and characteristics that have been previously associated with invasions, together with recommendations that both traits and community characteristics should be taken into account in future research.

For example, the paper by Richardson and Pyšek (2006) is titled “Plant invasions: merging the concepts of species invasiveness and community invasibility”. However, the paper is divided into fourteen sections, with the majority of those sections providing an overview of recent research in different approaches to invasion. There are sections on taxonomic patterns, the naturalization hypothesis, phenotypic plasticity, long distance dispersal, plant-seed invasiveness, propagule pressure, diversity, fluctuating resources and mutualisms. The merger is shown in a figure which classifies the various approaches into one schema (shown in figure 6.1).

The idea is that the theory of seed plant invasiveness and the theory of fluctuating sources are the two theories which successfully merge the two frameworks. Seed plant invasiveness merges particular traits with particular environments in which they are likely to invade. The general idea is that “widespread species are more likely to be dispersed because they occur in more locations and have higher chances to be dispersed, and they are more likely to be adapted to wider range of conditions” (ibid p. 419). However, invasibility does not really seem to be part of the process. Environments which have already been invaded are included in the model, yet this does not make any mention to their initial invasibility. There is also mention of disturbed habitats, yet it is not clear how they factor into the model. In addition, there are many traits that have been identified as potentially causing invasiveness, including small seed mass, small size, large size, flowering early, flowering late, non-dormancy or long dormancy.

Figure 6.1. Merger of invasion studies (reprinted from Richardson and Pyšek 2006)



The case for fluctuating resources is more promising. The idea is that “invading species must have access to available resources, e.g. light, nutrients, and water, and that an invading species will be more successful at invading a community if it does not encounter intense competition for these resources from resident species.” (ibid p. 423). Again, however, the generalizability of the framework is limited, as there are many different resources available in an environment, and many different resources that an invader may need. There are also many different cycles of resource fluctuation, which make the process even more indeterminate. In short, as it stands, this theory is still at an early stage and cannot do much more than prove a long list of potential environments which could be paired up to potential invaders at some stages of the resource cycles.

In response, there has been an increase calls for unification in invasion biology, rather than integration. For example, Gurevitch *et al.* (2011) argue that the way forward is to provide a general, synthetic conceptual framework for invasion biology. They define conceptual frameworks as theoretical

entities which “encompass the assumptions, laws and ideas that underlie the construction of a broad concept”. The point of such a framework is to generate new hypotheses, evaluate existing ones and lead to the development of models and experiments. Importantly, frameworks are unificatory, insofar as they “define connections and elements of knowledge in a general area of inquiry, giving coherence and direction to the study of empirical problems”.

According to Gurevitch and colleagues, a successful conceptual framework for invasion biology should first and foremost incorporate invasion research into the general context of evolution and ecology. In other words, invasion research should not be conducted in isolation from general biology, because the processes that give rise to invasions, such as population interactions (competition, predation), ecosystem processes, community and landscape ecology, and demographic factors, are all aspects of general biology. In addition, they believe that the conceptual framework should take into account time lags and effects of scale: “a unified approach must in some way accommodate a consideration of how spatial and temporal heterogeneity facilitate invasion, and the role of stochastic variation in space and time”. Finally, they explicitly state that the framework should focus on both individual and ecosystem-level processes (as defined by SCOPE), by facilitating explanations of “the success of invasives relative to natives in the novel environment, the success of invasives in the novel environment contrasted with their home environment and the success of invasive species or populations relative to non-invasive aliens in the novel environment”.

In the next section I will outline three examples of synthesis in invasion biology. In my analysis, I will use the term synthesis as a neutral term, to denote any account that the scientists themselves present as synthetic. I will argue that only one of the three kinds of synthesis is an example of unification.

## 5. Three Kinds of Synthesis

Since the SCOPE publication and the later landmark paper by Richardson & Pyšek (2006), there have been many attempts to integrate results of invasion studies in a general framework, with the hope of finding the key to predicting invasions. These typically take the form of theoretical or meta-analytic articles attempting to define a synthetic conceptual framework for invasion biology. In this section, I will take a



closer look at three different attempts to provide a synthetic framework.<sup>2</sup> The first synthesis merges aspects of trait-level, ecosystem-level research and other other causes and processes of invasion (merger synthesis). In this sense, it can be seen as a continuation of the type of research outlined by SCOPE. The other two syntheses are different. There are some scientists who believe that the SCOPE framework is not the right way to synthesize results (Romanuk et al., 2009). In light of this, some invasion biologists have set out to construct integrative or synthetic frameworks that might fare better in questions of invasion prediction and management. The second framework synthesizes invasions hierarchically (hierarchical synthesis), while the third focuses on the process of invasions instead of their potential causes (process synthesis).

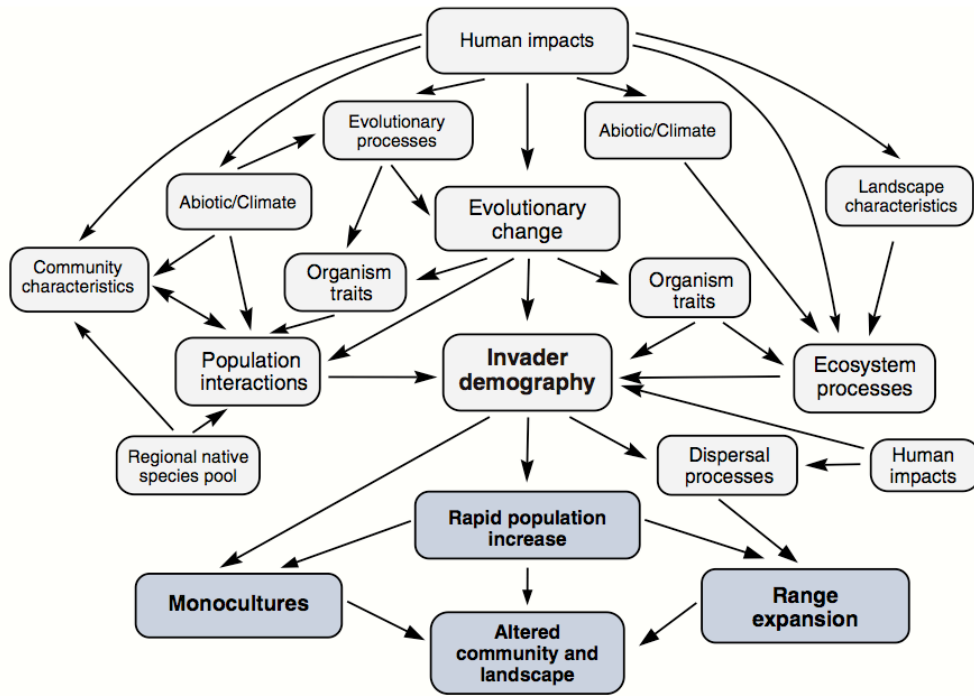
### 5.1. The *merger synthesis*.

After spelling out their notion of a conceptual framework for invasion biology, outlined in the previous section, Gurevitch and colleagues (2011) put forward their own synthetic framework, which satisfies the conditions they set out. They call their approach a ‘Synthetic meta-framework’ (SIM), as it synthesizes pre-existing conceptual frameworks. The SIM is constructed by first identifying the causes and process of invasions and then determining the causal interactions between them (see Figure 6.2). Any specific invasion, invasion hypothesis and prediction about invasion will be explained by some part of the SIM, not all of it. For example, the authors consider a hypothetical case where the invasion is caused by enemy escape. In this example, a invading plant is subject to reduced predation on its seeds and seedlings (population interactions). This leads to an increase in survival at particular life history stages (invader demography), which can then result in rapid population increase. Finally, “if conditions are favourable for establishment beyond the extent of the local population, and means exist for dispersal, this may ultimately lead to range expansion and in turn alter landscape characteristics”. This invasion is explained by invoking a particular set 5 causes/processes out of the 17 possible ones.

---

<sup>2</sup> These are by no means the only papers which provide synthetic frameworks for invasion biology. Yet I have chosen to focus on them because they are representatives of the three most common types of synthesis. The analysis of each type of synthesis can also be applied to other examples of synthetic frameworks.

**Figure 6.2. The SIM** (reprinted from Gurevitch *et al.* 2011)



According to the authors, the important point is the SIM provides a more in-depth understanding of invasions. For example, ‘enemy release’ is a recognized cause of invasion, yet the authors believe that it is a black box explanation. With the SIM, the box can be opened and broken down, so that a more accurate picture of the invasion emerges. Moreover, the SIM is a flexible framework which can be used in many different ways to illuminate biological invasions. One of the most important advantages is that it allows for the comparison of invasions. More specifically, it can be used to “compare different specific invasions in a particular system, or general cases of invasion processes in different systems (e.g. terrestrial plant communities and invasions by freshwater invertebrates) to highlight their essential similarities and differences” (p. 141). A second advantage is that it allows for the formulation of new hypotheses, as it facilitates the combination of processes and causes and the discovery of new connections between them.

The SIM is considered to be synthetic in two ways. First, because it merges a number of causes and processes into a single framework. For example, ‘organism traits’ and ‘evolutionary processes’ were traditionally thought of as having their own distinct frameworks, yet here they appear together, with causal

arrows between them. The second reason is that the framework also incorporates causes and processes from general ecology and evolution, such as ‘evolutionary change’ and ‘rapid population increase’. According to the authors, the SIM is a reconceptualization of invasion biology “in terms of basic ecology and evolution, rather than in terms of special processes and factors unique to invasion biology”.

This framework aims to be synthetic in the top-down unification sense, outlined in the previous section. The focus on ‘reconceptualization’ of invasion biology and attempting to subsume individual instances of invasions under the general theoretical framework, both point to this. However, before we accept this as a true synthesis, we should look more closely at the SIM and the way in which it is used to explain invasions.

The SIM is a merger of a number of different hypotheses, such as organism traits, demography, abiotic factors, ecosystem processes and so on. Some of the hypotheses come from trait-level research, others come from ecosystem-level research, while other still come from different areas of invasion research. Still, the result is a complex framework which reads like a list of most of the known causes of invasion. Consider again the example of enemy release mentioned above. Despite the claims of the authors, there are numerous existing studies which explain ‘enemy release’ in great detail. For example, (Bever, 2003) provides an extensive analysis of positive and negative feedback in plant communities which gives precise a explanation necessary for different types of feedback to occur. This study is firmly situated within the competition framework as it uses an amended version of the Lotka-Volterra model to determine the effects of competition. At the same time, there are numerous biogeographical studies (which fall in the ecosystem-level camp), which show that soil biota generally have more positive effects on invaders than on native species (Klironomos 2002).

What additional information does the SIM provide? The SIM does make the case that population interactions, demographical factors, population increase, range expansion and the altered community landscape must be taken into account. Yet these are factors which are already known to affect invasions, as they appear in studies like the ones of Bever (2003) and Klironomos (2002). In other words, the SIM simply links together factors which potentially cause invasions.

However, linking together is an integrative rather than a unificatory process. In order to count as a true unificatory explanation, a theory must provide “one (or more generally, a few) pattern(s) of argument which can be used in the derivation of a large number of sentences which we accept” (Kitcher 1981). The

SIM cannot be said give a pattern of argument which reduces the number of sentences which we accept, because it simply provides disjunctive set of potential causes of invasions. Moreover, the SIM does not explain why or how a particular subset causes an invasion, nor does it give us a general pattern of invasion. Therefore, a better way to understand this synthesis is as an example of integration.

## 5.2. The *hierarchical synthesis*

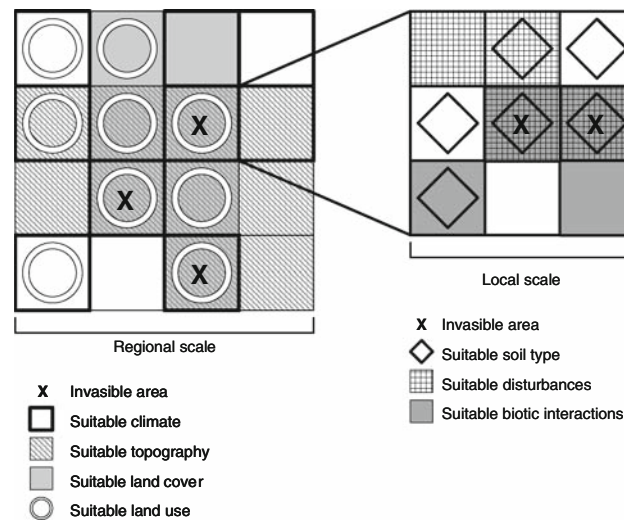
The next approach is the hierarchical synthesis which attempts to unify invasions under a theory based on scale. As stated above, invasions and methods of studying them are extremely diverse. Ann Milbau and colleagues (2009) identified no less than six factors distinguishing invasibility studies. Invasion studies can differ in terms of *scale* (micro, local, regional, continental), *factors studied* (abiotic or biotic), *stage of invasion* (colonization, establishment, spread), *measure of success* (e.g. germination of seeds, alien species richness, percentage cover of a particular introduced species), *status of invader* (native versus non-native) and *applicability* (explaining a case study versus looking for generalizations). This diversity makes the creation of a synthetic framework quite difficult. However, the authors identified scale as being more important than the other factors, but also as a factor which has the potential to unify all the others in a single hierarchical framework.

The authors saw correlations between particular types of factors studied and stages of invasion with particular measures of scale. In light of this, created a hierarchical framework which determines the invasibility of a particular environment. In this framework, “factors operating at a smaller spatial scale are subordinate to factors operating at a larger scale, and only if conditions at higher levels are satisfied, factors at a more local scale may become significant”. In other words, if a species is introduced to a new area, the most important barrier to its survival are the climatic conditions. If they are favorable, then other, smaller scale factors such as topography, “including topography-related effects on climate, such as rain shadows and temperature changes with elevation”, start to become important. Next, a species may be restricted by *local* factors such as soil type or biotic interactions, and so on.

At the same time, however, the smaller-scale factors are “usually essential for precise predictions at more local scales”. In other words, depending on the accuracy of the predictions required, a different level of the scale will be applicable. “If one is interested in the potential spread of a species at the landscape scale, only data on climate, topography, land cover and land use are needed. However, if one is interested in

the invasibility of a particular site, both factors affecting invasibility at site-scale and at all higher scale domains should be considered.” In this way, pre-existing invasion studies can be categorized depending on the scale of the study, either because they explicitly mention the scale at which they operate, or indirectly, by inferring the scale from other factors, such as the type of factors studied, the measure of success or the intended applicability of the study.

**Figure 3.** The hierarchical framework (reprinted from Milbau *et al.* 2009)



For example, the herbaceous summer-annual species *Impatiens glandulifera* is native to the Himalayas and is now present across Europe and North America. It has been studied extensively and research has shown that it is sensitive to frost. The hierarchical framework can therefore show that “regional climate data can be used to identify countries or regions where the species can potentially invade” (ibid p. 947). Also, “in regions where the climatic requirements are met, one can indicate waterways as the most vulnerable areas for invasion, since *I. glandulifera* is closely confined to riparian habitats” (ibid). Finally, “Because there is a strong positive correlation between total plant biomass of *I. glandulifera* and irradiance (Andrews *et al.* 2005), micro-climatic variables may be used to predict its invasion success at the smallest spatial scale” (ibid).

The new framework is synthetic, because it incorporates factors already mentioned in other invasion studies, yet re-conceives of invasions in a new, hierarchical manner. This type of synthesis differs

from the merger synthesis in that it is a departure from the framework set out by SCOPE. It does not aim to merge pre-existing frameworks, but identifies one factor of invasions (scale) and builds a framework based on that. At first glance, this framework seems more unificatory than the merger framework, as it aims to unify all invasions by classifying them in terms of scale. Re-conceptualizing in terms of scale seems like a top-down approach, which would count as unification instead of mere integration. The unificatory power of the framework is that the reconceptualization in terms of scale adds to the explanatory power of the framework. In other words, we understand more about invasions by classifying them in terms of scale: we learn that different factors operate at different levels of scale.

However, I think that this interpretation of the framework is slightly overstated. While scale is used to *classify* particular invasions, the authors still rely on the *factors* operating within each scale to explain and predict invasions. In other words, scale is used as an organizing concept, yet the each invasion is explained by the causes operating at a particular scale. For example, figure 2 shows that at a regional scale, invasions can be caused by climatic and topographical factors, whereas at the local scale they can be caused by soil type, disturbances and social interactions. Yet this classification is too coarse-grained to explain a particular invasion or predict an invasion in the future. In order to do that, a scientist must identify exactly which factor combination of factors caused an invasion. For example, if a scientist is interested in explaining local invasion, it is not sufficient to show that the invasion is local and not regional. She must identify whether the invasion is caused by suitable soil type, a particular kind of disturbance, particular soil interactions or some combination of these factors. In other words, even if the reconceptualization in terms of scale does increase our general understanding of invasions, we must still revert to examining the factors at each scale in order to truly understand and be able to predict particular invasions.

The ultimate reliance on the factors which cause an invasion for explanation and prediction is why this framework should be understood as integrative rather than unificatory. Moreover, in order to be unificatory, based on my previous definition of unification, the framework should reduce the number of factors which explain invasions. However, this framework adds the factor of scale to the framework without reducing the number of other factors. Thus, even though the addition of scale might add something to our overall understanding of invasions, it does so by creating an additional factor, yet it does not subtract any of the pre-existing factors.

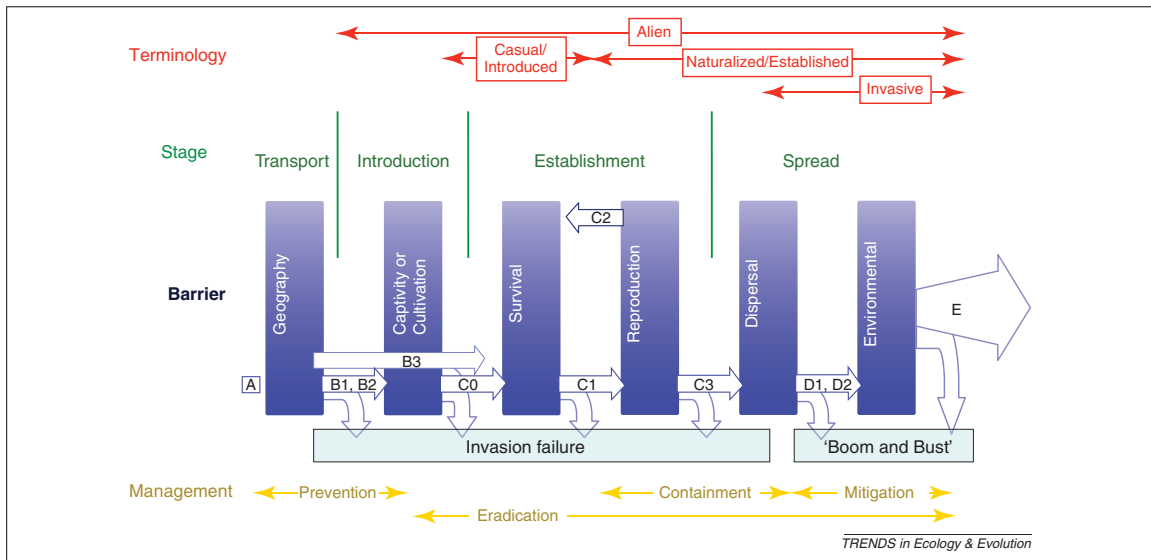
### 5.3 The *process synthesis*

The third synthesis is quite different to the previous two syntheses, in that it focuses on the process by which organisms invade ecosystems, instead of the causes for invasion. In other words, the point is not to determine why a particular organism becomes an invader but to explain how it does so. In the framework proposed by Blackburn et al. (2011), invasion is a process that can be divided into a series of stages (transport, introduction, establishment and spread) with barriers at each stage (geographical, captivity/cultivation, survival, reproduction, dispersal and environmental) that need to be overcome for a species or population to pass on to the next stage see figure 4. If, on the other hand, an invader is stopped by a barrier at any stage, then the invasion is considered a failure. For example, the budgerigar *Melopsittacus undulatus*, an Australian native is considered an introduced species (in the UK), as there are individuals released into the wild (i.e. outside of captivity or cultivation), which are nonetheless incapable of surviving for a significant period. In this framework the budgerigar has passed the *transport* and *introduction* stages, overcoming the geography barrier (by traveling to the UK, with some help from humans) and the captivity barrier (by outsmarting said humans and escaping captivity) and is about to enter the *establishment* stage. Unless the population overcomes the survival and reproduction barriers (by surviving long enough to reproduce in sufficient numbers), it will not pass into the next stage and the invasion will fail.

This framework is considered to be synthetic because it merges two previous frameworks. According to the authors, the two main frameworks of invasion biology are the 'Richardson framework', where invasions are seen as a "series of barriers that a species negotiates to become either naturalised or invasive" and the 'Williamson framework', which views invasions as "a series of stages that a species must pass through on the pathway from native to invasive alien". The authors also state that the former framework is adopted by plant ecologists while the latter is favored by animal ecologists. Blackburn *et al.* create an overarching framework which combines elements of both organism-focused stages and ecosystem-focused barriers. The main advantage of this framework, according to its proponents, is that all other invasion studies can be accommodated within it, even if they do not focus explicitly on a particular stage of the invasion process. This is because "different parts of this framework emphasize views of invasions that focus on individual, population, process, or species". For example, a study that identifies small seed size as a cause of plant invasiveness would be placed at the *transport* stage of invasion, as small seed size is a trait

that enables plants to disperse over wide areas (usually with the aid of wind or birds). The framework also affords an organized way to incorporate the various terms by which invasive species are known, as “species are referred to by different terms in the terminology depending on where in the invasion process they have reached”.

**Figure 4.** The unified framework (reprinted from (Blackburn et al., 2011))



Unlike the two previous synthetic frameworks, it seems that this framework is truly unificatory. It re-conceptualizes invasions in terms of their process rather than their causes. In fact, it does more than that, as it takes two already general frameworks and creates an overarching framework which combines them. Yet it does not simply add causal factors to an overarching conceptual framework, as it does not aim to explain particular invasions in terms of their causes. Factors that might explain a particular invasion, such as small seed size are not used to explain a particular invasion, but are used to identify the stage of the invasion. This means that these factors, are only used to classify invasions, but not used to explain them. In other words, when a study of a particular invasion identifies the factors that caused it, this framework uses the factors as proxies to determine what stage the invasion is at and uses those, in turn, to determine the type of barrier it will have to overcome next.



I will illustrate this point by contrasting it to the hierarchical synthesis mentioned above. The hierarchical synthesis aimed at both classification and explanation of particular invasions, as it identifies the scale of a particular invasion and the factors which explain the invasion at that scale. This means that it ultimately relies on the causes of the invasion to explain it. In contrast, the process synthesis does not aim to explain particular invasions, but is primarily classificatory. The classification is top-down, as it relies on the general overarching framework and situates a particular invasion within it. It does not rely on lower-level causes to explain the invasion. In this sense is it truly unificatory and not simply integrative.

Yet this framework is not merely classificatory. It does increase our understanding of invasions, albeit not in terms of their causes. That is, it does not explain why an invader might pass on to the next stage, but explains how a barrier can be overcome. For example, the framework does not aim to identify the reason why the budgerigar *Melopsittacus undulatus* might progress to the establishment stage, but does explain the type of barrier which needs to be overcome and how the organism can overcome it. Moreover, the framework is explanatory in the sense that it identifies the strategy for managing invasions at each stage. In the case of *Melopsittacus undulatus*, which is about to enter the establishment stage, this means that it is already too late to use the *prevention* strategy to manage the invasion, yet it is still possible to aim for complete *eradication*. The invasion is not already at the stage where the only possible management is *containment*.

This leads to another interesting difference between this synthesis and the previous two syntheses. In addition to explanation of invasions, the first two syntheses aim at the prediction of future invasions. In contrast, the the authors of the process synthesis seem to substitute talk of prediction with management. I surmise this from the fact that there is absolutely no mention of prediction in the paper, whereas there is a whole section devoted to managing invasions. Even though it might be possible to make predictions using the process synthesis, it seems that the authors are not interested in pointing them out. I suspect that this is largely because these types of predictions might not be particularly interesting. For instance, one could predict what the next stage of the invasion will be if a barrier is overcome, yet this sort of prediction is not very informative. If the framework gave us the tools to predict whether or not the invasion will pass onto the next stage or not, that would a very useful kind of prediction, one which the authors would presumably take care to point out.

In contrast, the previous two syntheses are supposed to be used to make more successful predictions of invasions. For example, Gurevitch et al. (2011) state that without an “adequate conceptual framework” such as the SIM “important aspects of invasion biology can be overlooked and predictions can be incomplete, ambiguous, inaccurate or misleading.” (p.417). In the case of Milbau et al (2009) the importance of predictions is even more clearly stated. They frame the problem of invasion in terms of predictability and offer their framework as a solution to that problem: “Although the complexity of natural systems presents fundamental limits to predictions, we think this framework can provide a useful tool for the identification of areas of risk for biological invasions” (p. 491).

In fact, it is quite uncommon to find a paper which does not focus on predicting invasions. It seems that the received view of science in invasion biology is that unified theories give better explanations of more phenomena and better predictions of the same phenomena. I will examine this view and its implications for invasion biology in the next section.

## 6. Symmetry Between Explanation and Prediction?

The idea that unified theories have both explanatory and predictive power, is not a new one. It was the received view of science for much of the 20<sup>th</sup> century, adopted by scientists and philosophers alike (Salmon, 2006). The view was presented clearly in Hempel and Oppenheim’s account of scientific explanation and named the thesis of ‘symmetry between explanation and prediction’ (Hempel & Oppenheim, 1948). On their view prediction and explanation have the same logical structure. In both cases, an explanation is composed of an explanandum (that which is being explained) and an explanans (that which is explaining the explanandum). The explanans is composed of laws and antecedent conditions, which logically entail the explanandum. If the explanandum takes place before the explanation has occurred, it is an explanation, whereas if it takes place after the explanation has occurred, it is a prediction. For example, an explanans could be made up of the laws of planetary motion, and auxiliary conditions concerning the location of two planets at a particular point in time. These could be used to explain an eclipse (explanandum) which has already occurred, or predict an eclipse in the future. According to Hempel and Oppenheim, an explanation is not fully adequate unless it could have served as the basis of a prediction of the phenomenon in question (Hempel & Oppenheim, 1948).

Many invasion biologists also take up the idea of symmetry between explanation and prediction. Given the progress that has been made in the explanation of particular invasions, they thought it only natural that good predictions would follow (Lodge, 1993). The way to go about making better predictions was to look for more general and more unified explanations. For example, Moles *et al.* (2012) state that better predictions will come about by incorporating different types of information. By so doing, they “hope that it will be possible to cut through the apparent idiosyncrasies and predict the circumstances under which species and ecosystems will respond in different ways.” (p. 120). In other words, they think that generalizing processes and results will help identify the general causes of invasions, which will increase the understanding of invasions and yield better predictions. This, they believe, will represent “a major advance” in invasion biology (*ibid*).

Nonetheless, there are reasons to think that, at least in invasion biology, more unified explanations will not yield better predictions. While it is too early to tell whether the synthetic frameworks presented in the previous section will yield successful predictions, I expect that they will not fare any better than previous frameworks. The crux of the problem for invasion biology, is that invasions themselves are extremely heterogeneous. More specifically, the causes of invasions are extremely heterogeneous. For instance, just within plant taxa, causes of invasion include climate, topography, land cover, land use, soil type, disturbance, resources, micro-climate, mutualism, competition, facilitation and pathogens (Milbau *et al.* 2009). According to Romanuk *et al* (2009), the Global Invasive Species Database (<http://www.issg.org>) suggests that there are no natural rules that govern the processes of invasion, which have any real predictive value (Bright 1998 in Romanuk *et al* 2009). This is because any “generalizations about invaders over too wide a taxonomic range, such as all species, or all insects, or all angiosperms, invariably (lead to) too many exceptions to be useful” (Williamson 1996 in Romanuk *et al* 2009).

From a philosophical standpoint, the invasion biologists’ adherence to the symmetry thesis seems a bit puzzling. While Hempel and Oppenheim’s view of scientific explanation gained some popularity it then came under heavy attack. The general consensus seems to be that the covering law model of explanation as a whole, is not an adequate account of scientific explanation (cite Kitcher, Strevens, Scheffler, Scriven). Much of the criticism came in the form of counterexamples, many of which have now become famous in their own right (Salmon 1992). The counterexamples were meant to show that there are cases of explanations which cannot be construed as predictions and cases where a successful prediction

cannot be constructed as an explanation. For example, we can explain why a patient who suffers from paresis, given that they had syphilis, yet we cannot predict that a person suffering from syphilis will develop paresis (as syphilis is a precondition for paresis, yet only a small fraction of those afflicted with syphilis develop paresis) (Scriven, in (Salmon 1999)). On the other hand, the fall in a barometer predicts a storm, yet it cannot count as an explanation of the storm (ibid).

It is important to note that the invasion biologists mentioned above do not necessarily espouse the symmetry thesis *per se*. That is, I do not think that they view explanation and prediction as symmetrical, in the Hempelian sense. Rather, the view is that better and larger scale explanations will yield better predictions. This view is implied by the symmetry thesis, yet it does not constitute the symmetry thesis. Moreover, the examples which are used to argue against the symmetry thesis cannot be used as arguments against the view that there is *some* relationship between explanation and prediction, albeit not one of symmetry.

The point of this discussion is to show that focusing on causes will not yield unification in invasion biology. On the other hand, when unification is achieved, as in the case of the process synthesis, this happens at the expense of prediction. This shows that, at the very least, the idea that larger scale explanations are necessary for more successful prediction should be abandoned. Nonetheless, this does not mean that any hope of successful predictions should also be abandoned. In the next section I will show that integration and mid-level generalizations are possible and have been shown to yield successful predictions.

## 7. The value of Integration and Mid-level Generalizations

Given the recent philosophical literature on the importance and value of integration, there is good reason to investigate the role of integration in invasion biology. In the philosophical literature, integration is seen as a an acceptable middle ground between unificatory but reductionist accounts of science and accounts of the disunity of science, such as Dupré (1993) (Brigandt 2013). Moreover, the fact that integration is problem-oriented instead of theory oriented is seen as a value which unification does not share (Plutynski 2013).

At a theoretical level, we have good reason to expect integration in the field of invasion biology, as it exhibits many of the hallmarks of integration cited by Plutynski (2013). For many years, invasion biology was also thought of as a *problem-oriented* field, as the focus of invasion studies were the prediction and

management of invasions. Also, even though the problem of biological invasions is global, invasions are often studied *locally*. For example, even large scale ecological niche models such as GARP and MAXENT, which are sometimes used to predict invasions, rely on local environmental variables such as temperature and precipitation (Sobek-Swant *et al.* 2012). With a few exceptions, invasion research usually focuses on identifying the *mechanism* of an invasion, such as the factors which increase the competitiveness of the invader or the invasibility of the community. There is also a history of *collaboration* in the field, with researchers from different backgrounds contributing data, methodology and techniques particular studies of invasion. Even though some invasion biologists lament the fact that there is little collaboration in the field (Milbau *et al.* 2009, Davis *et al.* 2001), there are numerous examples of collaborative projects in invasion biology (Gurevitch *et al.* 2011 Rejmánek & Richardson 1996; Richardson & Rejmánek 2004; Richardson and Pysěk 2006; Milbau *et al.* 2009; Moles *et al.* 2012).

Many of these collaborations have been very successful. Scientists now have a much deeper and more extensive understanding of the mechanisms of invasion and its management. More importantly, some of these collaborations have also produced successful predictions. One of the most important of these collaborations is a series of papers by David Richardson and Marcel Rejmánek, on conifers (Rejmánek & Richardson 1996; Richardson & Rejmánek 2004). Richardson and Rejmánek originally studied invasive and non-invasive species of pines and showed that invaders shared three important traits: small seed mass, small juvenile periods and short intervals between seed crops (Rejmánek & Richardson 1996). Plants which exhibit these traits have a high “index of invasiveness” (number of invasive species/number of rare/threatened species) (*ibid*). Small seed mass helps plants to disperse over long distances, without the aid of mutualists. It also helps them to compete against other plants because of the sheer number of seeds which germinate (*ibid*). Short juvenile periods and intervals between seed crops further increase a plant’s competitiveness as it reaches the reproductive stage earlier than its competitors and can proliferate more extensively (*ibid*). This is especially important in disturbed habitats, which are the kind of habitats most susceptible to invasion (*ibid*).

Rejmánek and Richardson used this very simple framework to identify species of pines which pose particular threats to ecosystems. For example, they identified *P. radiata*, *P. contorta*, *P. halepensis*, *P. patula*, *P. pinaster* as the five most invasive species of pine and also identified 14 species of pine which do not pose a threat to invasion. Richardson & Rejmánek’s work is not a one-off case study. Their framework

can be generalized to cover many gymnosperms and even some angiosperms. For example, the framework can be used to make predictions about other conifers and some woody angiosperms (Rejmánek and Richardson 2004). They state that “conifers are probably unequalled in their capacity to improve the robustness of predictions in plant invasion ecology” (p. 327), because of the simple framework which has been shown to yield successful predictions. Rejmánek and Richardson conclude that the dispersal of pine seeds from plantations into adjoining natural or seminatural habitats is a major threat, especially in the Southern Hemisphere. They also identify forests in Sweden as a potential danger zone, as it is susceptible to invasion by *Pinus contorta*, a North American pine. Another interesting point which they mention, is their surprise that a recent global assessment of criteria and indicators for sustainable forest management did not mention the invasiveness of many conifer species. It seems that based on their predictions measures should be taken to ensure that conifer invasions do not occur.

It seems that this project is respected in the scientific community, as it is widely cited and to my knowledge, there is no criticism of the methodology or conclusions of the studies. In fact, it is considered by some to be the most successful example of prediction in invasion biology (Lake, 2004). However, many of the scientists who argue for larger-scale conceptual frameworks in invasion biology, have described these kinds of predictive frameworks as too limited. For example, Milbau et al. (2009) state that it is not enough to explore mechanisms of plant invasions at a particular scale, but that generalizations should be made across different scales. Moles et al. (2012) think of local studies as idiosyncratic, state that even though they produce “seemingly clear results, the larger picture is one of idiosyncrasy and inconsistency” (p. 117).

Yet this point should be seen as a problem for invasion biology in general, rather than a criticism of a study within invasion biology. The same authors who put forth these criticisms also state that invasions are extremely heterogeneous (ibid), and acknowledge that invasions are very complex processes, which can be studied in many (at least 7) different dimensions (Milbau et al 2009). Moreover, potential invaders are extremely heterogeneous, so it is impossible to extrapolate traits across taxa. For example, it is meaningless to talk of small seed mass in the case of animal invaders. In fact, when studies have attempted to abstract these differences between organisms, the resulting generalizations have been criticized as “trivial” (Richardson & Rejmánek 2004).

In other words, the criticism of these types of projects is not that they are not generalizable, but that they are not generalizable enough. Yet I think that pushing beyond mid-level generalizations is not a useful move in invasion biology. I picked Richardson & Rejmánek's work as an example, because it is one of the few documented cases which yielded successful predictions. In contrast, the large-scale frameworks cover more taxa, yet they have not resulted in precise and successful predictions.

## 8. Conclusion

Invasion biology can be characterized as a field with two different aims. The first is to understand and explain the fundamental nature of invasions and the second is to predict particular cases of invasion. The recent trend in invasion biology is the attempt to satisfy both aims, by creating unifying frameworks for invasion. The aim of this paper has been to raise some concerns about the feasibility of such a project. As invasions are extremely heterogeneous, for a framework to be truly unifying it cannot be based on factors which cause invasion. In order to unify invasions a framework must classify invasions in a different way, such as the invasion process. The problem then is that *predictions* must focus on causes, hence a unifying framework must give up the aim of prediction.

The point here is not to minimize the significance of unifying frameworks, as they are essential for providing explanations for invasions, both in terms of particular invasions and of the the nature of biological invasion as a whole. Instead, my aim is to show that these frameworks cannot accomplish both aims of invasion biology. The aim of prediction must be left to those integrative frameworks which aim for low to mid-level generalizations. In short, there can be a unified theory of invasion biology, but it will not be predictive.

## 7. Case Study

### Analysis of Target Systems in Invasion Biology

#### 1. Introduction

The extent of the turmoil in the field of invasion biology is such that it has generated a host of meta-analyses, ranging from methods and results (Davidson *et al* 2011) to basic concepts and definitions (Humair *et al* 2014). In the latter case, a group of scientists conducted a study on 26 academic experts, asking them to define key terms in invasion biology, such as ‘native species’, ‘invasive species’, ‘alien species’, and to evaluate the effect of invasions on ecosystems. The experts came from two groups, invasion biologists and landscape experts. The results of the study were very interesting as there was very little consensus on any of the questions, but more importantly, there was as much variation within the group of invasion biologists as there was between the groups.

Results like these should provide some perspective for the issues discussed in the previous chapter. If invasion biologists cannot even agree on when an organism has invaded a community, then how will they be able to agree on the best way to study invasions? These results also give some indication for why some invasion biologists have put aside the study of actual invasions, and are spending their time trying to provide a coherent conceptual framework for invasion biology.

I am in complete agreement with these invasion biologists, that their field can benefit from an investigation into its theoretical and conceptual framework. In fact, I think that these scientists are actually engaged in a form of target system analysis. The most obviously theoretical of these papers is the one by Gurevitch and colleagues (Gurevitch *et al*, 2011). They define conceptual frameworks as theoretical entities which “encompass the assumptions, laws and ideas that underlie the construction of a broad concept”, whose main point is to generate new hypotheses, evaluate existing ones and lead to the development of models and experiments. At the same time, they identify a number of reasons which they think explain the lack of consensus in the discipline.



For example, they argue that different approaches to the study of invasions often means that scientists end up studying ‘different phenomena’, by which they mean that invasiveness of organisms and invasibility of communities have important differences. These differences lead to difficulties integrating the two existing frameworks. Gurevitch et al. also have an astute diagnosis about the problem of generalizability. They state that “One of the challenges to creating useful frameworks for invasion is the tension between generality and specificity: overly general efforts risk explaining nothing very well, while overly specific contributions risk explaining only a limited range of cases.” (p. 408).

I think that in both cases of integration and generalizability, they have hit the nail on the head. It is precisely because of the differences in conceptualization of phenomena that integration is difficult, while it also the case that there is a tradeoff between generality and explanatory power in invasion research. However, I also think that it is possible to give a much more specific diagnosis of these two problems, and to provide some indications of how they can be overcome. In this chapter, I will examine the issues of integration and generalizability in turn, through an analysis of various target systems of invasion studies. I will show that on the one hand, integration is not as difficult as Gurevitch and colleagues imply, but that on the other hand full generalizability is beyond the scope of most invasion biology.

## 2. Integration

Integration is considered a valuable goal for science. In recent years there has been a lot of focus on integration, both in philosophy and scientific practice. In philosophy, it is often seen as a more attractive alternative to unification accounts of scientific explanation, as it retains the inter-disciplinary and collaborative aspect of unification without being reductionist (Brigandt, 2013b; O’Malley, 2013). In science, it is sometimes understood as a successful methodological paradigm, with initiatives such as ‘integrative biology’ acquiring journals, departments and disciplinary status (O’Malley, 2013).

There are many ways to define integration and to practice integrative science. In the previous chapter, I made use of Plutynski’s account of the characteristics of integration, which included localized problem solving and collaboration (Plutynski, 2013). Other philosophers of science different ways of categorizing the types of integration. For example, O’Malley distinguishes between *data integration*, *methodological integration* and *explanatory integration* (O’Malley, 2013). The first type of integration refers to the creation of common data sets used by different groups of scientists. In contrast, methodological

integration “involves directing a range of methods at a particular biological phenomenon or research problem in order to achieve multiple perspectives on how a system works or what the dimensions of the problem are” (ibid 552). Explanatory integration is the synthesis of ‘previously unconnected theories’ and the application of explanatory models to new disciplines. In a paper on integration in the plant sciences, Leonelli distinguishes between *inter-level integration*, which aim to “aims to acquire an interdisciplinary understanding of organisms as complex wholes”, *cross-species integration*, which is the integration of data from different species and *translational integration*, which combines data from within and outside academia, and which aims at improving human health (Leonelli, 2013).

Many of these types of integration are present in invasion biology. For example, in the previous chapter I gave the example of the studies on pine invasions, and how they exemplify many of Plutynski’s characteristics. They are also examples of cross-species integration, in Leonelli’s sense. The synthetic framework provided by Gurevitch et al. (Gurevitch *et al.*, 2011) is an example of explanatory integration, as it aims to integrate invasion biology with general ecology and evolution. It can also be understood as an example of methodological integration, as it combines a number of different methods for studying the problem of invasion.

However, producing successful integration in invasion biology is not as straightforward as it may seem. On the one hand the 1982 SCOPE publication explicitly called for integration between the two existing conceptual frameworks, as the third question asked how the results from the competition and community invasibility frameworks could be combined, in order to better predict and manage invasions. On the other hand, the general consensus is that there has not been sufficient integration in the field (Blackburn et al., 2011; Hayes & Barry, 2007; Williamson, 1999). In addition, there is no clear consensus about the type of integration which invasion biologists should aim for. The SCOPE publication seems to call for data integration, which may also leave room for methodological integration, while some invasion biologists seem to think that nothing short of explanatory integration is sufficient (Blackburn et al., 2011; Hayes & Barry, 2007; Williamson, 1999).

The main point of this section is to argue that target system analysis can help to disentangle the issues concerning integration, in two important ways. First, it can help us understand if integration is possible, and second, it can help to determine which type of integration is useful. For the remainder of the section I will examine both of these issues in turn, using successful examples of integration from invasion

biology. Throughout my analysis, I will be using O'Malley's distinction between data integration, methodological integration and explanatory integration (O'Malley, 2013).

### 3.1 When is Integration Possible?

Invasion biology is a heterogeneous field which has at least two distinct conceptual frameworks. In the previous chapter, I showed how the history of invasion research has led to the creation of two sub-fields, individual-level research which focuses on the traits which make an organism a good invader and community-level research, which focuses on the traits which make a community able to resist invasion. These two conceptual frameworks have distinct histories, theoretical frameworks and methodologies, as they use different types of experiments and models. Given this diversity, it is not surprising that the integration outlined by SCOPE has been difficult to achieve.

The main barrier to integration is that the two frameworks seem to be operating at very different levels. This means that the two frameworks have different partitions, as individual-level partitions have organisms as their main units, while community-level partitions have entire communities or ecosystems as their main units. At first glance, this implies that methodological and explanatory integration will be extremely difficult to achieve, if not impossible. These two types of integration go beyond the simple combination of data, and require a common ground of methods and explanation. However, as I showed in the previous chapter, the models and experiments of the two approaches differ significantly. Moreover, it can also be argued that the two frameworks are actually studying different problems, as one focuses on invaders and the other on invasibility.

I will leave aside the question of data integration for the moment. Operating at different levels does not preclude data integration in principle. In fact, this type of integration became quite common in invasion research, following the SCOPE publication. However, it is this type of integration which was not considered sufficient by the researchers, mentioned in the previous section. I will argue in section 3.3 that the kind of data integration achieved was not particularly useful for invasion biology. Instead, I will tackle the problem of more whole-scale integration head on. A closer look at the target systems of the two conceptual frameworks, reveals that methodological and explanatory integration is in fact possible, as both frameworks can be reinterpreted as operating on the same level.

It is community-level research, which should undergo the most radical reinterpretation. The poster-child of community-level research was biodiversity. This is because biodiversity is thought to be an emergent property of communities (Goldstein, 1999), which is not reducible to interactions between its constituent parts. Following Elton, it was thought that biodiversity was negatively correlated with invasibility, as low diversity communities that could be more easily upset by invaders (Dunstan & Johnson, 2007). This claim was based on the idea of the “balance of nature”, where communities exist in a state of equilibrium until they are disturbed (ibid). Communities which were highly diverse were also thought to be more complex, in the sense that they had many more ecological interactions. Thus, a disturbance would upset a lower percentage of interactions in a highly diverse community, whereas the same disturbance in a non-diverse community would constitute a much more extensive disruption (ibid).

The main problem with this claim is simply that it is not empirically supported. While there are examples of highly diverse communities which have resisted invasion (Kennedy et al., 2002), there are many studies which reveal a positive relationship between biodiversity and invasion (Levine & D'Antonio, 1999). In response to these empirical results, some researchers have focused on ‘species richness’ a more deflationary interpretation of biodiversity.<sup>1</sup> While there are many conflicting notions of biodiversity and its measurement, species richness is thought to be a much more straightforward notion to define and measure (Maclaurin & Sterelny, 2008). Unfortunately, in the case of invasions, the empirical problem remains, as there is no evidence that species-diverse communities are better at resisting invasions (Dunstan & Johnson, 2007).

These empirical results led some researchers to look at the relationship between diversity and invasibility more closely. This gave rise to the theory of resource fluctuation as the actual cause of invasibility (Davis et al., 2000; Dunstan & Johnson, 2007). The theory states that at least in plant communities, a community becomes susceptible to invasion when it has an increase in the amount of unused resources (Davis et al., 2000). It is based on the idea that the intensity of competition is inversely correlated with the amount of unused resources, simply because there are more resources to go around

---

<sup>1</sup> There is an important debate in philosophy of biology concerning the term ‘biodiversity’ and how it should be interpreted. For some, species richness is not a viable substitute for biodiversity, as biodiversity is more than just species richness {Maclaurin:2008ww}. However, an examination of this debate is beyond the scope of this chapter. I will assume, for the sake of the argument that species richness is a viable alternative, for the purpose of determining the invasibility of a community. This is because species richness faces the same problems as biodiversity in terms of explaining invasibility.

(ibid). Thus an invader will be more successful if it does not have to compete so intensely with the resident community.

Interestingly, the theory of fluctuating resources can explain why some highly diverse communities can resist invasion and why some cannot. In some cases, highly diverse communities also utilize a large percentage of their resources and there is little fluctuation, hence those communities are resistant to invasion. However, this can also happen in low-diversity communities. At the same time, communities which have high resource fluctuation will be more susceptible for invasion irrespective of their level of diversity (ibid).

The interesting point for target system analysis, is that unlike biodiversity, fluctuation of resources is not an irreducible property of communities. Of course, it is a property of communities, as we can measure the amount of resources in a community at any given time. However, the resources themselves not emergent properties of communities but parts of those communities. For example, the amount of nitrogen in a system can be thought of as a property of a community, but it can also be understood as a part of that community. The upshot of this is that at least in the case of resource-fluctuation, community-level invasion research does not have communities as units in its partitions.

Before I examine the partitions of resource-fluctuation target systems, I will address a potential worry of this picture of community-level research. It can be argued (though it has not, to my knowledge) that resource-fluctuation should not be considered a case of true community-level research, as it does not have communities as its main units. There are two counterarguments which can be offered. The first is that there is no reason to think that community-level research must have community-size units. It is sufficient that research focus on properties of communities, for their research to fall into that camp, without stipulating that the properties themselves are *emergent* properties of communities. As I argued above, the amounts of resources in a system are properties of that community, even though they are not emergent. The second counterargument is that there are no viable emergent properties of communities which explain invasibility. The main other factor in community-level research is disturbance, as it is often the case that invasions occur a disturbance event. Yet disturbance cannot be understood as an emergent property of a community, either. Therefore it seems that the theory of fluctuating resources remains the most viable option for community-level research.

I will now turn to an example of a study which examines fluctuating resources, in order to illustrate the types of partition which the theory uses. In their paper “Mechanisms of invasions: can the recipient community influence invasion rates?” Dunstan and Johnson (2007) use two individual-based models to determine the relationship between resource fluctuation and invasibility. In the interest of brevity, I will focus only on the first model, which aimed to determine whether resource variability or species richness had a greater effect on invasion resistance. The model contained a number of individual agents, organized in groups of species. Each individual consumed  $r$  from the total pool of resources  $K$ . In addition, each individual had a birth rate ( $b_i$ ) and a death rate ( $d_i$ ), which are the probabilities of the agent producing a single offspring or dying. Finally, individuals from different species varied in terms resource consumption ( $r_i$ ) the amount of resources they needed in order to survive.


Dunstan and Johnson ran the model over two different scenarios. In the first, the mortality (death rate) of each individual was independent of resource consumption, while in the second, it was inversely related to the resource consumption of the species. In other words, increasing the amount of the resource in the second scenario reduced the probability of mortality. In both scenarios, the original set up included individuals from 20 different ‘native’ species and a  $K = 1000$ . These individuals interacted for 2000 iterations utilizing resources and establishing a particular population structure through competition for these resources. At this point the number of species was uncontrolled and varied from 2 to 20. Then, an ‘invader’ was introduced and the model was left to run to 2500 iterations. In the meantime, the variability of the free resources was calculated as the variance in unconsumed resources between 1000 and 2000 iterations.

The models produced some interesting results. In the first scenario, where mortality was independent of resource utilization species richness varied negatively with invasion success (see Figure 1 (A)). This result is consistent with the standard view that rich communities are less susceptible to invasion. However, the second scenario showed the opposite result, i.e. A positive correlation between species richness and invasion success (see Figure 7.1(A)). According to the authors, this should not be surprising, as resources are the underlying causes of these different results. This is because in species-rich communities there are by definition, more species, yet each species consumes fewer resources ( $K = 1000$  for all cases). This means that the death of an individual in a species-rich system releases fewer resources than it would in a species-poor system, hence each death results in a smaller increase in available resources. In

addition, when mortality is independent, for the same death rate, fewer resources are freed in species-rich systems than in species-poor systems. Hence there are fewer available resources for invaders in a species-rich system. This is also the case when mortality is dependent on resource utilization. However, in this case, there is a feedback effect in the system, so that the overall death rate for species-rich systems is higher than for species-poor systems. Therefore, every death actually frees more resources which the invaders can utilize.

Interestingly, the results also showed that the variability of resources was always positively correlated with invasion success (see Figure 7.1(B)). This can be explained by the fact that when resource availability is stable in an ecosystem, then the native organisms are usually well-adapted to that environment (and/or phenotypically plastic), hence they can utilize the available resources efficiently. This means that when a new organism invades it simply cannot get a foothold in the ecosystem. This result should not be surprising, given that approximately only 1% of all invasions are successful (Weidenhamer & Callaway, 2010). On the other hand, if resources fluctuate, then at times when resources are high, invaders will be able to utilize all excess and become established in the ecosystem.

**Figure 7.1.** Correlation between invasion success and (A) species richness (B) resource variability



There are a number of target systems which this model can represent. According to the authors, the model is particularly suited to the investigation of seaweed communities. This is due to the the fact that

many species-rich seaweed communities have been susceptible to invasion, yet all of these communities are also characterized by resource fluctuation, due to seasonal oceanic and tidal effects. While the authors themselves do not apply the model to any particular community, they make a number of suggestions which show how they think a target system should be specified so that it can be represented by the model.

Specifying the partition of a target for this model is quite easy, as the model is an individual-based model, hence individual organisms must feature as units. In this sense, the model is a departure from traditional community-level invasion research, as the partition does not individuate communities as units. Properties of these organisms such as births, deaths, biomass and metabolic rate will be used to calculate the average birth rate, death rate, resource consumption and growth rate for each species. These, in turn will be used to determine the various competitive strategies for each species.

Yet this model retains some aspects of community-level research, which is typically absent from individual-level research. The first is the importance of resources, as units which are relevant causal factors. In other words, this partition individuates resources such as nitrogen, phosphorus and light, and these resources are not omitted in the abstraction. They are represented in the individual-based model as 'resource units', which are consumed by the organisms, yet they also vary. The reason why they are more important than in traditional competition models, is that they also interact with the organism agents, as they would in a predation model.

The second is the importance of space as a limiting factor. In community-level research, this is usually represented as a dimension of the organism's niche, as two organisms cannot occupy exactly the same space. As seaweeds are relatively immobile, each organism needs a minimum amount of space to become established and grow. In the individual-based model, space is represented by the patches, as each organism can occupy only one patch. In contrast, space is typically not individuated in traditional competition models, as the carrying capacity acts as a proxy for space.

Dunstan and Johnson consider their research as part of the community-level tradition, as they think that the most important factor which determines an invasion is the fluctuation of resources, which is a property of communities. Yet its partition is very similar to the partitions traditionally associated with the individual-level tradition. The remaining difference between this model and traditional individual-level models is the inclusion of resources and space. However, this is a difference of abstraction, not partitioning, as the target of this model includes resources, which are usually omitted.



In fact, the Dunstan and Johnson's model is an example of methodological integration between the two frameworks. It incorporates aspects of both traditions, as the model is essentially an individual-level model, yet with the addition of causal factors from community-level research. It is also considered to be an example of a successful merger between the two frameworks by other invasion biologists (Richardson & Pyšek, 2006).<sup>2</sup> It also aims to achieve a new perspective on how systems of invasion work, and explores new dimensions of the problem of invasion.

The point of this discussion is to show that analyzing the target systems of two different conceptual frameworks can help us determine if integration is possible, whether it will be valuable, and how to achieve it. In this example, given that the integrative model is already in existence, the analysis serves as an explanation for why the integration was achieved. The analysis showed that the partitions of the two frameworks are not very different (provided that the community-level framework was reinterpreted in the appropriate way). It also showed how the integration was achieved, by locating the issue at the level of abstraction instead of partitioning.

Finally, the analysis highlighted the fact that the integration was relatively uncomplicated, as it was achieved by changes in the level of abstraction rather than partitioning. Individual-based models are structured so that adding factors is relatively simple. This does not mean that integration of two frameworks with different partitions is impossible to achieve. However, it is often much more complicated. Hence, in this case target analysis showed not only that integration was possible, but also that it would be relatively simple to accomplish.

Yet there are cases where integration turns out not to be as simple or as useful. In the next section I will give another example of integration in invasion biology and use target system analysis to show extent and limitations of this type of integration.

### 3.2 When is Integration Useful?

In my discussion of the history of invasion biology (chapter 6), I mentioned that many of the original attempts to integrate the two frameworks of invasion biology have been largely unsuccessful. In particular, I mentioned the theory of seed plant invasiveness, which attempted to pair up particular traits of plants with particular environments in which they are likely to invade. The idea is that while particular

---

<sup>2</sup> See also chapter 6 pp. xx

traits make organisms better competitors, most invasions fail, hence some communities are less able to compete with individuals with those traits (Facon et al., 2006). For example, small seed size is a factor which leads plants to become more widely dispersed and because of this wide dispersal, they are more likely to be adapted to a wider range of conditions (Richardson & Pyšek, 2006). The idea is to find environments which are particularly susceptible to plants with that trait, so as to better predict invasions. This would be an example of data integration, as it would combine data from individual-level and community-level invasion research.

There are numerous other pairings which could be made. For instance, the aquatic cattail *Typha angustifolia*, which I mentioned in chapter 5 managed to invade Indian Creek because it was allelopathic. As toxins are much more easily distributed in small aquatic environments such as creeks or lakes, another pair of individual and community traits could be allelopathy small aquatic environments.

The list of available pairing between traits of individuals and communities is long. For each particular organism, we can identify an environment which has been susceptible to invasion by that particular organism. However, one problem with this method is that these results are not generalizable beyond each particular instance. That is, identifying a combination of a trait with a particular environment as the joint causes of a particular invasion will not automatically be more applicable to a larger number of systems. If anything, it will probably be even be less applicable, as there will be fewer systems that have both characteristics. In other words, combining potential causes of invasion may lead to a decrease in the instances of false positive predictions of future invasions, as it will rule out cases where a particular trait does not promote invasiveness because it is not in the right sort of environment, or cases where a particular community is not susceptible to invasion because the invader does not have the right sorts of traits. Nonetheless, this does not mean that such a combination will increase the instances of true predictions of future invasions. Invasions will probably continue to occur in environments that we had not even thought of as susceptible to invasion and by organisms that we had not envisaged as invaders. Moreover, there are probably a great many combinations of traits and environments that will result in invasions, which again, we have not even imagined.

The underlying problem with this method is not that the integration itself is impossible, but that it is not clear how it is useful. All of these cases are examples of *data integration*, as they combine results from individual-level and community-level research in order to explain (and predict!) invasions. However,

the integration is limited to data integration, and there is no methodological or explanatory integration. That is, the combination of traits and environments does not involve a range of methods in order to achieve multiple perspectives on how a system works, nor does it create new explanatory models.

I will illustrate this point by revisiting the example of *Typha angustifolia* from chapter 5. In my analysis of the target system of the experiment, I showed that Jarchow and Cook were firmly situated within the individual-level framework. The partition individuated organisms as units, with allelopathy as the trait which caused the invasion of Indian Creek. Of course, they were not aiming for any sort of integration, as they did not include any community-level characteristics. Yet the question is, what would the integration of community-level characteristics add to the explanation of the invasion? Jarchow and Cook successfully showed, with their experiment, that *Typha angustifolia* succeeded in invading Indian Creek, because it was allelopathic (Jarchow & Cook, 2009). The explanation of this particular invasion is complete. The only potential gain from including community-level characteristics would be the generalization of this particular result to other domains. However, as I argued above, the generalizability of these types of results is severely limited. Therefore, it seems that while data integration is possible, it is not always useful.

This characterization of the problem mirrors the arguments of the invasion biologists, mentioned in section 2, who are calling for greater integration of the field. In the case of invasion biology, data integration is not sufficient for better predictions of invasions, nor does it add much to existing explanations of particular invasions. Yet the important point for my purposes is that target system analysis offers a way to diagnose the problem. In this particular case, data integration does not involve any radical changes in the target systems of existing models. But this shows that the existing framework and resulting target systems are apt for the purposes of explaining particular invasions. Merely integrating data will not provide a better explanation, hence it merely complicates the procedure without providing any additional benefits. Still, as I showed in the previous section, target system analysis also shows that other forms of integration are possible and can often be very useful.

In the end, the most important goal for invasion is prediction. While integration is one part of the puzzle for achieving successful predictions of invasion, another part is generalizability. In the next section I will show how target system analysis can help explain why it is so difficult to generalize results and explanations of invasion.

## 4. Generalizability

The analysis of the targets in invasion biology can also help to untangle the other puzzle of invasion biology, namely why it is so difficult to generalize the results of invasion studies. More specifically, while it is possible to make grand generalizations in invasion biology, these tend to be generalizations that are rather trivial, and thus not particularly useful for explaining or predicting invasions. For example, a factor which correlates very strongly with invasions of both plants and animals is facilitation by humans. However, humans facilitate invasions in many ways. They physically introduce organisms to new environments (consciously or unconsciously) (Levine & D'Antonio, 1999), they disturb environments directly or indirectly (Facon et al., 2006), they prey on organisms and so on. While all of these activities facilitate invasions, they do so in different ways, which have very different effects on ecosystems. In fact, there are important differences even within these types of facilitation. Hydraulic fracturing (fracking) and human-induced wildfires are both examples of extreme disturbance, yet the effects that they have are very different. In order to be useful, an explanation of an invasion must be more specific, by including the mechanisms of the invasion. However, doing this immediately sacrifices the generality of the explanation, as the mechanisms of invasions differ greatly across different cases.

While this may seem like an insurmountable problem, analyzing the target systems of invasion can be very informative. First, it gives us a precise way of explaining when generalizability fails and why it is the case. Second, it can reveal the extent to which a result from an invasion study can be generalized. In the previous chapter I gave an overview of the field of invasion biology and highlighted the problem of generalizability. Here, I will show how target system analysis helps to resolve the issue, by bringing to the forefront the conflict between causal heterogeneity and generalizability which characterizes biological invasions.

One of the most important tools in the target analysis toolbox, is the identification of difference-makers. In the case of invasions, a difference-maker is a causal factor, present in the domain which makes a difference (or in some cases *the* difference) between the success and the failure of the invasion. In the invasion of Indian Creek by *Typha angustifolia* (chapter 5), I showed that the difference maker was allelopathy, the exudation of toxins from the roots of the plant. In the pine studies by Rejmánek & Richardson (Rejmánek & Richardson, 1996; Richardson & Rejmánek, 2004), small seed mass is the difference-maker in invasions that involve pines. Other factors which have been shown to be difference-

makers in particular plant invasions include phenotypic plasticity (Daehler, 2003; Geng et al., 2006), plant-soil feedback (Callaway et al., 2004; Klironomos, 2002), disturbance through fire (Buckley, Bolker, & Rees, 2007), disturbance through climate change (Hayes & Barry, 2007), biological inertia (Joost, 2003), small seed size, (Richardson & Rejmánek, 2004), propagule pressure (Van Kleunen, Dawson, Schlaepfer, Jeschke, & Fischer, 2010).

These difference-makers are a very diverse group. Some of them are particular traits, yet they are very different types of traits. For example, small seed mass and allelopathy are traits that affect competitive ability at very different stages of the invasion process. In addition, a trait which is a difference-maker in one case, might not be a difference-maker in another. Thus, while small seed mass is a difference-maker in pines, it was not a difference-maker in the *T. angustifolia* invasion, even though *T. Angustifolia* does have small seed mass and produces a large number of seeds.<sup>3</sup> Other difference-makers are relations between organisms, such as plant-soil feedback. This is a particularly interesting example, as there can be different kinds of feedback which affect invasiveness in different ways. That is, soil microorganisms which fix nitrogen and thus increase a normally limiting factor for plant growth can coexist symbiotically with other microorganisms which are harmful to the plants.

As I argued in the previous chapter, in order to be apt and to provide the setup for a successful explanation, the target system of an invasion study must identify the difference makers of that invasion event. On the other hand, a general explanation or framework for invasion must be able to account for most if not all of the particular instances of invasion. It should be easy to see that a general explanation cannot focus on some difference-makers and leave out others, as it would only be an explanation of some subset of invasions. Nor can an explanation be truly general or unificatory if it identifies each and every difference-maker separately.

Of course, another option available is to to dispense with difference-makers and causal relations entirely and unify invasions in another way. This is exactly the strategy employed by Blackburn et al. with their reconceptualization of invasions as processes (Blackburn et al., 2011) (see chapter 6). However, there are good reasons for wanting at least some explanations which take causes into account. We still need to

---

<sup>3</sup> We can assert this with confidence, because we know that *T. Angustifolia* is much more invasive than other members of its genus, with which it shares the traits of small seed mass, fast growth rate etc.

identify the causes of particular invasions, even if the results do not generalize. In addition, as I argued in chapter 6, causes are necessary for the successful prediction of invasions.<sup>4</sup>

The third option is to find a way to group together the difference-makers in broader categories, which still retain their causal power. I will argue that while this is possible, it can only be achieved within certain limits, that of mid-level generalizations. It is impossible, given the heterogeneity of the difference-makers to generalize beyond that while retaining the causal character of an explanation. I will illustrate with the example of the pine studies by Rejmánek & Richardson (1996, 2004) mentioned in the previous chapter. These studies are generally considered to be successful, in that they have generated important predictions, and they are also modestly general, as the results have been generalized to other conifers and some woody angiosperms. In order to show why these results cannot be generalized further, I will first show how they managed to achieve even this level of generalizability. I will do this, once again, by looking more closely at their target systems.

The partitions in Rejmánek & Richardson's studies are all individual-level partitions. That is, they individuate organisms (pines) as units, and the difference-makers are identified from the properties of these units. Rejmánek & Richardson did not know beforehand, which of these many factors were difference-makers, hence they made a shortlist of 10 potential difference makers, and then conducted a discriminant analysis on those to reveal three: small seed mass, short juvenile period and a short interval between seed crops. These three traits *together*, explain why some pines are invasive and others are not. Rejmánek & Richardson hypothesized that the reason why these simple factors are difference-makers in invasions, is that pines have quite simple 'regeneration requirements' (Richardson & Rejmánek, 2004). They are gymnosperms, do not rely on other organisms for their dispersal, and invest little into each seed. Given this reproductive strategy, the species with these three traits are the most successful when it comes to invading new areas.

Rejmánek & Richardson hypothesized that these results could be applied to other conifers and some woody angiosperms. The reason for this, in the case of conifers is that they have some important similarities with pines (which are also conifers). The most important of these is that they are gymnosperms,

---

<sup>4</sup> It is interesting to note that in other disciplines explanation is commonly tied in with causes, while predictions can be generated without reference to causes. This has also become part of many accounts of scientific explanation. For example, a standard counterexample to Hempel's account, the 'barometer and the storm' concerns a situation where the drop in the barometer predicts the storm's coming, but it does not explain the storm. Hence, it is a case where prediction is 'easier' than explanation. Nonetheless, this is not the case for invasion biology as the systems in which invasions occur are extremely complex and not similar enough to each other, hence even when an explanation is successful, it does not give us good grounds for prediction.

and therefore do not rely on mutualisms in the way that angiosperms do. Of the various conifers, cypresses (Cupressaceae) were found to be the most invasive, while within the family, those species which had the difference-makers were the most invasive. This is because still being conifers, the strategy that works well for pines works well for cypresses too. However, it is important to note, that there were more exceptions within the cypresses than there were within the pines. That is, there were some cypresses which were not invasive despite having the three difference makers. This shows that the generalizability of the results is already starting to be limited, as there may be a different set of difference makers which are the best predictors of invasions for cypresses.

The situation becomes more complex in the case of angiosperms. First of all, they differ from gymnosperms in an important way, as they rely on other organisms to reproduce and disperse their seeds. In fact, the results were much less robust for the angiosperms than they were for the pines or even the cypresses. Nonetheless, some patterns emerged, namely that woody angiosperms (trees and shrubs) tend to be much less invasive than conifers (esp. pines) and that the rarest of these species have none of the difference-makers. Still, these results are much less robust than the results for the pines. Again, this shows that the generalizability of the results is limited to small groups of organisms.

The explanation for this limitation is that the difference-makers for invasion, are only difference-makers within a certain context. In this case, the context is the general reproductive apparatus and strategy. These three causes (small seed-mass, short juvenile period and short interval between seed crops) are only difference-makers given reproductive organs which produce these kinds of seeds frequently, and within a reproductive strategy which allocates few resources to the production of each seed. If a plant had the same characteristics within a different reproductive strategy, then the same factors would not make it more likely to invade. This does not mean that they would not confer other advantages to the plant, just not advantages for invading.

The general point is that the factors which are difference-makers in one group of organisms or species are often completely absent in others. This means that it is very unlikely that the knowledge of the difference-maker in one species will help us determine anything about behavior of a species which does not have that trait. This is the main reason why plant invasions and animal invasions are so different. Yet even within smaller taxa, identifying a known difference-maker in another species, does not mean that it will be

a difference-maker in that species. Whether or not a factor counts as a difference maker depends, in part, on the context within the organism itself, its other traits, strategies and interactions with other organisms.

The upshot of this discussion is to show the reasons for which the generalizability of results is limited, in the case of invasion biology. Moreover, these reasons are good reasons, and scientists will not gain much from trying to ‘overcome’ these limits. Invading organisms and invaded communities are so heterogeneous, that focusing on the cause of a particular invasion will not give us the tools for a general explanation of invasions, or general theories for predicting invasions. However, target system analysis can help us determine which organisms are sufficiently similar to each other, so that we can know *when* the identification of a factor is difference-maker.

## 5. Conclusion

The main of this chapter was to show that the nature of integration and generalizability in invasion biology can be elucidated by target system analysis. Some invasion biologists have already started conducting a form of target analysis, by re-examining the conceptual frameworks of invasion biology and looking for ways that they can be improved. However, a target analysis which is both scientifically informed and philosophically grounded can be more informative. In this case, it showed that it is possible to integrate the two main conceptual frameworks of invasion biology, provided that they are re-conceptualized to some extent. In addition it showed that if a study in invasion biology focuses on the causes of the invasion, the results can only be extended to mid-level generalizations.

These results are important both for invasion biology but also for the role of target analysis. The argument in this chapter highlights just how useful and informative target system analysis can be. An important point to keep in mind is that if target systems were not considered to be real parts of the world, the type of investigation presented in this chapter would be beyond the scope of target system analysis. It is only because of the clear distinction between models and targets which comes from my view of targets as real parts of the world that target system analysis can reach its full potential, informing scientific practice.



## 8. Conclusion

The protagonist of this dissertation is the humble and generally overlooked target system. Unanalyzed, a target system is simply a vague term referring to whatever it is that a model models, a shorthand which philosophers of science use to denote the uninteresting aspect of model-world relations. I hope to have shown through this dissertation, that target systems can offer much more to philosophers and scientists alike. They are the final piece in the puzzle of model world relations, as their analysis reveals part of the relationship between models and the world. Target system analysis comes in two parts, the first being the specification of target systems, which is comprised of four parts: identifying the domain of study and the phenomenon of interest, partitioning the domain and abstracting to reveal the difference makers of that phenomenon. The key to understanding the importance of target systems is to realize that they are real parts of the world, which are therefore easily distinguishable from the models which represent them. In other words, they are not intermediaries between models and the world, but literally are parts of the world.

This view of target systems as real parts of the world, seemed (at first) to be at odds with my account of partitioning and abstraction. Defining both partitioning and abstraction for target systems involved some complexity and required investigation of other areas in philosophy of science. In the case of partitioning, this was the debate on natural kinds. Different views of natural kinds have different implications for the metaphysical status of partitions, yet I hope to have shown that many moderate views are compatible with an understanding of target systems as real parts of the world. In the case of abstraction, the issue was that common views of the notion of abstraction run together the concepts of omission, de-concretization and generalization. I hoped to have shown that the core notion for abstraction is omission and that omission alone does not change the ontological nature of target systems.

The second aspect of target system analysis is the theory of target system evaluation. The main difficulty in providing a theory of target system evaluation is that, target systems are real parts of the world, hence they are not good or bad in themselves, but must be evaluated with respect to a particular scientific purpose (often a model of the phenomenon of interest). Hence I devised a theory of evaluation based on aptness, which is itself a relativistic notion, tied to a function or goal. This led to another difficulty, namely the provision of a theory which was sufficiently pragmatic to accommodate the importance of context for

the choice of partitions, but also sufficiently strict provide a criterion for abstractions which cuts across contextual factors. Thus was born a theory which combined aspects of van Fraassen's account of the pragmatics of explanation and Strevens's kairetic account of causal explanation.

An integral part of this dissertation were the ecological models which I used extensively to explain the various intricacies of target system analysis. This reliance on ecological examples was not entirely accidental. Ecology is a fascinating discipline. Its content is captivating, as it deals with extremely complex systems, comprising of numerous parts, properties and a multitude of interactions between them. In addition, ecology is a microcosm of science in general, both in terms of the numerous methods ecologists used to study phenomena, to the problems associated with each one. In fact, it was ideally suited to the investigation of the role of target systems, as it furnished my investigation with examples suited to each of the issues associated with target systems.

Yet the importance of target systems need not be limited to models in ecology. It is possible that other fields might also benefit from rigorous target system analysis and evaluation. To some, this may seem unnecessary, especially in disciplines where there are well-established norms governing model selection and where these models are generally regarded to be successful. In fact, this can lead to a more general argument against the importance of target systems. The argument rests on the idea that scientists have been using models to explain phenomena in the world for a very long time, with sufficient success. Moreover, the successful practice of modeling already includes what I am calling target systems and target system analysis. This gives rise to a second, stronger criticism, which focuses on the philosophical aspect of modeling. The idea is that philosophers have provided a number of sophisticated accounts of model-world relations, which adequately explain how models relate to the world and can account for any difficulties which arise in the application of models to real-world systems. Proponents of this criticism would argue that accounts of model-world relations are complete and that the inclusion of target systems simply introduces additional complications to the picture.

The response to the first criticism is to agree that the practice of modeling already includes many aspects of target systems. Scientists identify domains and phenomena, categorize these domains into units and make judgments about which units are relevant. It is also true that scientists think of these activities as aspects of modeling. As I showed in the introduction, with the example of the Vancouver Island marmots, scientists often spend a significant amount of time and effort doing this preliminary work before they run

their models. I actually think that this is actually an advantage of my account. One of the reasons why it is important to give a *philosophical* account of target systems and their analysis is precisely because these activities are part of scientific practice. In other words, target system analysis is an aspect of science which philosophers have not paid enough attention to.

Ultimately, the point of this dissertation is to highlight the importance of this particular aspect of science. I think that distinguishing this aspect of science from other parts of modeling has important advantages, yet the importance lies in the particular scientific practice itself, not the name we give it. In other words, what matters is that the process of identifying a domain and a phenomenon, partitioning and abstraction are given the requisite attention from scientists and philosophers. If it does not really matter if we do not call this aspect of science target system specification and analysis, but simply think of it as a part of modeling.

However, proponents of the second, stronger criticism might not be satisfied with this answer. The point of this criticism is that my account simply complicates things, as existing philosophical accounts of modeling adequately describe the scientific practice of modeling. Here, the worry is that talking about target systems does not add anything of real value to the existing picture.

There are a number of responses to this criticism, yet I will focus on two. First, I suspect that the criticism rests on a misunderstanding of why target systems are important and useful. Target systems are not an ontologically distinct entity, as they are parts of the world, which means that the inclusion of target systems does not add anything new in the metaphysical sense. That is, even with target systems in the picture, there are only two types of entities, systems in the world and models of those systems. Again, this is an advantage of my view. Critics would be right in arguing that the introduction of a new kind of entity would complicate matters without adding any benefits to existing accounts of model-world relations. In fact, my argument against other conceptions of target systems as fictions or as models is based on this point.

However, the importance of target systems stems from the *epistemic* benefits that they have for conducting and understanding science. Specifying, analyzing and evaluating target systems are all ways of making these aspects of science more explicit, so that (i) we achieve a better understanding of the parts, properties and mechanisms of a real-world system in which a phenomenon manifests and (ii) we achieve a

better understanding of how the model we are using represents the system in the world. In other words, understanding target systems gives us a better grasp of how a phenomenon should be modeled.

This leads to the second counterargument, namely that existing accounts of model-world relations are strengthened by the inclusion of target system specification and evaluation. Even though some of these accounts make some reference to aspects of target system specification, there are important gaps which can be filled. I will illustrate by focusing on the case of partitioning. Existing accounts of model world relations sometimes make implicit reference to partitioning, as in many cases particular models or scientific disciplines require a certain type of partition. For example, Weisberg's account of model-world relations is based on the notion of similarity between models and targets (Weisberg, 2013). Determining the extent to which a model is similar to a target is achieved through identifying common mechanisms and attributes between the model and the target.

However, these mechanisms and attributes presuppose a particular partition of the system from which the target is specified. That is, in order to be able to identify aspects of a system as mechanisms or attributes, we must partition the system. Thus, partitioning is implicit in the view of similarity yet not explicitly addressed. Yet, there are cases where explicit analysis of partitions is necessary and can help the similarity account of model-world relations. The problem is that each domain can be partitioned in many ways, and the model itself might not always fully determine which is the most useful partition. For example, if we are using an individual-based model in ecology, we can include a particular factor as an agent or as a property of other agents. In a model of plant competition, resources such as nitrogen and phosphorus could be simply regarded as part of the general resource pool, or they could be actual agents (or patches). These two alternatives correspond to different partitions of the system, yet the model itself does not dictate which is more useful. The particular choice could lead to differences in the predictions of the model which can only be understood and solved if we pay close attention to the partition we are using.

In addition, sometimes models contain features which do not have analogues in real-world systems. I am referring to examples such as the growth rate of the Vancouver Island marmots, which is not similar to any part of the world or to any property of individual marmots. Hence, in order to assess the similarity of the model to key features of the marmot population there are a number of other steps which need to be taken, that is, identifying the factors from which the growth rate is calculated. Paying explicit attention to these features, by determining that they are relevant for the scientific purpose and including

them in the target system is epistemically valuable as it can give us a better picture of the phenomenon taking place in the system.

The upshot of this point is that existing accounts of model world relations are not complete. While some of them include some of the aspects of target-system specification, they do not include all of them. My account provides the final piece in the picture of model-world relations, as it explains the nature of target systems and their relationship to systems in the world. While this might not seem to be especially problematic, there are cases where a model's results are not successful, yet the problem lies with the target system. Examples of cases like these were presented in chapters 6 and 7 in my discussion of models in invasion biology.

I will now address a different issue, which demonstrates another way in which target systems are important. There are some clarifications which need to be made, so that some of the advantages of my theory of aptness are made more palpable. These are distinctions between aspects of aptness as a theory of target system evaluation, which need to be reiterated and highlighted in order to dispel confusion about the nature and importance of target systems. I will focus on two clarifications: (i) the difference between maximally and sufficiently apt targets and (ii) the difference between aptness at a local and global scale. These can also be seen as further evidence of the importance of epistemic issues and benefits in the discussion of target systems.

The first distinction is between maximally and sufficiently apt targets. A maximally apt target system is the best possible target system for a particular scientific purpose. This means that the target has the most useful partition and contains only the difference makers for the phenomenon and no other factors. In most, if not all cases, scientists and philosophers are not in a position to determine the maximally apt target for a particular scientific purpose. Instead they try to identify a target which is sufficiently apt for the purpose. But how can they determine if their target is sufficiently apt without being able to identify the maximally apt target for each scientific purpose? The answer is that the maximally apt target for a particular purpose is best understood as a regulative ideal, something which we can strive for, even though our epistemic limitations will not allow us to fully determine it. Trying to understand what a maximally apt target would look like can help us understand *the way in which* target systems are apt, that is, by having useful partitions which individuate potential difference-makers and which omit everything which is not a difference-maker.

An implication of this point is that it shows that there is a fact of the matter about the aptness of a target given a particular purpose. This is important because it shows that target system analysis is not overly relativistic. The fact that in practice we are often unable to determine the maximally apt target for each scientific purpose does not mean that there is no such thing as a maximally apt target. This in turn means that the standard of target system evaluation is objective, even if scientists' actual judgements about the aptness of particular target systems are limited by what they know at the time. This can be demonstrated by the comparison of two target systems. If we have two targets which differ in terms of partition and/or abstractions there is a fact of the matter about which is closer to the maximally apt target. In some cases we will be able to determine this, yet in others we will not. There might also be cases where we are completely mistaken in our evaluation and the favored target system is actually less apt. The important point is that it is our epistemic limitations as human scientists and philosophers which keep us from being able to make the correct judgement.

This leads to the second distinction between local and global conceptions of aptness. Strictly speaking, this is not a distinction between two concepts, but a scale with local and global at the two extremes. On the local end of the spectrum a target is considered apt for a very specific scientific purpose, entirely determined by a single model applied to a single system. Here I am referring to mathematical or computational models such as the logistic growth model which was applied to the population of Vancouver Island Marmots. In this case, the context is given entirely by the model, and the target is apt if it picks out the best partition and set of abstractions for that model. In some sense, this is the easiest way to determine the aptness of a target, as it is relatively easy to determine the context for usefulness and relevance, when that context is determined just by the model. However, in many other cases, the determination of context is much harder, as it is broader than what is given by model itself. For instance, a scientist might specify a target system with a family of models, rather than a particular model, in mind. This means that the 'scientific purpose' which provides the context for evaluating aptness will be more broad. In these cases the context is determined by analyzing the concepts, methods and norms of a discipline or sub-discipline.

At the most broad level, we can think of the aptness of target systems in science as a whole. That is, as science progresses and we learn more about the systems from which targets are specified, we are able to make more general judgements about the aptness of target systems. For example, the invention of microscopes and telescopes make it the case that partitions of systems into units that are smaller or larger

than can be distinguished by human senses can still be useful. On a less general level, as biologists now know that soil biota affect plant growth, it is possible to make relatively general claims about the aptness of a group of target systems: soil biota are difference-makers. Moreover, these general or global evaluations of aptness have an interesting relationship with the more local evaluations of aptness. In some cases, the general evaluations are used to determine the context for the local evaluations. In other cases, examination at the local level reveals an exception to the general level, and leads to a re-evaluation of the general notion of aptness. The point is that all of these evaluations are possible and important, as they are what helps us get a better understanding of the real-world systems and phenomena that are the subjects of scientific inquiry.

This point should not be taken to undermine the relative objectivity of target system evaluation. That is, the fact that the context itself can change should not be taken to imply that target system evaluation becomes overly relativistic. In other words, I am not claiming that the objective standard of aptness changes as science progresses. What counts as a maximally apt target for a particular purpose is fact of the matter and does not change. Of course, context is related to the scientific purpose for which the target system is being used. Yet this simply means that no target system is good or bad in itself, but apt for a certain purpose. Once the purpose is determined then context does not affect the *standard* of aptness, what counts as a maximally apt target system.

At the same time however, I do not wish to imply that contextual and pragmatic factors are not important. Even though the standard of target system evaluation is objective, as there is a fact of the matter about what is the maximally apt target for a particular scientific purpose, contextual and pragmatic factors affect our *judgements* of aptness. Each scientific purpose is associated with a certain type of technology, computational power, availability of data etc. All of these contextual factors affect our epistemic ability to determine whether a particular target is sufficiently apt. Thus, determining the context is an essential aspect for helping us get better at making judgements about the aptness of particular target systems, which is also an extremely important aspect of scientific practice, one which often needs philosophical examination in order to be better understood. In this type of philosophical analysis, it is not sufficient to simply state that context determines evaluation. It is possible to provide a much deeper and informative analysis, by showing the ways in which context determines the standard and method of evaluation. I hope to have shown that it is possible to achieve this at various levels of specificity and generality.

This concludes the final clarifications of my view of target systems, which I hope has shed some more light on their importance for scientists and philosophers alike. Still, there is a lot more work to be done, so that target system analysis reaches its full potential. This involves extending the investigation of target systems beyond ecological modeling, in two directions: beyond modeling and beyond ecology. The main area of interest in the dimension going beyond modeling, is experimentation. It is quite common in philosophy of science to distinguish between modeling and experimentation as two radically different aspects of science. This distinction has sparked a debate concerning the relative virtues of modeling versus experimentation. However, there has been a recent shift in focus, as some authors have started looking for commonalities between modeling and experimentation so that the two aspects of science can become integrated. I think that paying attention to the specification of target systems can help with this integration, as experiments also have targets. Just like a modeler, a successful experimenter must choose a particular framework in which to work and identify the factors relevant to the phenomenon being studied. Target systems and their evaluation can therefore help in the creation of successful experiments but also provide a common framework for understanding natural phenomena, whether they are studied through models or experiments.

An area of particular interest for experimentation is plant science. Invasion biology provides ample evidence that plants are special, as they are influenced and limited by different combinations of factors than other organisms, most importantly light, CO<sub>2</sub>, nitrogen, phosphorus and soil biota (Klironomos, 2002). This means that they respond differently to density and overcrowding (Yoda, 1963) and behave differently as competitors and invaders (Jarchow & Cook, 2009). At the same time, Another interesting distinction is often made between different types of experiments (laboratory, field and natural), and their relative merits (Diamond, 1983). According to Diamond, there is a tradeoff between realism and manipulability (control of variables), where lab experiments offer full control but are highly unrealistic, natural experiments are the opposite and field experiments are situated between the two extremes. Thus, each of these methods has virtues, weaknesses and limitations of scope, that vary with the subject of the experiment and the type of question being asked. I understand this investigation as a preliminary target analysis, which should be deepened and extended. Specifying an appropriate target system for a particular phenomenon will help determine which type(s) of experiment is best suited to it. This kind of analysis has an important advantage. The method of starting with the target system (instead of the type of experiment) is less prone to choose an



inappropriate framework, because of a desired value (for example, desiring an experiment that is realistic, even though it is extremely difficult to design a natural experiment).

Finally, another way to extend target system analysis is towards a different discipline, social science. Social scientists use models to investigate phenomena in the world, thus target systems are also present in social science. Being clear and precise about the targets of models and experiments in social science can support and reinforce investigations. There are two issues which can be illuminated by the investigation of target systems in the social sciences. The first is the ongoing debate concerning the optimal level at which societies and social phenomena should be studied, namely at the level of the individual or the group. I think that target system analysis can help illuminate the issue, by providing a principled way to distinguish between the types of questions which are best suited to individualism and those best suited to holism. The second issue is that of generalizability. Social scientists face the same problems as biologists when it comes to the scope and generalizability of their results. This is particularly important for the social sciences, as there is a recent debate concerning the nature and explanatory power of case studies in the social sciences (Baden-Fuller & Morgan, 2010; George & Bennett, 2005; Morgan & Morrison, 1999). As in the case of invasion biology, target systems can aid in the identification of the optimal level of generality which allows for both explanation and prediction.

Whichever direction target system analysis takes, one thing is certain. Target systems are important and ignoring their role in scientific practice will prevent a number of scientific and philosophical investigations from reaching their full potential.

## Bibliography

- Baden-Fuller, C., & Morgan, M. S. (2010). Business models as models. *Long Range Planning*.
- Bever, J. D. (2003). Soil community feedback and the coexistence of competitors: conceptual frameworks and empirical tests. *New Phytologist*, *157*(3), 465–473.
- Bird, A., & Tobin, E. (2008). Natural kinds.
- Blackburn, T. M., Pyšek, P., Bacher, S., Carlton, J. T., Duncan, R. P., Jarošík, V., et al. (2011). A proposed unified framework for biological invasions. *Trends in Ecology & Evolution*, *26*(7), 333–339. doi: 10.1016/j.tree.2011.03.023
- Brashares, J. S., Werner, J. R., & Sinclair, A. R. E. (2010). Social “meltdown” in the demise of an island endemic: Allee effects and the Vancouver Island marmot. *Journal of Animal Ecology*, *79*(5), 965–973. doi:10.1111/j.1365-2656.2010.01711.x
- Brigandt, I. (2013a). Explanation in Biology: Reduction, Pluralism, and Explanatory Aims - Springer. *Science & Education*.
- Brigandt, I. (2013b). Integration in biology: Philosophical perspectives on the dynamics of interdisciplinarity. *Studies in History and Philosophy of Biol & Biomed Sci*, *44*(4), 461–465.
- Buckley, Y. M., Bolker, B. M., & Rees, M. (2007). Disturbance, invasion and re-invasion: managing the weed-shaped hole in disturbed ecosystems. *Ecology Letters*, *10*(9), 809–817. doi:10.1111/j.1461-0248.2007.01067.x
- Callaway, R. M., Thelen, G. C., Rodriguez, A., & Holben, W. E. (2004). Soil biota and exotic plant invasion. *Nature*, *427*(6976), 731–733. doi:10.1038/nature02322
- Cartwright, N. (1994). Nature's Capacities and Their Measurement. *OUP Catalogue*.
- Catford, J. A., Jansson, R., & Nilsson, C. (2009). Reducing redundancy in invasion ecology by integrating hypotheses into a single theoretical framework. *Diversity and Distributions*, *15*(1), 22–40. doi: 10.1111/j.1472-4642.2008.00521.x
- Cooper, G. (1993). The competition controversy in community ecology. *Biology & Philosophy*, *8*, 359–384.
- Costa, N. A., & French, S. (2003). Science and partial truth: A unitary approach to models and scientific reasoning.
- Courchamp, F., Clutton-Brock, T., & Grenfell, B. (1999). Inverse density dependence and the Allee effect. *Trends in Ecology & Evolution*, *14*(10), 405–410.
- Daehler, C. C. (2003). JSTOR: Annual Review of Ecology, Evolution, and Systematics, Vol. 34 (2003), pp. 183-211. *Annual Review of Ecology*.
- Davidson, A. M., Jennions, M., & Nicotra, A. B. (2011). Do invasive species show higher phenotypic plasticity than native species and, if so, is it adaptive? A meta-analysis. *Ecology Letters*, *14*(4), 419–431. doi:10.1111/j.1461-0248.2011.01596.x

- Davis, M. A. (2011). Invasion biology 1958-2005: The pursuit of science and conservation. In D. M. Richardson, *Conceptual ecology and invasion biology: reciprocal approaches to nature* (pp. 35–64). Springer.
- Davis, M. A., Grime, J. P., & Thompson, K. (2000). Fluctuating resources in plant communities: a general theory of invasibility. *Journal of Ecology*, 88(3), 528–534.
- Davis, W. (n.d.). Implicature. *Plato.Stanford.Edu*.
- Diamond, J. M. (1983). Ecology: Laboratory, field and natural experiments. *Nature*, 304(5927), 586–587. doi:10.1038/304586a0
- Donnellan, K. S. (1966). Reference and Definite Descriptions. *The Philosophical Review*, 75(3), 281. doi: 10.2307/2183143
- Dunstan, P. K., & Johnson, C. R. (2007). Mechanisms of invasions: can the recipient community influence invasion rates? *Botanica Marina*, 50(5/6), 361–372.
- Dupré, J. (1995). *The disorder of things: Metaphysical foundations of the disunity of science*.
- Dupré, J. (2009). It is not possible to reduce biological explanations to explanations in chemistry and/or physics. *Contemporary Debates in Philosophy of Biology*, 32–47.
- Facon, B., GENTON, B., SHYKOFF, J., JARNE, P., ESTOUP, A., & DAVID, P. (2006). A general eco-evolutionary framework for understanding bioinvasions. *Trends in Ecology & Evolution*, 21(3), 130–135. doi:10.1016/j.tree.2005.10.012
- Frigg, R. (2009). Models and fiction. *Synthese*, 172(2), 251–268. doi:10.1007/s11229-009-9505-0
- Frigg, R., & Hartmann, S. (2009). Models in Science.
- Geng, Y.-P., Pan, X.-Y., Xu, C.-Y., Zhang, W.-J., Li, B., Chen, J.-K., et al. (2006). Phenotypic plasticity rather than locally adapted ecotypes allows the invasive alligator weed to colonize a wide range of habitats. *Biological Invasions*, 9(3), 245–256. doi:10.1007/s10530-006-9029-1
- George, A. L., & Bennett, A. (2005). Case studies and theory development in the social sciences.
- Giere, R. (2004). How Models are used to Represent Reality. *Philosophy of Science (Proceedings)*, 71(5), 742–752.
- Giere, R. N. (1988). *Explaining Science*. University of Chicago Press.
- Godfrey-Smith, P. (2008). Models and fictions in science. *Philosophical Studies*, 143(1), 101–116. doi: 10.1007/s11098-008-9313-2
- Godfrey-Smith, P. (2009). *Theory and Reality*. University of Chicago Press.
- Goldstein, P. Z. (1999). Functional Ecosystems and Biodiversity Buzzwords. *Conservation Biology*, 13(2), 247–255. doi:10.1046/j.1523-1739.1999.013002247.x
- Goodman, N. (1978). *Ways of Worldmaking*. Hackett Publishing.
- Gurevitch, J., Fox, G. A., Wardle, G. M., Inderjit, & Taub, D. (2011). Emergent insights from the synthesis of conceptual frameworks for biological invasions. *Ecology Letters*, 14(4), 407–418. doi:10.1111/j.1461-0248.2011.01594.x
- Hayes, K. R., & Barry, S. C. (2007). Are there any consistent predictors of invasion success? *Biological Invasions*, 10(4), 483–506. doi:10.1007/s10530-007-9146-5

- Hempel, C. G., & Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of Science*, 15(2), 135–175. doi:10.2307/185169?ref=search-gateway:0c1dbed607f660b0c1569ddf721ed555
- Higgins, S. I., Richardson, D. M., & Cowling, R. M. (1996). Modeling invasive plant spread: the role of plant-environment interactions and model structure. *Ecology*, 2043–2054.
- Hull, D. L. (1965). The effect of essentialism on taxonomy--two thousand years of stasis (I). *British Journal for the Philosophy of Science*, 314–326.
- Hulme, P. E. (2007). Phenotypic plasticity and plant invasions: is it all Jack? *Functional Ecology*, 22(1), 3–7. doi:10.1111/j.1365-2435.2007.01369.x
- Humair, F., Edwards, P. J., Siegrist, M., & Kueffer, C. (2014). Understanding misunderstandings in invasion science: why experts don't agree on common concepts and risk assessments. *NeoBiota*, 20, 1–30. doi:10.3897/neobiota.20.6043
- Humphreys, P. (1995). Abstract and concrete. *Philosophy and Phenomenological Research*, 55(1), 157–161.
- Jarchow, M. E., & Cook, B. J. (2009). Allelopathy as a mechanism for the invasion of *Typha angustifolia*. *Plant Ecology*, 204(1), 113–124. doi:10.1007/s11258-009-9573-8
- Jeschke, J., Gómez Aparicio, L., Haider, S., Heger, T., Lortie, C., Pyšek, P., & Strayer, D. (2012). Support for major hypotheses in invasion biology is uneven and declining. *NeoBiota*, 14(0), 1–20. doi:10.3897/neobiota.14.3435
- Jones, M. R. (2005). Idealization and abstraction: A framework. In M. R. Jones & N. Cartwright, *Idealization XII: Correcting the model. Idealization and abstraction in the sciences*. (pp. 173–217). Amsterdam Rodopi: *Idealization XII: Correcting the model Idealization ....*
- Joost. (2003). Von et al 2003, 1–8.
- Kennedy, T. A., Naeem, S., Howe, K. M., Knops, J. M. H., Tilman, D., & Reich, P. (2002). Biodiversity as a barrier to ecological invasion. *Nature*, 417(6889), 636–638. doi:10.1038/nature00776
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science*, 507–531.
- Kitcher, P. (1984). 1953 and all That. A Tale of Two Sciences. *The Philosophical Review*, 93(3), 335. doi:10.2307/2184541
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. *Scientific Explanation*, 13, 410–505.
- Kitcher, P. (1993). *The Advancement of Science : Science without Legend, Objectivity without Illusions*. Oxford University Press.
- Kitcher, P. (1999). Unification as a regulative ideal. *Perspectives on Science*, 7(3), 337–348.
- Kitcher, P. (2003). *Science, Truth, and Democracy*. Oxford University Press.
- Kitcher, P. (2013). Toward a Pragmatist Philosophy of Science. *Theoria. Revista De Teoría, Historia Y Fundamentos De La Ciencia*, 28(2), 185–231.
- Klironomos, J. N. (2002). Feedback with soil biota contributes to plant rarity and invasiveness in communities. *Nature*, 417(6884), 67–70.

- Knuuttila, T. (2005). Models, Representation and Mediation. *Philosophy of Science (Proceedings)*, 72(5), 1260–1271.
- Korta, K., & Perry, J. (n.d.). Pragmatics. *The Stanford Encyclopedia of Philosophy*, (Winter 2012).
- Kuhn, T. S. (1996). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Lake, J. (2004). Invasion success of exotic plants in natural ecosystems: the role of disturbance, plant attributes and freedom from herbivores. *Biological Conservation*, 117(2), 215–226. doi:10.1016/S0006-3207(03)00294-5
- Leonelli, S. (2013). Integrating data to acquire new knowledge: Three modes of integration in plant science. *Studies in History and Philosophy of Biol & Biomed Sci*, 44(4), 503–514.
- Levine, J. M., & D'Antonio, C. M. (1999). Elton Revisited: A Review of Evidence Linking Diversity and Invasibility. *Oikos*, 87(1), 15. doi:10.2307/3546992
- Lodge, D. M. (1993). Biological invasions: lessons for ecology. *Trends in Ecology & Evolution*, 8(4), 133–137.
- Lucas, J. F. (1990). *Introduction to Abstract Mathematics*. Rowman & Littlefield.
- Maclaurin, J., & Sterelny, K. (2008). What Is Biodiversity? - James Maclaurin, Kim Sterelny - Google Books.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science Part A*, 16(3), 247–273. doi:10.1016/0039-3681(85)90003-2
- Mitchell, S. D. (2003). Biological complexity and integrative pluralism.
- Moles, A. T., Flores-Moreno, H., Bonser, S. P., Warton, D. I., Helm, A., Warman, L., et al. (2011). Invasions: the trail behind, the path ahead, and a test of a disturbing idea. *Journal of Ecology*, 100(1), 116–127. doi:10.1111/j.1365-2745.2011.01915.x
- Morgan, M. S., & Morrison, M. (1999). Models as mediators: Perspectives on natural and social science.
- Odenbaugh, J. (2007). Seeing the forest and the trees: Realism about communities and ecosystems. *Philosophy of Science*, 74(5), 628–641.
- O'Malley, M. A. (2013). When integration fails: Prokaryote phylogeny and the tree of life. *Studies in History and Philosophy of Biol & Biomed Sci*, 44(4), 551–562.
- Plutynski, A. (2013). Cancer and the goals of integration. *Studies in History and Philosophy of Biol & Biomed Sci*, 44(PA), 466–476. doi:10.1016/j.shpsc.2013.03.019
- Psillos, S. (1999). Scientific realism: How science tracks truth.
- Railton, P. (1981). Probability, Explanation, and information. *Synthese*, 48(2), 233–256. doi:10.2307/20115660?ref=search-gateway:55ccbe36a9db3c3483acbc695dd09843
- Reichard, S. H., & White, P. S. (2003). Invasion biology: an emerging field of study. *Annals of the Missouri Botanical ...*
- Rejmánek, M., & Richardson, D. M. (1996). What attributes make some plant species more invasive? *Ecology*, 77(6), 1655–1661.
- Richardson, D. M. (2011). *Fifty years of invasion ecology: the legacy of Charles Elton*. John Wiley & Sons.

- Richardson, D. M., & Pyšek, P. (2006). Plant invasions: merging the concepts of species invasiveness and community invasibility. *Progress in Physical Geography*, 30(3), 409–431. doi:10.1191/0309133306pp490pr
- Richardson, D. M., & Rejmánek, M. (2004). Conifers as invasive aliens: a global survey and predictive framework. *Diversity and Distributions*, 10(5-6), 321–331.
- Ricklefs, R. E., & Miller, G. L. (2000). *Ecology*. Macmillan.
- Romanuk, T. N., Zhou, Y., Brose, U., Berlow, E. L., Williams, R. J., & Martinez, N. D. (2009). Predicting invasion success in complex ecological networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1524), 1743–1754. doi:10.1073/pnas.192448799
- Rosen, G. (2012). Abstract Objects. *The Stanford Encyclopedia of Philosophy*.
- Rosenberg, A. (2008). *Darwinian Reductionism*. University of Chicago Press.
- Salmon, M. H. (1999). *Introduction to the Philosophy of Science*. Hackett Publishing.
- Salmon, W. C. (2006). *Four Decades of Scientific Explanation*. University of Pittsburgh Pre.
- Schulze, E.-D., Beck, E., & Müller-Hohenstein, K. (2005). *Plant Ecology*. Springer.
- Shea, K., & Chesson, P. (2002). Community ecology theory as a framework for biological invasions. *Trends in Ecology & Evolution*, 17(4), 170–176.
- Sobek-Swant, S., Kluza, D. A., Cuddington, K., & Lyons, D. B. (2012). Potential distribution of emerald ash borer: What can we learn from ecological niche models using Maxent and GARP? *Forest Ecology and Management*, 281, 23–31. doi:10.1016/j.foreco.2012.06.017
- Stanford, P. K. (1995). For pluralism and against realism about species. *Philosophy of Science*, 70–91.
- Stephens, P. A., & Sutherland, W. J. (1999). Consequences of the Allee effect for behaviour, ecology and conservation. *Trends in Ecology & Evolution*, 14(10), 401–405. doi:10.1016/S0169-5347(99)01684-5
- Sterelny, K., & Griffiths, P. E. (2012). *Sex and Death*. University of Chicago Press.
- Strevens, M. (2004). The causal and unification approaches to explanation unified—causally. *Noûs*, 38(1), 154–176.
- Strevens, M. (2008). *Depth*. Harvard University Press.
- Suarez, M. (2003). Scientific Representation: Against Isomorphism and Similarity. *International Studies in the Philosophy of Science*, 17(3), 225–244.
- Suppes, P. (1969). Models of data.
- Taylor, C. M., & Hastings, A. (2005). Allee effects in biological invasions. *Ecology Letters*, 8(8), 895–908. doi:10.1111/j.1461-0248.2005.00787.x
- Van Fraassen, B. C. (1980). *The Scientific Image*. Clarendon Press.
- van Fraassen, B. C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford University Press.
- Van Kleunen, M., Dawson, W., Schlaepfer, D., Jeschke, J. M., & Fischer, M. (2010). Are invaders different? A conceptual framework of comparative approaches for assessing determinants of invasiveness. *Ecology Letters*, no–no. doi:10.1111/j.1461-0248.2010.01503.x

- Weidenhamer, J. D., & Callaway, R. M. (2010). Direct and Indirect Effects of Invasive Plants on Soil Chemistry and Ecosystem Function. *Journal of Chemical Ecology*, 36(1), 59–69. doi:10.1007/s10886-009-9735-0
- Weisberg, M. (2007a). Three kinds of idealization. *The Journal of Philosophy*, 639–659.
- Weisberg, M. (2007b). Who is a Modeler? *The British Journal for the Philosophy of Science*, 58(2), 207–233. doi:10.1093/bjps/axm011
- Weisberg, M. (2013). *Simulation and Similarity*. OUP USA.
- Weisberg, M., Needham, P., & Hendry, R. (2011). Philosophy of Chemistry.
- Williamson, M. (1999). Invasions. *Ecography*, 22(1), 5–12.
- Wimsatt, W. C. (1974). Complexity and organization, 67–86.
- Woodward, J. (2009). Scientific Explanation.
- Yoda, K. (1963). Self-thinning in over-crowded pure stands under cultivated and natural conditions.(Intraspecific competition among higher plants.XI.). *J Biol Osaka City Univ*, 14, 107–129.