

Appeared as Chapter 3 in A. Cassini and J. Redmond, eds. (2021),
Idealizations in Science: Artifactual and Fictional Approaches, Springer, pp. 71-85.

Informative Models: Idealization and Abstraction

Mauricio Suárez
Department of Logic and Theoretical Philosophy,
Complutense University of Madrid
msuareza@ucm.es

and

Agnes Bolinska,
Department of Philosophy,
University of South Carolina
bolinska@mailbox.sc.edu

Final 2021 Version

7,288 words

Abstract: Mauricio Suárez and Agnes Bolinska apply the tools of communication theory to scientific modelling in order to characterize the informational content of a scientific model. They argue that when represented as a communication channel, a model source conveys information about its target, and that such representations are therefore appropriate whenever modelling is employed for informational gain. They then extract two consequences. First, the introduction of idealizations is akin in informational terms to the introduction of *noise* in a signal; for in an idealization we introduce ‘extraneous’ elements into the model that have no correlate in the target. Second, abstraction in a model is informationally equivalent to *equivocation* in the signal; for in an abstraction we ‘neglect’ in the model certain features that obtain in the target. They then conclude that it becomes possible in principle to quantify idealization and abstraction in informative models, although precise absolute quantification will be difficult to achieve in practice.

Keywords: Information – Idealization – Abstraction - Content

1. Introduction

Scientific models are often employed to gain information regarding their targets. The building of the model is guided by preliminary knowledge of some phenomenon, and the model aims to provide further information regarding unknown aspects of the phenomenon, or its underlying causes. This is evident in most models with predictive power that we know, but it applies to most other models whenever some cognitive gain in understanding is sought (for compilations of an array of case studies in physics, economics, and other sciences, see e.g., Morgan and Morrison, 1999; Jones and Cartwright, 2005; Suárez, 2009). A key feature of scientific representations is their capacity to enable users to draw informative inferences from what we may call representational *vehicles*, or *sources*, to their *target* systems. This has been studied in depth and is by now widely accepted (Boesch, 2017; Bolinska, 2013; Contessa, 2007; Suárez, 2004).¹ What is not so established is the nature of the information that is conveyed by models. In this paper we focus on the nature of the information provided, and apply the tools of communication theory in order to draw a few interesting lessons regarding some approximation techniques typical in scientific modelling. When a model is represented as analogous to a communication channel, the sources of informational *noise* and *equivocation* have correlates in different forms of idealization and abstraction in modelling practice. This sheds some light on some common methods for minimizing idealization and abstraction, as well as their rationale.

The relevant information theoretical concepts are introduced in section 2: They originate in Shannon's classic (1948) and have been discussed in an epistemological context by Dretske (1981).² The central analogy between communication channels and scientific models is laid out in section 3. The following section 4 expounds on a case study within the kinetic theory of gases, and it argues for a role for informational noise and equivocation in standard understandings of idealization and abstraction. The concluding section 5 wraps up the main claim, and raises some questions prompted by the analogy that deserve further exploration.

2. The Mathematical Theory of Communication (MTC)

In a communication system, a message travels from an emitting *source* to a *receiver* through some medium, such as a radio signal. Communication channels are *noisy*: part of

¹ So is the concomitant terminology of 'sources' and 'targets' (Suárez, 2004); yet, for reasons that will become apparent, it is best for the purposes of this article to refer to representational *vehicles* and *targets*, to aptly distinguish them from the terms employed in information theory.

² See Bolinska (2015) for an application of these notions in the philosophy of science.

the message can be obscured, for instance, through crosstalk in the radio signal, interference, or a random noise generator mixing into the signal. On the other hand, there is always less than perfect quality in the transmitted message, i.e., every signal suffers from loss, also known as *equivocation*, due to impurities in the transmission channel, or in the coding and decoding of the message. In other words, no communication channel is ever 100% efficient. The goal of effective communication is thus rarely, if ever, to completely eliminate or eradicate the inefficiencies in the form of either noise or equivocation – since to bring those inefficiencies to 0% in practice is an impossible task. Rather the goal of effective communication is to maximise the efficiency of the signal within the bounds of what is in fact possible for any given channel. For a complete array of possible signals, every channel will have some limit to what is capable of transmitting from emission to reception. The limit is the average efficiency of the channel. In other words, the informational efficiency of any communication channel is an average property of the channel relative to all possible transmissions through that channel. The goal of communication efficiency is to maximize this quantity – and it entails choosing a channel with as great a ratio of information to noise, or equivocation, as is possible for transmissions of information from a given source to a receiver.

A communication system comprises minimally five separate parts: an emitting source; an encoder or transmitter; a signal; a decoder or receptor; and a receiver. (See figure 1: Shannon’s information theory). The source possesses or generates certain quantitative properties, which the encoder or transmitter (some sort of machine or mechanism) codifies in a signal. The signal carries the information over to a decoder that extracts the relevant information regarding the source in a form that is adequate for the purposes of a receiver. In practice there is a degree of information loss at every stage (e.g., Pierce, 1961): the source properties may not all get adequately codified, or not codified at all in the signal. The signal may lose some of its properties or resolve. The decoding may be deficient and fail to extract all the information in the source. The receiver may in principle be inefficiently geared towards the information received. In other words, there are many sources of what we call *equivocation*: relevant information regarding the source that is not transmitted over to the receiver.

Dretske (1981, pp. 16ff.) represents equivocation formally as follows. Let us refer to the source as s and the receiver as r . And let us refer to the total amount of information contained in s as $I(s)$, and to the total amount of information received at r as $I(r)$. Then we can denote the total amount of information about s that is received at r as $I_s(r)$. A straightforward measure of equivocation is then given as:

$$E(r) = I(s) - I_s(r) \quad \text{[Equivocation]}$$

That is, the equivocation of a communication system is the amount of information that gets lost in the transmission from the source to the receiver, i.e., the amount of information generated at the source that fails to be transmitted. Dretske’s formula applies to average equivocation in a particular communication channel between a given source s and a given receiver r . It is clear that in order to compute it quantitatively we need to possess measures of the information contained in the source $I(s)$, and of the part of this information that is in effect transmitted to the receiver, $I_s(r)$. Dretske points out that in communication theory we can quantify the information contained in any system by calculating the reduction in the number of possible states of the system, as: $I(s) = \log n$, where n is the number of possibilities that get reduced to 1. (This is sometimes known as the information entropy of the source, and assumes that the source is some kind of process or phenomenon endowed with some dynamics that reduce a large space of possibilities into one – an assumption that I shall return to later on in the discussion of the case study).

Reducing the equivocation in a channel increases completeness – a channel that equivocates a lot provides an impoverished rendition at reception of the qualities of the source. Dretske suggests (1981, p. 25) that we can calculate the equivocation in a communication system as follows. Suppose that there are only eight possible events at the source $\{s_1, s_2, \dots, s_8\}$, and correspondingly eight possible events at the receiver $\{r_1, r_2, \dots, r_8\}$. Now, suppose that the signal is such as to generate each event at the receiver with a given probability given each event at the source. That is, there are well defined values of $P(r_i / s_j)$ for every couple $\{r_i, s_j\}$. We may focus on a particular event, say r_7 , and work out the equivocation for that event as follows: $E_s(r_7) = - \sum P(s_i / r_7) \log P(s_i / r_7)$. This is just the probabilistically weighted average of the equivocation for each of $\{s_i\}$ with respect to r_7 . To compute the average equivocation for the channel we need only sum up the contributions made to it by each of the events, weighted by their probabilities:

$$E(r) = \sum_j P(r_j) E(r_j) \quad \text{[Average Equivocation]}$$

In other words, the equivocation of a communication channel is an average quantity computed over each of the possible values of the properties of the source, weighted by the conditional probability that each of those values generates a particular value of some property in the receiver. We calculate the equivocation for each value of the receiver property by estimating the probability that information loss may occur for this value of the receiver property. And then we sum over all the values weighted by their corresponding probability. Whenever applicable, the procedure yields quantitative values for the equivocation and this constitutes a measure of the channel’s efficiency. The larger the equivocation, the larger share of information is lost in the transmission from source to receiver, since: $I(r) = I(s) - E(r)$. The smaller the equivocation, by contrast, the larger proportion of the information contained in the source is transmitted to the receiver. At the limit, when the equivocation is zero, all the information contained at the source is completely transmitted: $I(s) = I_s(r)$.

The other important source of inefficiency in a communication channel is *noise*. Roughly, the noise of a communication channel is whatever extraneous information is picked up, and thus added to the signal as it travels from source to receiver. It may be added at the stage of coding, e.g., at the source, or it may get added later on during the transmission or at the stage of decoding at the receiver's end. At any rate, the noise in a communication channel is defined as whatever information is transmitted to the receiver which does not originate in the source:

$$N(s) = I(r) - I_s(r). \quad \text{[Noise]}$$

Hence the greater the noise the larger share of the information received was not actually generated at the source; the smaller the noise the greater proportion of the information did actually originate at the source. At the limit where noise is null, all the information received faithfully originates at the source: $I(r) = I_s(r)$.

Reducing the noise over a communication channel increases faithfulness, or reliability. We can calculate the average noise over a communication channel in converse fashion to equivocation as follows. First calculate noise for every possible event in the source. Thus, for instance, for event s_j at the source, its contribution to the average or overall noise is given as: $N(s_j) = -\sum_i P(r_i/s_j) \log P(r_i/s_j)$. Then the noise of the channel is simply the statistical average of each of these contributions, i.e. the contribution of each weighted according to its probability:

$$N(s) = \sum_j P(s_j) N(s_j). \quad \text{[Average Noise]}$$

In many cases, noise detracts from signal transmission, preventing some of the information from the source to be transmitted to the receiver, i.e., increasing equivocation. But this need not be the case: it is possible for the receiver to contain additional information, information that didn't originate in the source, which doesn't detract from signal transmission. Thus, an increase in noise need not logically or conceptually imply a corresponding increase in equivocation (Dretske 1981, pp. 20-21).

Most communication channels operate on signals of much greater complexity, with a much larger and more complex space of possibilities at the source than here described. But the basic notions stand, and they will suffice for the purposes of this paper. It is particularly relevant that informational quantities like equivocation and noise are averages and therefore properties of the communication channel, not of any particular signal. While

this averaging feature of informational notions is not particularly suited to Dretske's purposes (he chooses to concentrate instead on concrete signalling actions), it is well suited for the purposes of the analogy with scientific modelling that shall be explored here. A complex communication channel comprises five elements, and its most relevant informational quantities must be computed as averages pertaining to the channel as a whole.

3. An Analogy: Models as Communication Channels

The practice of model building displays considerable sensitivity to some notion of fit between a representational vehicle and its target. The fit, however, very rarely involves matching pairwise structural properties; more often than not what is at stake is how relevantly *informative* the source is as an inferential whole about certain aspects of the target.³ However, the notion of information at work here is suitably thin, for any kind of model. A richly informative model describes its target in great detail; yet not all the detail need be a guide to what the target really is like – sometimes the detail serves pragmatic purposes in understanding or predicting the ensuing phenomena. Many models idealise their targets a great deal, presupposing point mass particles, frictionless planes, and the like. These are detailed and are in some sense informative, but for all we know they are not faithful to their targets. By contrast, a powerful or deep model may not be rich in detail but is nonetheless informative in some sense about the main or critical characteristics, those that are more central to the production of the phenomena. What is the sense of informativeness that is involved in these models? From an inferential point of view, the richer model is most informative in the sense that it licences a large number of inferences to many different aspects of the phenomenon. A deep model may licence fewer inferences, and to fewer aspects of the phenomenon, but they are inferences to aspects that we have reason to suppose are more central, including some of its putative causes.

In other words, models encode information. The kind of information will differ in different types of model, and the nature of the information involved will also differ depending on the account of representation endorsed. But in any case, information transmission is often a main aim of modelling.⁴ This observation suggests an analogy: the target of the model acts often as an information source for the representational vehicle to

³ In fact, a 1-1 copy of the target – as in Borges' (1954) beautiful parable – would be a 100% informative 'channel' and, yet, a perfectly useless model.

⁴ Indeed, one of us (Bolinska, 2016) has explicitly understood models as tools for conveying information, identifying features responsible for their informativeness. See also the discussion of 'objectivity' in Suárez (2004, 2024).

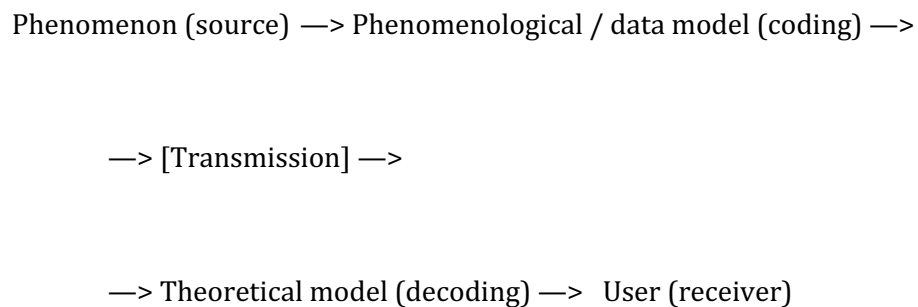
transmit to the user of the model. In other words, modelling often is in some sense a communication practice, and a model may be thought of, at some level of abstraction, as a communication channel. It could help understand the practice of modelling to make it more precise what this sense is.

Nevertheless, a cogent application of the analogy meets some challenges. In a communication channel as described in Shannon's theory (figure 1), the source's information is first encoded into a transmitter, which emits a signal. The signal is carried by some means to a receiver, which decodes the information and provides it over to its destination. There are thus five objects laid out in a communication system, and we need to fix on the respective analogues in the case of modelling if Shannon's theory is to apply. Now, we are not suggesting that models are just communication channels, but only that there are helpful analogies that may be exploited to better understand modelling practice. There may well be different ways of helpfully laying out the analogy, depending on both types of models, and the underlying account of representation. In other words, we make no claim that the proposal advanced here will always be applicable or helpful, in every instance of modelling, regardless of how the case is understood or interpreted. And we are certainly not claiming that the analogy presented here exhausts everything that may be claimed about modelling practice, even from an informational point of view. Our more modest claim is that the analogy provides some tools to better understand the kind and nature of the information transmission that takes place in some instances of modelling, as pertains their inferential function in particular.

The proposal is thus to treat the system, process, or phenomenon of interest (the representational "target") as the informational source in a communication channel. The system or phenomenon of interest is typically a dynamical process, or it involves one, and it is often represented to us already in some preliminary or antecedent description. (The apparent circularity is a well-known issue in the representation literature, which there is no space to broach here – see Van Fraassen, 2008, Ch. 11 and part IV, for further discussion). The information about this dynamical process is then *encoded* in at least either of two ways. It can be studied empirically, on the basis of the static data that it elicits in some experimental trial, and the resulting information can be built into what is known as a data-model (in the sense of e.g., Suppes, 1962). Or it can be modelled dynamically, in terms of a given parameter set, in what is a phenomenological model. Different ways of encoding the same information may be more or less appropriate for the purposes at hand.

In either case, the model then acts as a courier of information, a tool to compactly convey codified information regarding the system of interest. At the other end of the process, the model needs to be interpreted in ways that make the information salient for the purposes of prediction, understanding, explanation, or generalisation. This will typically require a scientist to employ a theory or sets of theories (sometimes high-level theories, such as, in physics, the kinetic theory of gasses, the theory of general relativity, or quantum

mechanics; other times medium-level theory, such as diffraction optics, radiation theory, or the molecular hypothesis for gases) to interpret the phenomenological or data model. The information at the source is thus finally transmitted over to its final destination, the model user. It is important to note the multiple steps involved in this communicative act, and the concomitant judgements along the way: The source has to be competently and aptly described, the information it contains must then be codified / transposed into an appropriate model that will act as an information-carrier, and this model must be correctly interpreted in the light of some theoretical knowledge for the information to be relevant, comprehensive and / or apt for the purposes of the model user. The overall communication channel may be schematically described as follows:



In this picture the communication system also has five steps, the middle transmission step being the immersion of the phenomenological description or model into the theoretical explanation or interpretation of the phenomenon. Information loss, in the shape of either noise or equivocation, is possible in every step in this chain. There is bound to be information loss in the choice of data or phenomenological model (here understood in informational terms as the selection of the coding system); in the actual embedding of the data or phenomenon into the theoretical description (understood informationally as the transmission of a signal); and in the choice of the theoretical model that interprets the phenomenon (the choice of the information decoding system). There are no doubt judgements in all three cases as to what is most likely to preserve the largest amount of information or the most critical kind of information. But such judgements are also involved in any communication channel (for some insight into the kinds of practical choices those working in information theory have to routinely make, see MacKay, 2003). All these choices must be fit for purpose, so a lot will depend on the actual goals pursued by the user of the model – and these may vary greatly depending on the context of use.

The kind of contextuality of use involved in modelling practice is by now widely accepted (cf. Bailer-Jones, 2003; Giere, 2004, 2009; Mäki, 2009; Teller, 2001), and it cannot be easily algorithmically or automatically done away with, if at all. This severely constrains the analogy in at least two different ways – which explain why modelling practice cannot be simply reduced to the building of effective communication channels. In a communication

system, the goal is for the message at the receiver end to identically reproduce the message at the source, or at least to do so with minimal information loss. Yet, as noted, a scientific model rarely aims to reproduce the target system in its entirety exactly. More often than not a scientific model aims to capture certain consequences of central features of the target system – those consequences that are of importance for the purposes of prediction, explanation, generalization, etc.

There are at least two forms of helpful ‘distortion’ in modelling practice that need to be explained from an informational point of view. They correspond roughly to *abstraction*, and *idealization*.⁵ Consider, for instance, the oft-discussed example of the simple harmonic oscillator as a model of a pendulum subject to no friction (Giere, 1988). There are two ways to understand the absence of friction in these models. The first one assumes the model is an abstract rendition of the phenomenon that ignores some of the complexity: the real target phenomenon possesses friction, but the model does not. In informational terms, this means that some of the information available at the source is not represented in the model of the phenomenon in any way that can be interpreted theoretically as friction. As a result, the information does not get transmitted over to the receiver. Thus, in informational terms, there is ‘equivocation’ involved in the simple harmonic oscillator model of a real pendulum.

Now consider the alternative reading of “frictionless” according to which the simple harmonic oscillator model is an idealization. On this reading the model includes a property (“lack of friction”, or “frictionlessness”) that the phenomenon that it models does not in fact possess. The model idealizes the phenomenon by introducing properties that are not there in the source phenomenon. This understanding is of interest too, and rather typical in modelling: it entails introducing properties in the model description for ease of calculation, manipulation, prediction, etc. The “frictionlessness” of the ideal pendulum is informationally akin to ‘noise’ – it is present in the model yet does not originate in the source but is “extraneous”. One can thus see, following these two interpretations of “frictionless,” that abstraction will typically be analogous in informational terms to ‘equivocation’, while idealization is analogous to informational ‘noise’.

It is clear that a lot in the building and applying of models depends on controlling and monitoring both “equivocation” and “noise.” As regards equivocation, reducing its presence in a model is essential to the aim of comprehensiveness – a model with 100% equivocation is a model where none of the information originating in the source is transmitted ($I_s(r) = 0$) and is hence perfectly useless. As regards noise, a model is helpful only in as much it transmits faithfully the features of the source. Otherwise we run the risk

⁵ Our use of the terms is inspired by the definitions in Cartwright and Jones (2008), and Weisberg (2012); and is reflected most precisely in those used in Pero and Suárez (2016).

of incorrectly ascribing to the phenomenon properties that it does not possess - a model where 100% of the information received is 'extraneous' is again useless, since $I_s(r) = 0$, i.e. none of the information received originates in the source.

4. Modelling and the Transmission of Information: Some Examples

One of the paradigmatic case studies in the contemporary literature is the billiard ball model of gases. It is first treated as part of an extensive discussion of the kinetic theory of gases in Campbell (1920) and thereafter in great depth in Mary Hesse's masterly (1963).⁶ In the kinetic theory of gases, the dynamics of molecules in a gas is modelled as if it were a system of perfectly elastic microscopic balls in collision - a set of miniature 'billiard balls' in constant motion. Let us refer to the real properties of actual gas molecules as $\{G_1, \dots, G_n\}$ and those of billiard balls as $\{B_1, \dots, B_m\}$. Then the model sets up a correspondence between a subset $\{\{G_1, \dots, G_i\}, \text{ with } i < n\}$ of the properties of gas molecules and a subset $\{\{B_1, \dots, B_i\}, \text{ with } i < m\}$ of the properties of billiard balls. We say that gas molecules are billiard balls as regards their collision dynamics, but this does not mean that they share all the properties of billiard balls. Billiard balls are coloured and shiny and reflect light, but gas molecules possess none of these properties. Even as regards their dynamical properties there are significant differences. Billiard ball motion is subject to limited friction against the surface on which they move, while gas molecules presumably interact freely.⁷ And conversely, a system of gas molecules exhibits macro-properties that no system of billiard balls can ever display, such as viscosity or free expansion.

In other words, the model omits certain properties of gas molecules, while including others that molecules don't in fact have. In accordance with the analogy laid out in this paper, the model may for the purposes of information be taken to be a communication channel transmitting information from the source - a gas - encoded and transmitted in accordance with some phenomenological model (in the form of the billiard ball model), interpreted in the terms of kinetic theory of gases, for the sake of the receiver's information (the physicist able to interpret or decode the information in the signal). Properties of the gas that are ignored in the model will contribute equivocation in the signal, impoverishing

⁶ More recently, Pero and Suárez (2016), Suárez and Pero (2019) and Suárez (forthcoming) contain an extended historical discussion of this example that corrects some philosophical misconceptions.

⁷ It is actually worse than that: contrary to assumptions billiard balls are not perfectly elastic, but of course experience minor energy loss in the form of heat in collisions; but then again neither are gas molecules perfectly elastic, since they also experience loss of (kinetic) energy in collision. See the discussion in Pero and Suárez, 2016, pp. 75-76.

the communication, while features in the model that do not correspond to properties of the gas will contribute noise. In line with the discussion in the previous section, while we would ideally reduce both noise and equivocation to zero, if we could, their introduction turns out to be unavoidable for any practicable model at all.

The main difficulty in applying the informational framework to the billiard ball model is in laying down the appropriate probability distributions. We shall have to make a number of assumptions at this point – including some prior distributions for the different combinations of values. Fortunately, however, the assumptions are warranted by the physical model itself as well as the philosophical discussions regarding the billiard ball model.⁸ We shall therefore assume a normal distribution over the values of the properties of the source system (the gas molecules). This is warranted by the physics, since the standard assumption regarding the velocity of free molecules in a gas is that they are distributed in accordance to Maxwell-Boltzmann statistics. To be more precise this says that, in a vessel containing m particles in equilibrium, the proportion of particles with a particular velocity v is n (where $n < m$) and is given by the Maxwell-Boltzmann distribution as:

$$f(v) = \sqrt{\left(\frac{m}{2\pi kT}\right)^3} 4\pi v^2 \exp\left(\frac{-mv^2}{2kT}\right) \quad [\text{Maxwell-Boltzmann distribution}]$$

This probability distribution function depends only upon the initial velocities of the gas molecules, since the so-called Boltzmann constant k and the thermodynamic temperature of the gas T are both constants of motion. We shall assume that all these velocities have correlative properties in the velocities of the billiard balls in a system of billiard balls. However, there is no reason in principle why the velocities of an equally large group of billiard balls should also obey a Maxwell-Boltzmann distribution. In fact, the notion of equilibrium itself makes no sense for billiard balls. We may assume by fiat that the set of billiard balls in our model obeys the Maxwell-Boltzmann distribution, and this is indeed commonly done. What this means, from the point of view of information theory, is that we assume that there is no information loss in the description of the molecules' velocities in the billiard ball model.

But, of course, these are not the only properties. Let us begin with all the noise introduced into the system: the properties of billiard balls that are irrelevant to the kinetic

⁸ Although, certainly, these assumptions can be contested. But nothing much hinges on the particular values. The only claim that needed for the present proposal to go through is that there are some values for these probability distributions – whether they are not within our reach to know is not essential. As Dretske points out (ibid, p. 55) the probability distributions that go into communication theory, and in particular the conditional probabilities in the definition of equivocation, are objective, and may be very hard to get to know.

theory of gas, such as their colour, their shine, and of course their rigid solid structure. Take colour, and assume for the sake of argument that billiard balls can be any of seven colours from the deep red to the violet end of the spectrum, and that each colour is as likely as any other (this is just an assumption about the entities in the model, for which it is not relevant whether or not it contradicts standards in e.g. ordinary sets of billiard balls in the game of pool!). This already entails that, in the equation for the average noise in the channel, the value of the probabilistic distribution is uniform over all the colours:

$$N(s) = \sum_j P(s_j) N(s_j), \text{ is such that } P(s_i) = P(s_j) = a \text{ for all } i, j.$$

But now, a further substantial assumption must be made, namely that the colour of the balls is not correlated with any of the properties of the gas molecules that they represent. The assumption seems intuitive, and it would be a strange model that correlated the properties across in this way, but it is not in principle impossible model.⁹ So we shall just have to assume that the model does not work that way and that the colour of the balls is completely uninformative with respect to any of the physical properties of the gas molecules. There is no correlation. This entails that the cross or conditional probabilities in the expression for the noise are equal, and the conditional probability distribution is flat: $P(r_i / s_k) = P(r_j / s_k) = b$, for any i, j , and any property s_k of the source.

Now, once we have established that the probabilities are constant numbers across the average, we can easily see that the contribution to the noise for each event at the source is: $N(s_k) = -\sum_i P(r_i / s_k) \log P(r_i / s_k) = -b \sum_i \log P(r_i / s_k)$. But as we had already found out the noise to be a constant of the average of the noise contribution from each value, we obtain: $N(s) = a \sum_j N(s_j) = a \cdot b \sum_{ik} -\log P(r_i / s_k)$. In other words, the noise is just a constant function of each of the conditional events. If those are zero, then the noise goes down to zero and the signal achieves maximal efficiency.

What does this mean for our billiard ball model? It means that the information transmitted by the model about the system of gas molecules is only as efficient as the 'noise' by spurious variables in the models is low – and this depends only on how much every one of the possible values of any of the spurious variables is correlated with the relevant physical variables in the source. Bear in mind that the correlation is objective, so even if we lack any knowledge – even if we assume there to be no correlation – the actual noise in the signal depends on the existence of the correlations, independently of our knowledge. So, we

⁹ For instance, one could imagine a model where the colour of the elastic 'billiard' balls is taken to represent the initial velocity of each corresponding molecule in the gas, with purple representing the larger speeds and red representing the smaller speeds and all the other colours representing intermediate speed ranges in the prescribed order in the electromagnetic visible spectrum. Such model would not be very useful, but it is perfectly possible, for any given gas.

are never in a position to rule out informational noise (because we can never completely de-idealize the model).

Let us now consider equivocation. Very similar considerations will apply, even though we are now interested in the properties in the source that are ignored or abstracted away in the model. As mentioned, all the macroscopic properties of the gas (free expansion, viscosity) are derived from theory on the basis of the model, but do not appear in the analogy itself: systems of billiard balls exhibit neither viscosity nor free expansion. However, the same reasoning we applied to noise will also apply to the equivocation function: $E(r) = \sum_j P(r_j) E(r_j)$, even if equivocation depends upon the probabilities of the receiver, not the source. If the model is deterministic, we can assume all the probabilities to be zero or one for all the values of all the dynamical variables of interest.¹⁰ Then the equivocation depends only on the contribution made by the value that actually obtains, suppose r_7 . This in turn is given by $E(r_7) = - \sum P(s_i / r_7) \log P(s_i / r_7)$. Yet, if the value at reception is uniquely picked out, it should then not depend statistically on the values of any other variables at the source (the position of a billiard ball in the model ought to depend only on the position of the corresponding gas molecule in the gas). So once again the conditional probability function is one or zero ($P(s_i / r_7) = 1$ if $i = 7$, and $= 0$ otherwise). So, we obtain the result that the equivocation can only go to zero if every variable in the source has a correlated variable in the model-signal. Patently this is not the case for the macroscopic variables, in which case $E(r) \neq 0$ and the model displays a degree of inefficiency.

While it is not possible to quantify the inefficiency in detail, it should be clear that the way to reduce equivocation is to tightly correlate every dynamical variable in the source to a variable in the model. This ensures completeness in the description the model provides of the phenomena (e.g., by building additional properties into billiard ball systems that account for the macroscopic properties of the gas), so not surprisingly equivocation inefficiency goes down. In other words, from an informational point of view, idealization introduces noise into the signal transmission provided by the model; while abstraction introduces equivocation; and the way to reduce the inefficiency generated by both is to tightly match the properties in the model to those in the source, and vice-versa. No model is ever perfect, in the sense of ever achieving this kind of one-to-one matching. On the contrary every model contains some degree of each kind of inefficiency. The art of modelling, in informational terms, involves a trade-off between noise and equivocation.

¹⁰ Determinism is also arguably an assumption for systems of billiard balls, if we assume an initial probability distribution over the dynamical variables of interest (as is done, e.g., in the tradition of the method of arbitrary functions). We shall ignore this complication and assume deterministic Newtonian dynamics.

5. Conclusion

Many scientific models aim at conveying information regarding their targets. When a model does so – and in so far as it does so – a model functions as a communication channel. We have in this paper endeavoured to take this insight seriously, and to apply the rudiments of communication systems theory to scientific modelling. The result is no doubt just an analogy – since models are *per se* not built as communication channels, nor may they be entirely treated as if they were. But it is an instructive analogy, which already at a preliminary stage sheds considerable light on some aspects of modelling practice. In particular, we have shown that some typical modelling techniques have correlates in informational terms. The analogy shows that often tractability is gained at the expense of informational faithfulness or completeness. There is therefore a certain trade off taking place in modelling which is best explained in terms of informational cost. Philosophers interested in modelling have not so far appreciated this informational cost, nor have they considered its diverse forms. We try to provide a first approximation in terms of informational *equivocation* and *noise*. These are technical terms – and we have employed the definitions in Shannon’s mathematical theory of communication. On this informational analogy, roughly, what is known as idealisation can be understood as introducing noise; while abstraction introduces a form of informational equivocation.

Shannon’s theory moreover provides precise ways to quantify over informational loss, by measuring the informational content in the source, and detracting noise and equivocation in the signal. It then becomes possible to derive a quantitative measure of information effectively transmitted. The application of such severe quantitative methods to scientific modelling is limited, and this shows some of the limitations of the analogy. Measuring the informational content of a dynamical system or phenomenon is far from trivial, as it depends on the description of the parts and their interrelation. In other words, the informational content of a system or phenomenon often sensitively depends upon what we call the phenomenological model. Yet, once this model is in place, it becomes possible to establish relations between its parts (which are genuinely probabilistic correlations in the case of statistical physics models) and compare them to those in a higher-level theoretical description. We have illustrated how this would work in the case of the billiard ball model of gases – a phenomenological description of a gas within Maxwell’s kinetic theory. The result is a rendition in informational terms of the idealizations and abstractions that operate in the model. This analogy is sufficiently robust to allow us to draw some conclusions regarding what a more realistic (either less idealized, or more concrete) description would involve; it also provides a better understanding of the trade-offs involved between tractability and information efficiency.

Furthermore, the analogy between scientific models and communication channels is suggestive of a number of further methodological and epistemological issues that deserve

to be explored – although there is no space in a single article to address them in any detail. The more obvious methodological questions concern the goals that trump informational efficiency – and, in particular, whether they carry an expectation of greater informative efficiency down the line. If so even the divergences from the goal of informational efficiency described here – in terms of idealization and abstraction in scientific models – would ultimately be accountable by recourse to presumed informational gains further on. The analogy would become more than just suggestive: It would provide the rudiments of an account of modelling as a branch of information theory. From an epistemological point of view a question that deserves to be studied is the status of the so-called *veridicality thesis*, which assumes all information to be true. The thesis has as many defenders as detractors (Floridi, 2007; Scarantino and Piccinini, 2010), and we have attempted to keep all our claims neutral in this paper. (In other words, the hope is that every claim in this paper is acceptable to epistemic realists and antirealists alike). Yet, it is reasonable to suppose that the veridicality thesis will apply in some cases of informative modelling, but not in all cases. If so, it would be worth figuring out a principled way to draw the relevant distinctions within modelling practice itself – thus rendering the realism-antirealism debate about scientific models an internal issue in science itself.

Acknowledgements: Mauricio Suárez would like to thank Tarja Knuuttila and the audience at the 2016 Valparaíso workshop on Models and Idealizations in Science for their comments and reactions, as well as financial support from the Spanish Ministry of Science and Innovation project PGC2018-099423. Agnes Bolinska would like to thank Anna Alexandrova and Joseph D. Martin for helpful feedback.

References:

Bailer-Jones, D. (2003), “When scientific models represent”, *International Studies in the Philosophy of Science* 17(1), 59–74.

Boesch, B. (2017), “There is a special problem of scientific representation”, *Philosophy of Science* 84 (5): 970-981.

Bolinska, A. (2013), "Epistemic representation, informativeness and the aim of faithful representation", *Synthese* 190: 219–34.

Bolinska, A. (2015), *Epistemic representation in science and beyond*. PhD Dissertation. Toronto: University of Toronto.

Bolinska, A. (2016), "Successful visual epistemic representation", *Studies in History and Philosophy of Science Part A*, 56: 153–160.

Borges, J. L. (1954), "Del Rigor en la Ciencia", in Dutton, E. P. (ed.) *Historia Universal de la Infamia*, Buenos Aires: Emecé. Translated as "On Exactitude in Science", in: *A Universal History of Infamy*, E. P. Dutton (1972).

Campbell, N. (1920), *Physics: The Elements*. Cambridge: Cambridge University Press.

Contessa, G. (2007), "Scientific Representation, Interpretation and Surrogate Reasoning", *Philosophy of Science* 74 (1): 48-68.

Dretske, F. (1981), *Knowledge and the Flow of Information*, Cambridge: Cambridge University Press.

Floridi L. (2007), "In Defence of the Veridical Nature of Semantic Information", *The European Journal of Analytic Philosophy* 3(1): 1–18.

Giere, R. N. (1988), *Explaining Science: A Cognitive Approach*, Chicago: Chicago University Press.

Giere, R. N. (2004), "How Models Are Used to Represent Reality", *Philosophy of Science*, 71(5): 742–752.

Giere, R. N. (2009), "An agent-based conception of models and scientific representation", *Synthese*, 172(2): 269–281.

Hesse, M. (1966), *Models and Analogies in Science*, Notre Dame: University of Notre Dame Press.

Jones, M. R. & Cartwright, N. (2005), "Idealization XII: Correcting the Model", in M. R. Jones and N. Cartwright, eds., *Idealization and Abstraction in the Sciences*. Rodopi.

MacKay, David J. C. (2003), *Information Theory, Inference, and Learning Algorithms*, Cambridge: Cambridge University Press.

Mäki, U. (2009), "MISSing the World. Models as Isolations and Credible Surrogate Systems", *Erkenntnis*, 70(1): 29–43.

Matthews, M. R. (2004), "Idealisation and Galileo's pendulum discoveries: Historical, philosophical and pedagogical considerations", *Science & Education* 13: 689-715.

McMullin E. (1985), "Galilean idealization", *Studies in History and Philosophy of Science* 16: 247–273.

Morgan, M. and M. Morrison (1999), *Mediating Models*, Cambridge: Cambridge University Press.

Pero, F. and M. Suárez (2016), "Varieties of Misrepresentation and Isomorphism", *European Journal for the Philosophy of Science* 6 (1): 71-90.

Pierce, J. R. (1961), *An Introduction to Information Theory: Symbols, Signals and Noise*, Dover Publications.

Scarantino, A., and G. Piccinini (2010), "Information Without Truth", *Metaphilosophy*, 41(3): 313–330.

Shannon, C. E. (1948), "The Mathematical Theory of Communication", 1963. *M.D. Computing: Computers in Medical Practice*, 14(4): 306–17.

Suárez, M. (2004), "An Inferential Conception of Scientific Representation", *Philosophy of Science*, 71(5): 767–779.

Suárez, M., ed., (2009), *Fictions in Science: Philosophical Essays on Modelling and Idealization*, London: Routledge.

Suárez, M. (2024), *Inference and Representation: A Study in Modeling Science*, Chicago: The University of Chicago Press.

Suppes, P. (1962), "Models of Data", in *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, Ernst Nagel, Patrick Suppes and Alfred Tarski, eds., pp. 252-261. Stanford: Stanford University Press.

Teller, P. (2001), "Twilight of the Perfect Model Model", *Erkenntnis* 55 (3): 393–415.

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

Weisberg, M. (2007), "Three kinds of Idealization", *Journal of Philosophy* 58: 207-233.

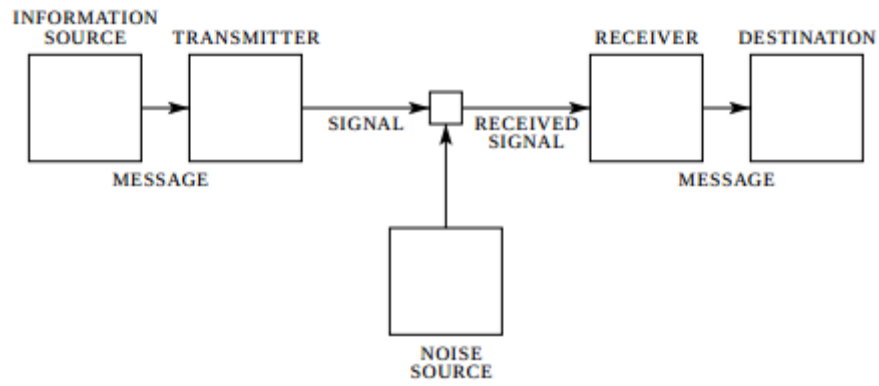


Fig. 1—Schematic diagram of a general communication system.

Figure 1: Shannon’s mathematical theory of communication (Shannon, 1948)