BELIEF REVISION IN SCIENCE: INFORMATIONAL ECONOMY AND PARACONSISTENCY

Daniel Credico de Coimbra[1]

Abstract: In the present paper, our objective is to examine the application of belief revision models to scientific rationality. We begin by considering the standard model AGM, and along the way a number of problems surface that make it seem inadequate for this specific application. After considering three different heuristics of informational economy that seem fit for science, we consider some possible adaptations for it and argue informally that, overall, some paraconsistent models seem to better satisfy these principles, following TESTA (2015). These models have been worked out in formal detail by TESTA, CONIGLIO, & RIBEIRO (2015, 2017).

Keywords: Paraconsistent Belief Revision. Ockham's Razor. Principle of parsimony. Minimality. Informational economy. Rationality. Scientific method.

1. Introduction and Outline

The standard model in the field of belief revision was created in 1985 by Carlos Eduardo Alchourrón, Peter Gärdenfors, and David Makinson, and has been appropriately named "AGM." Its basic characteristics are as follows.

The model was designed to show what a rational revision of an agent's beliefs in response to new information may, or should, look like. An agent is an idealized entity that accepts and rejects belief-representing sentences. The set of an agent's accepted sentences is called an *epistemic state*. Agents and sentences are given general formal definitions, so that they can be put to multiple uses depending on one's objectives as a theorist: agents can be understood as databanks, individual humans, artificial intelligences, collections of human beings such as scientific communities, and otherwise, whereas sentences or beliefs can be interpreted as data, facts, norms, rules, objectives, hypotheses, assumptions, and much else besides; and collections of sentences can be interpreted as worldviews, models, or scientific theories, among others. An epistemic state can be operated upon so as to add and remove sentences in response to incoming information, and to this end a few operations have been defined.

This paper focuses on the application of belief revision models to scientific rationality, and it examines how well the AGM model fares in this task. Section two will be dedicated to an elementary exposition of AGM, during which we attempt to show

1 Undergraduate student of Philosophy at the University of Campinas. E-mail: d155077@dac.unicamp.br.

how the model is consonant with principles we selected as ways of capturing intuitions regarding *informational economy*.[2] In section three and section four, we argue that AGM falls short of being an adequate model of the scientific process and that informational economy plays an important role in scientific inquiry. In sequence, we explore adaptations of AGM that may enhance its informational economy adequacy and its virtues as a model of science.

In conclusion, we will not provide a formal model of scientific development but, rather, provide informal considerations regarding this development along with informal arguments for the suitability of *AGM-like paraconsistent systems of belief revision* as models of science, as presented in TESTA, CONIGLIO, & RIBEIRO (2015, 2017). These models turn on *paraconsistent logic*, a kind of logic that was developed in the twentieth century as an alternative to classical logic, in which the *law of non-contradiction* and the *principle of explosion* are not theorems. Such a logic seems perfect to capture what has been called the *learning power* of contradictory states (cf. TESTA 2015), and on this hinges our aforementioned conclusion. All historical information regarding belief revision, unless otherwise noted, comes from HANSSON 1999.

2. The Basics of AGM and the Concept of Informational Economy

The AGM model has many features that can be seen as closely abiding to informational economy, or so we shall argue. Perhaps its three builders were quite aware of how scant information might be and how valuable it is, as one of AGM's chief characteristics is its classical logical closure, which extracts as much new information from the current information as is deductively possible — as if no stone was left unturned. This means the epistemic state, which we may denote by “K”, is posited to be identical with the set of logical consequences of its elements. Formally, this is written as “K = Cn(K)”, where Cn is an operator that outputs a set with all the logical consequences of its argument.

Since this logical closure is understood classically, if K contains any pair of contradictory sentences it “explodes” and trivializes itself: any sentence becomes deductible and K becomes identical to L, where L is the object language the system is

2 Informational economy is one of several *desiderata* of belief revision models and of science, as will be argued.

formalized in. Such a set would be maximally uninformative, since true or reliable sentences (for instance) are undifferentiated from false or unreliable sentences. Due to the *principle of explosion*, then, *non-contradictoriness* is classically equivalent to *consistency.*

New information is presented to the system through an input apparatus, the *epistemic entry*, and there are two main ways this apparatus has been configured. The standard AGM model adopted something called *prioritary revision*, in which new information is always held in higher regard than old information: in the case old information conflicts with new, generating contradictions, the former is scraped in benefit of the latter, which will be duly incorporated. Other theorists prefer the idea of being able to submit new information to close scrutiny before incorporation — and perhaps reject it, if it is deemed not good enough according to appropriately set measures. For this purpose, models adopting *non-prioritary revision* were built. These models put new information on hold while the benefits of updating the system (scraping old information) to accomodate it are evaluated. If these old pieces of information seem more useful, more reliable, or in any other way more *well-regarded* than the new piece of information, then the latter will be ignored so that the former can be kept.[3]

In AGM there are three basic operations that dictate how a given change in the epistemic state is to proceed: *expansion*, *contraction*, and the eponymous *revision*. (Other operations, such as *consolidation* and *merging*, will not be covered, in favor of an examination of the basic operations.) The three basic operations seem to work so as to best satisfy four criteria of rationality. We say first that operations upon an epistemic state must result in a *consistent* state, meaning it has not been triviliazed, and we follow saying that they should be closed under logical consequence, which in standard AGM means *classical* logical consequence. Third, we say they must account for the fact that some beliefs are more certain, informative, explanatory, or otherwise more *well-regarded* than others, and we will explore this notion further in section 2.2.[4]

The three previous criteria were plausibly seen by TESTA (2015) as hinging on a fourth, overarching criterion of informational economy: operations should minimize

3 Examples of *non-prioritary revision* approaches are *screened revision* and *credibility-limited revision*. These classically-oriented approaches for dealing with contradictory new information will not be covered here and, for reasons that will become clear, we will instead favor the exposition of paraconsistent models, in which conflicting information is not always a problem.

4 There are various ways of of capturing the intuitive ranking of sentences according to regard, such as *epistemic entrenchment*, *Grove's spheres*, and *selection functions*. See HANSSON (1999) for a detailed account.

information loss and maximize information gain. Closure maximizes information gain, *consistency* minimizes information loss, and ranking sentences according to regard is, arguably, *usually* done so that the more informative and certain sentences will be favored. We will see that there are further requirements in information economy beyond such processes of minimization and maximization, such as parsimony and perhaps, in science, explanatory power.

2.1. AGM expansion and the razors

We begin this section by proposing three principles, embodying general ways of capturing intuitions regarding informational economy, and we proceed to show how each of the operations may be interpreted in light of such principles. It is, however, only in section four that we provide an informal and exploratory analysis of how AGM's could better satisfty these intuitions, as there are some consideration we must make before that. It should be noted that these principles will be called *razors*, for the three of them have some conceptual or historical relation to Ockham's razor; for ease of exposition too we have unified them under this label. These principles will be critically examined and defended in section four. They are:

- **The razor of silence.** Do not assent to anything more than is necessary — in Latin, with a somewhat different literal meaning, we may say *entia non sunt multiplicanda praeter necessitatem.* What counts as "necessary" is up to some interpretation and, as we shall see in section four, the interpretation given is important to evaluate how well *expansion* satisfies informational economy in science. The name of this razor is due to SOBER (2015), and it is generally considered to be Ockham's razor and a parsimony principle.
- **The razor of economy.** TESTA (2015) has pointed to what he has called the *dual* of Ockham's razor: do not discard anything more than necessary — in Latin, perhaps we may say *entia non sunt subtrahenda praeter necessitatem*. We propose to call this the *razor of economy*, since through it one economizes already held information, beliefs, or assertions like one saves money. Likewise, we could say the *razor of silence* economizes silence, or suspension of judgment. In this respect, these two razors are easily seen as duals of one another, as one preserves current *judgment* and the other preserves current *lack*

of judgment.

- **The razor of explanation.** A way of putting it is: assent that which is necessary to explain whatever it is that needs explaining (which may be, conveniently to our purposes, observational data). This razor enshrines the spirit of science, and it should be respected by any model of scientific rationality. SOBER (2015) pointed out this is a principle at least as old as Ockham's razor, as will be seen later on.

Expansion is an operation designed to add a single piece of information to an epistemic state, along with all its classical logical entailments, neither adding beliefs beyond that nor removing any. It is designed so as to respect the three criteria provided above, and we gather it is structured in close agreement with the *razor of silence*, making it so it at least partially respects the fourth criterion. *Expansion* also seems to abide to the *razor of economy* by not removing any beliefs during its operation. However, this means that, unfortunately, this operation fails to always preserve *consistency*: if one attempts to add some belief contradictory to another belief in the epistemic state, *expansion* will provide no mechanism for preventing the explosion of the state through classical closure. In this scenario, *expansion* fails to respect the *razor of silence*.[5] (We will see how *revision* does not share this problem.)

Now, onto the inner workings of *expansion*. If K and L are to be epistemic states, α a belief, and K+α the result of the *expansion* of K by α, then this operation can be defined using the following six postulates (cf. HANSSON 1999): *Closure*, which states K+α should be closed under logical entailment. *Success*, which states α should be an element of K+α. *Monotonicity*, which guarantees *expansion* preserves any relation of containment. *Inclusion*, which states the original epistemic state should be contained in K+α. It should be noted that Inclusion makes *expansion* respectful of the *razor of economy* by blocking any removals. *Vacuity*, which states K+α should be no different from the original epistemic state if α was already there — though, as it turns out, this already follows from the other five postulates.

Finally, we've got *Minimality*, which deserves detailed comment. It states that K+α should be the smallest set that satisfies Closure, Success, Inclusion, Vacuity, and Monotonicity. The importance of this postulate is its guarantee that nothing was added

5 It may be argued that this makes *expansion* disrespectful of the *razor of economy* also, for trivialization leads to no differentiation between (for instance) true and false sentences, and losing any such differentiation plausibly amounts to (perhaps total) loss of information.

to K besides α and its classical logical entailments. Minimality makes it so *expansion* respects the *razor of silence* in cases where ¬α is not in the epistemic state, as it avoids comittment to anything but that which regards the new information-piece.

To conclude this section, we will note that a central result in AGM theory is that the operations of *expansion*, *contraction*, and *revision* can be equivalently defined both as a certain set of postulates and as a certain set-theoretic constuction; this result that has come to be called *representation theorem*. (Though there is generally room for some variation in the choice of postulates.) The following is the remarkably simple construction for *expansion* that satisfies our initial requirements and is logically equivalent to the six postulates exposed here.

- **The construction for *expansion*:** $(K + \alpha) = Cn(K \cup \alpha)$.

2.2. AGM contraction and the hierarchy of regard

Contraction was built to assure the definite removal of some belief-representing sentence from the original set of belief-representing sentences. It's set-up so as to output a *consistent*, logically closed epistemic state, and to neither add nor remove more sentences than necessary. The minimization of removals seems to be one half of its biggest challenge, while removing sentences without disrespecting their hierarchy of regard is the other. A toy example may be instructive on why this may be challenging.

Suppose our epistemic state is $E = \{\alpha, \beta, \beta \rightarrow \delta, \delta \rightarrow \alpha, \beta \rightarrow \alpha\}$, and our wish is to perform a *contraction* of K by α. If α alone was removed from E, it would promptly be added back again through *modus ponens*, defeating our purpose. In order to prevent this, some of the other four beliefs in E need to be removed, but no more than necessary. One would naturally think we should remove the smallest amount necessary, but it is also informationally economic to remove the set that aggregates the least total *regard* in its beliefs even if it is the biggest of candidates. Suppose we have more confindence (for instance) in $\{\beta\}$ than in $\{\beta \rightarrow \delta, \beta \rightarrow \alpha\}$. Both sets, if removed, would be sufficient for our task, and removing the latter, the bulkier the two, will prevent us from removing beliefs we are *ex hypothesi* more confident in. The *desiderata* of minimization of removals and of maximization of regard walk together.

The regard attributed to beliefs may be a function of their informativeness, their certainty, their explanatory power, or some other metric, depending on one's modeling purposes as a theorist. As we have learned from GÄRDENFORS (1992), the logical form of belief-representing sentences do not capture the various *desiderata* of regard and, thus, in order to suit the theorist's aims in constructing a belief revision model, extra-logical criteria will be needed to perform a useful ranking. This makes it so a hierarchy of regard is highly valuable to the preservation of desired information — if *contraction* manages to satisfty it, then this operation will fulfill the *razor of economy*. All this must be done, we should recall, without violating *consistency* or closure.

If K is an epistemic state and α a belief, then the following six postulates characterize the *contraction* of K by α, to be denoted by 'K–α' (cf. HANSSON 1999): *Closure*, which states K–α should be closed under logical entailment. *Success*, which states non-tautologies should have been definitely removed from K–α. *Extensionality*, which guarantees that logically equivalent formulas will behave equally during *contraction*. *Vacuity*, a postulate that follows from the others, stating that K–α is identical to the original epistemic state if α was not in it to begin with. *Inclusion*, stating K–α is included in K; this assures no information was added to K during the process of *contraction*, in clear conformity the *razor of silence*.

Finally, we've got *Relevance.*[6] It was designed in order to ensure that only those sentences which entailed the sentence we wish to remove are, in fact, removed. Formally, it is stated as follows: $\beta \in (K \backslash K - \alpha) \rightarrow \exists K'(K - \alpha \subseteq K' \subseteq K \land \alpha \notin K' \land \alpha \in (K'+\beta))$. This seems to assure nothing is removed in *contraction* besides the minimum sufficient to definitely get rid of the target information-piece, which makes it respectful of the *razor of economy* for not throwing out anything unnecessarily.

By the representation theorem, we have that Construction can be defined equivalently to the previous postulates through a set-theoretic construction. This construction will make explicit how it is that sentences are selected for removal in a way that respects the dual of Ockham's razor: the *razor of economy*.

- **The construction for *contraction*:** Some subset K' of K, which does not contain or entail α, should be chosen as the output of the *contraction*, with the difference between K' and K being the information that has been thrown out. That being the case, if there are multiple such subsets, then, for the sake of the *razor of*

6 Originally a postulate named *Recovery* was used instead, but this postulate is surrounded by substantial controversy in the literature and, thus, following recent work, we adopt the *Relevance* postulate in its place. It should be noted that this discussion is relevant to discussions regarding informational economy, for postulates such as *Recovery* and *Relevance* are added to engender minimal change in *contraction* but, for a lack of space, we will unfortunately not dwell on it.

economy, we should pick the biggest of those which do not entail α. Formally, we need a K' such that, for any subset K* of K, (K' ⊂K* ⊆K) → (α ∈Cn(K*)). However, multiple K's might fit this description, constituting what is called a *remainder set* K⊥α. To chose amongst these, a *selection function* γ is defined so as to prioritize sentences and sets according to their regard. Even then, multiple subsets might receive the highest "score," and the best way to work around this economically is, the literature has found, by removing the intersection of all the highest-ranked subsets — if there is any K' to rank at all. Formally, ((K⊥α) ≠ ∅) → ($K{-_\gamma}\neg\alpha$ = ∩γ(K⊥α)), where $K{-_\gamma}\neg\alpha$ is contraction of K by α relative to the gamma function.

2.3. AGM revision

We have mentioned that *expansion* adds an information-piece to an epistemic state no matter what happens, with the possibility of this information-piece being contradictory with something in the state and an explosion resulting. The namesake operation of belief revision is built so as to avoid this failure of *expansion*: in accordance with the primacy of new information, if such a contradiction exists between new and old sentences, *revision* proceeds by tinkering with the epistemic state, i.e. removing information, so that the new information-piece stops being contradictory with it, and only then is it added to the state. (An obvious requirement is that the information-piece *itself* doesn't embody a contradictin.) Similarly to other operations, the objective is to do so neither losing more information than necessary nor assenting to more than is required. Once more, as we want to show, a conformity with the two razors can be seen.

If K is an epistemic state and α a belief, and 'K∗α' is to denote the *revision* of K by α, then six postulates that characterize *revision* are (cf. HANSSON 1999): *Closure*, which states that K∗α is to be closed under logical entailment. *Success*, which states that α is to be an element of K∗α. *Consistency*, stating that the *revision* of K by α will output a *consistent* state if α itself is *consistent*. *Extensionality*, assuring two logically equivalent senteces will be treated equivalently.

Then we've got *Vacuity* and *Inclusion*. The former states that if α is *consistent* with the original epistemic state, then no tinkering with K must be done, and thus K∗α = K+α. The latter, in turn, states that in all cases K∗α should be contained in K+α, which means nothing more will be added in *revision* than would be added in *expansion* — α and its logical entailments. These two postulates bind the operation of *revision* to the operation of *expansion*, which in the latter's case means binding it to the postulate of Minimality and, thus, successfully binding its addition procedures to the *razor of silence*.

There is another interesting result that ties *revision* to informational economy. HARPER (1977) proved a theorem that has come to be called *Harper Identity*, in his honor. It states the following: $(K\ast\neg\alpha) \cap K = K\text{–}\alpha$. This identity proves *revision* is bounded by the workings of *contraction* and, since our *contraction* removes sentences while respecting the hierarchy of regard, then so does *revision* in its removal procedures (which we have called "tinkering"). That being the case, *revision* also upholds the *razor of economy*. It also establishes that *contraction* is not more fundamental than *revision*, as the former is definable in terms of the latter. The converse is also true — *revision* is definable in terms of *contraction* —, given that the following equality holds: $(K\ast\alpha) = (K\text{–}\alpha)+\alpha$.

This latter theorem has been dubbed *Levi Identity* in honor of its discoverer, Isaac Levi. This provides a nice segue into *revision*'s set-theoretic construction, since the Levi Identity plays a role in it. Following the representation theorem, this construction is equivalent to the previous postulates:

- **The construction for *revision*:** The construction equivalent to the postulates above is based on the aforementioned Levi Identity, and it is: $(K\ast\alpha) = (K\text{–}_{\gamma}\neg\alpha) \cup \alpha$, where γ is the selection function defined for *contraction*. If the function was defined so as to make the operation a *partial meet contraction*, then we call this operation *partial meet revision*.

3. Application Problems in Classical Models

In this section, we wish to argue for two theses. First, that classical logical closure makes AGM able to model only extremely idealized agents. It falls short in accounting for many important aspects of the kinds of rationality exhibited by human agents, supercomputers and, more relevantly to our purposes, scientific communities — in short, the rationality of non-idealized agents, which have limited computing capacities, short time-windows for churning out conclusions, and finite memory. (This topic has been extensively covered by CHERNIAK (1986).) Second, that the classicality of AGM's logical closure makes it unable to attend to demands of informational economy in science to a satifactory degree, due to its inability to adequately exploit the *learning power* of contradictory information. In section four we will suggest that paraconsistent models seem better adequade in these two aspects.[7]

3.1. Logical omniscience, irrelevance, and computational intractability

The first problem that comes to mind is that even when we're dealing with a language that contains very few sentences, the number of logical consequences of any given belief is infinite, with endlessly many sentences with soaring levels of complexity. Tautologies are also infinite and also exhibit indefinitely big levels of complexity. Deducing and knowing all these sentences amounts to what has been called *logical ominscience* (cf. TESTA 2014).

The first problem comes with limited memory, which is unable to hold an infinite amount of information-pieces. The second comes from limited computational power, which is unable to parse indefinitely long and complex sentences and to efficiently search indefinitely long databases, however the information may be organized and nested. Thus, no finite physical system would be capable of being logically omniscient.

A third problem is that many tautologies and many logical entailments of already-held beliefs are *pragmatically irrelevant* for any actual system. For instance, one can derive no practical application for an immensely long disjunction of which only one disjunct is known or believed to be true. Thus, demanding rational systems to be closed under logical entailment may be, at best, futile. At worst it would be counter-productive, if the system has limited processing power and memory capacity: even a system with an infinite memory and formidable search heuristics would be clogged up by the vast majority of irrelevant sentences in the database, and thus be hendered inoperant, unable to revise its beliefs efficiently or act upon the world effectively: quite plausibly, the vanishing minority of relevant beliefs just won't be reliably found in short notice.

LEVI (1991) argues that in AGM an agent can be interpreted as being only

7 We want to reaffirm that the AGM model is constructed in a way that is sufficiently general to assure wide applicability, and that it was not specifically tailored for this application. Thus, the model is not being criticized in general, but rather being checked to see if its versatility includes modeling science.

commited to these tautologies and entailments, rather than actually *accepting* them — a proposal that would alleviate the problems of logical omniscience for AGM. (Also see GARDENFÖRS 1988.) This is a minor point, but it seems to us that a relation of commitment to sentences can only be captured by *belief bases*. For example, suppose an agent is commited to β because the belief α implies it. This agent would not believe β itself, but rather holds it *on the condition* that α is held: β would be removed if α was removed. We could say β has no standing of its own, and this is what makes β a commitment rather than a belief. In *belief sets*, however, all information-pieces stand by themselves. It is only in belief bases that we could have sentences being held conditional upon other sentences being held: sentences in the *implicit set* depend upon sentences in the *explicit set*.

The standard version of AGM, however, models epistemic states as *belief sets*, making it liable to the problems outlined in this sub-section. Thus, it seems that in application of belief revision models to science we should scrape classical logical closure if we are to keep belief sets, or scrape belief sets for belief bases if we are to keep classical logical closure. In section 3.3 we will see a further, and perhaps more serious, problem with classical closure in the modeling of science.

In the next section, we explore another shortcoming of AGM which spins not on its classical logical closure, but on its deductivism.

3.2. Non-idealized rationality: humans and scientific communities

The past 50 years of cognitive science has seen the rise of a field of research on the multiple ways humans reason heuristically, often leading to irrational patterns of belief-formation and behavior called *cognitive biases*, which were extensively studied by scientists such as Daniel Kahneman and Amos Tversky (see KAHNEMAN 1982). We argue that, despite these shortcomings, not only heuristic reasoning can lead to rational behavior, but it can even be pragmatically *necessary* to it when we are dealing with non-idealized agents. Heuristic processes can output reliable and approximately true conclusions in much less time and with much less data than deductive processes, which allows for decisions within reasonable time and for successful theory-building in a scenarios in which there's limited time for data-collecting and data-processing. AGM, however, is commited to deductivism, meaning it cannot model agents whose rationality depend on heuristics.

We will quickly lay out a few possible examples of how heuristics may be useful in human life, and then analogize to science. First, it is possile that we arrive heuristically at useful models of the behavior of physical objects (i.e. folk physics). For instance, BAILLARGEON (1994) reports evidence that there are “highly constrained, innate learning mechanisms” (p. 135) about physical phenomena such as collision, size constancy, and gravitation, which allows children to get a grip on the workings of the physical world at quite an early age. To give another example, CHOMSKY (1980) details how human infants learn language stunningly quickly and accurately even though their observations underdetermine the correct grammar of the language being spoken (the “poverty of stimulus” argument). These two examples suggest humans may have flexible pre-set programs for making the right kind of *guesses* (in hypothesis-formation and in conclusion-making) during theory-building so as to result, after a few years of experience, in a host of accurate and reliable theories of the world. (It is interesting to think how a tendency to develop such mechanisms could have been installed in us via evolutionary processes.)

It is also sure that we have reliable heuristics for rapid decision-making. For example, it seems to us there must be localized heuristic procedures giving us the ability to quickly deduce the trajectory of incoming objects from scant visual data, allowing us to make snap, *life-saving* decisions, such as some act of avoidance.

In a world with temporal constraits for action and limited resources for data-gathering, trading logical rigour and some precision for speed and low information requirements can *heighten* rationality, instead of *hindering* it. Heuristics are *sine qua non* for non-idealized rational agents, and scientific communities are no exception: science too has time and resource-gathering constraints.

First, it is a *desiderata* to output reliable conclusions within a few human generations, which cuts short available time. Second, it posits unobservable entities and universal generalizations, which are not deductible from science's source of information: the *observation* of *particulars*. Third, it seems that science would not manage to advance so quickly if it did not pick up presently non-deductible theories and assumed them as true for the purposes of research (i.e. as working hypothesis), since such practice often leads to new discoveries. Perhaps science would be mostly stagnant without it.

Given these three facts, there must be some non-deductive, heuristic processes regulating the creation of hypotheses (containing generalizations and unobservables), which hypotheses will be picked up for reseach, and which will be presently held to be probably approximately true. The process of theory-building and theory-confirmation, then, has ineliminable heuristic elements.

3.3. Contradiction, learning, and scientific progress

We want to argue in this section that scientific communities have, both as a matter of fact and as a matter of *epistemic* necessity, internal contradictions. Having contradictions in one's epistemic state is essential to assure a reliable learning process, so long as trivialization is avoided — or so we shall argue.

We should first note that some degree of *epistemic resilience* is rational, so as to let enough data sink in before passing judgment and rejecting long-held theories. If one's model has a history of wide-ranging predictive success and empirical adequacy to what seem like reliable data, it seems plausible we would act rationally if we took incoming contradictory evidence as statistical flukes or only as *apparently* contradictory, as in the case of Uranus's orbit which only seemed impossible because we lacked data on Nepture, or as in the case of the incompatibility between quantum mechanics and general relativity, which we hope is only apparent.

This epistemic resilience is not absolute, however. We gather the following metaphor will enlighten the logic of epistemic resilience. The strenght of the contradictory evidence must get above a certain *critical mass* (whose value should be set according to the reliability of our model) before it warrants a model revision, or before it makes rational for us to throw out our model and attempt to erect a brand new model upon the new evidence (a scientific revolution). It is widely believed that scientific revolutions only occur once a certain critical mass of problems accretes, even if this critical mass may not be as big as some have thought it to be (cf. LAUDAN 1990). This makes sense when one considers that discarding a reliable, promising, and thus *good* model upon insufficiently strong evidence is *not* a rational move.

This line of thought has unfavorable consequences to the adequacy of AGM as a model of rationality in science. The AGM model obliges one to either incorporate or reject each incoming contradictory evidence; if we take its *prioritary revision* version, it always incorporates new evidence. If *non-prioritary revision* is adopted instead, each contradictory evidence is either incorporated or thrown out, and there is never an *accretion* of pieces information to reach a critical mass before revision. All reform is piece-meal, which, as we have seen, will not do. Thus, substantial model revisions in AGM cannot be rational.

An alternative model of rationality is needed, one that can temporarily acommodate masses of contradictory statements without trivialization, until a rational theory revision can be incurred. We can now see how contradictory states can have the sort of *learning power* mentioned before. It is precisely to capture the *learning power* of contradiction that AGM-like paraconsistent models of belief revision will come in, since they can handle contradictions well while preserving the dynamic character of epistemic states, as has been shown in TESTA (2015) and detailed in TESTA, CONIGLIO, & RIBEIRO (2015, 2017).

4. Information and Paraconsistency: Rationale and Solutions

The *razor of silence*, which we have seen is Ockham's razor, has reached contemporaneity as a strong motivating force in analytic philosophy, theoretical physics, and evolutionary biology (cf. SOBER 2015). There is a general conviction that some version of it must be true: perhaps nature is inherently simple, or perhaps parsimony is a truth-conducive methodological principle for other reasons. String theory, for instance, has received considerable acceptance in part by appeal to its perceived elegance and simplicity. Modal realism, on the other hand, has crucial and wide-ranging applications, but has been generally neglected for being as unparsimonious as a theory could possibly be. Too high a price to pay, it is said.

In this section we intend to explore further the virtues and shortcomings of this razor and of its dual, the *razor of economy*. Furthermore, SOBER (2015) has mentioned that the original formulation of Ockham's razor had two opposing aspects, one negative and one positive, but that the latter has been forgotten over time. We may call this positive side the *razor of explanation*, and we will argue that our other two razors would be improved by an interplay with the *razor of explanation*. It is worth quoting Sober (2015, p. 7) in full on this:

> The maxim that often comes to mind when people now think of Ockham's razor is negative, but it is usually understood to have a positive complement. There is "do not accept a postulate if it is not needed to explain anything," but there is also "accept a postulate if it is needed to explain something.

4.1. The razor of silence

We should start with a defense of the *razor of silence*. First of all, as mentioned above, simplicity has great intuitive content. Second, it prevents the clogging-up of a system with limited search capacities with misleading, groundless, and “unnecessary” beliefs. Third, it is a theorem of Bayesian inferential statistics that, unless B is a sentence known with *certainty* (or has an associated prior of *one*), it is always the case that (i) the probability of A being true (without passing judgment on the truth value of B) is higher than (ii) the probability of both A and B being true. Formally, it is a theorem in Bayesian probability theory that, in such cases, $P(A) > P(A\&B)$. If we are to take the most probable one as our working hypothesis, we ought to pick the simplest one. However, as our objective is being scientifically rational, we should not fall into the mistake of accepting a theory so simple it is not explanatorily adequate. If we have some theory with supposition A and is explanatorily adequate, we should only add the uncertain supposition B if there is to be a counteracting gain in explanatoriness. Otherwise, we should rather suspend judgment on whether supposition B is true.

A problem arises, however, when comparing complexity among theories that do not share suppositions. While it might be possible to justify the adoption of a Keplerian astronomy because it postulates less orbits in comparison with Ptolemaic and Copernican theory and their myriad of orbits orbiting orbits, it seems it would not be possible (or, at least, not with the concepts we have hitherto deployed and developed) to compare, in terms of complexity, two theories utterly different in their ontology and postulated mechanisms. So it may be the case that the above bayesian theorem gives us little insight in many cases, and this is why the qualified debate concerning simplicity and Ockham's razor still rages on.

To tighten up a loose-end alluded to in section 2.1, there are at least two possible readings of “necessity” in “do not assent to anything not necessary.” In a deductive reading, what is necessary is only what comes in directly as information and whatever is entailed by it; this is how AGM is constructed to work. However, science aims to constructing good explanations, as Newton has put it when stating four principles of reasoning which have guided his theorizing in his *Principia Mathematica*: *“We are to admit no more causes of natural things, than such as are both true and sufficient to explain their appearances, (...) for Nature is pleased with simplicity, and affects not the* *pomp of superfluous causes."* (Cf. SOBER 2015.) Note how this is a quite direct statement of the interplay we will suggest between the *razor of silence* and the *razor of explanation.*

For this reason, science builds general hypotheses and postulates unobservable entities and mechanisms, going well beyond its set of observations to generate explanations for these observations. This means adding only individual observations to our epistemic state will not be enough. It is *necessary* to add beliefs that adequately explain the incoming data, if our aims are scientific. We suggest reading "necessity" as being explanatory necessity, and to make this clear we should amend the *razor of silence* to say: do not assent to anything not necessary for our best theoretical explanations, unless it is a direct observation. We should call this the *explanatory razor of silence*. Strict minimality of additions (*razor of silence*) should be traded in for explanatory prowess.

As a side note, perhaps we can say the *razor of explanation* is also directly related both to maximization of information gain and to parsimony. First, if a general theory is provided support, the theorist acquires knowledge about a wide range of entities, well beyond the observed ones — maximizing information gain. Second, general theories explain a wide array of phenomena with only a few postulates, and perhaps it is more parsimonious to state the simple and unified *"Every x is y"* than the unwieldy conjuction *"x_1 is y & x_2 is y & x_3 is y &. . ."*, with a conjunct for every observed x. This is a path to be explored in further work.

4.2. The razor of economy

In non-idealized situations, information is usually pricy or scant, and *prima facie* it seems it never has *negative* value if it does not generate contraditions. Therefore, it is reasoned, it should be kept as often as possible, and from this comes a good rationale for following the *razor of economy*. However, we should note that non-idealized agents must, as a practical necessity, often throw out information, for the sake of resource-management, in the ways explored in section 3.1.

However, such throwing out violates the *razor of economy*, which suggests that this razor needs updating. Perhaps we should re-state it as: *do not throw out information without a good reason to do so* — and resource management does seem like a good reason to do so. There is yet another such good reason for information-removal. During scientific model revisions and scientific revolutions, a lot of beliefs are thrown out. Suppose some theory T_1 deductively warranted some belief α, and scientists believed in α because they believed in T_1. Suppose further that, through a scientific revolution, T_1 was superseded by T_2, and that T_2 is neutral as regards α: it neither implies nor denies it. Should scientists keep their belief in α? They surely wouldn't, and we gather they shouldn't. In science, there is no reason to accept some sentence which is not even *recommended* by our current theory or our direct observations, and there is a strong reason against it: it can be misleading, leading scientists to mistake beliefs they have for good reasons and beliefs they have as remnants of outdated theories.

Strict minimality of dismissals should be traded in for the sake of resource-management and the avoidance of deception. Now a criterion of removal is needed, if we want to keep some beliefs and not others. The *razor of explanation*, in fact, seems to be the ideal criterion: throw out whatever current beliefs are neither pieces of observation nor elements of our best theoretical explanation of the observational data. We should call such razor the *explanatory razor of economy*: do not throw out that which is useful for our best theoretical explanation or a direct observation.

One should note that, in their explanatory overhauls, the *razor of silence* and the *razor of economy* seem to become even closer twins, since the former forecloses the addition of sentences beyond explanatory necessity and observation, and the latter enables or leads to the removal of sentences beyond explanatory necessity and observation. Perhaps they recommend the very same epistemic state at any given time, which suggests explanation was the bridge between intuitions regarding parsinomy and economy.

4.3. Conclusion: The razors and the paraconsistent informational economy

> *In formal logic, a contradiction is the signal of a defeat: but in the evolution of knowledge it marks the first step towards a victory.*
> *(WHITEHEAD, 1925)*

In conclusion, we have seen that the razors of silence, economy, and explanation are adequate in capturing intuitions about informational economy, and that, albeit each has its own difficulties, sometimes numerous, they have substantial grounding in intuition and reasonable grounding in articulated reasoning. Perhaps this is why many theorists, such as Sven Ove Hansson, have included and argued for operations of belief revision that cared for informational economy. Informational economy is an important ingredient of rational and truth-conducive investigations.

We have seen that the *razor of explanation* accounts for the theory-building objective of science, and that it can profitably interact with the *razor of silence* and the *razor of economy*, tying it up to the long-standing intuitiveness of parsimony in the former case, and evading an unnecessary shedding of information in the latter case. If a belief revision system purports to model the rationality of scientific communities, it should abide to the interplay of razors detailed above. The *razor of explanation* may be optional for some purposes, but it seems obligatory to anyone aspiring to do serious empirical science.

We have also seen how non-deductive, heuristic reasoning is vital to a non-idealized agent, and how it heightens rationality, instead of hindering it.

Furthermore, we have seen that gathering sizeable chunks of information contradictory to one's beliefs is necessary for rationality, as it allows for rational and reliable theory revision. An interesting historical example may come from mr. Charles Darwin. We recall reading about his habit of writing down in red ink all evidence contradictory to his theory of gradual evolution (for he could not trust his memory for such a cognitively uncomfortable task). He did it so he could create the opportunity of amassing enough contradictory evidence to compose a convincing refutation of his theory, which had much confirming evidence on the other side *that had to be outweighed*. By taking in contradictions and keeping them as so, one can accumulate enough of them to harvest their *learning power* and bring about a profitable *paradigm change*, as philosophers of science call it, and reach better theories of reality.

This makes it so *non-contradiction* is an informationally expensive principle, — what Testa refers to as *the cost of consistency* (cf. TESTA 2015), — as such contradictory states are necessary temporal stages of a rational scientific inquiry. The AGM model, however, is not constructed to deal with contradictions and makes contradictory systems explode.

TESTA (2015) argues formally that, while some clasically-oriented models are able to deal with contradictions by either isolating or suppresing them, this comes at the cost of losing the fruitful dynamic aspects of belief revision. Paraconsistent models, in turn, were tailored to deal with contradictions, and do so without any cost besides sacrificing the intuitive *law of non-contradiction*. These models keep closure, which extracts as much information as possible from the current state of knowledge (something which may be only desirable up to a point, as we have argued), but throws away its classicality and denies theoremhood to the *law of exposion*.

As has been mentioned in previous sections, paraconsistent models are able to preserve new information as well as old, even if they may be contradictory. This does not mean such models will keep contradictions permanently, but rather that they keep it in for long enough to profit from them. GIRARD & TANAKA (2016) noted that an operation that updates the state so as to resolve contradictions is still needed, a consideration which we gather is on the mark. Unless one accepts that the world may contain irresolvable contradictions, a quite unpopular position in metaphysics, one will hold that a non-contradictory state is ultimately desirable in truth-oriented inquiry. Previous models just failed to see how such non-contradictoriness is not *always* desirable.

Contradiction-tolerant models are able to maximize information gain, and thus fulfill informational economy, by harvesting the *learning power* of contradictory states, which is precisely what scientific communities do. Contradictory states seem to be necessary temporal stages of rational inquiry. It follows that paraconsistent models are to be preferred when one desires to model scientific rationality.

5. References

ALCHOURRÓN, C. E.; GÄRDENFORS, P; MAKINSON, D. C. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. In: *Journal of Symbolic Logic*, v. 50, pp. 510–530, 1985.

BAILLARGEON, R. How Do Infants Learn About the Physical World? In: *British Journal of Development Psychology*, v. 12, n. 1, pp. 9-33, 1994.

CHERNIAK, C. *Minimal Rationality.* Cambridge, UK: MIT Press, 1986.

CHOMSKY, N. *Rules and Representations.* Oxford, UK: Basil Blackwell, 1980.

GARDENFÖRS, P. *Knowledge in Flux:* Modeling the Dynamics of Epistemic States. Cambridge, MA: MIT Press, 1988.

HANSSON, S. O. Changing the Scientific Corpus. In: *Belief Revision Meets Philosophy of Science* (eds. OLSSON, Erik J. & ENQVIST, Sebastian), pp. 43-58, 2011.

______. Ten Philosophical Problems of Belief Revision. In: *Journal of Logic and Computation*, v. 13, pp. 37-49, 2003.

______. *A Textbook of Belief Dynamics:* Theory Change and Database Updating. Kluwer Academic Publishers, 1999.

HENDRICKS, V. F. *Mainstream and Formal Epistemology*. Cambridge, UK: Cambridge University Press.

KAHNEMAN, D.; TVERSKY, A.; SLOVIC, P. *Judgment Under Uncertainty:* Heuristics and Biases. Cambridge, UK: Cambridge University Press, 1982.

LAUDAN, L. *Science and Relativism:* Some Key Controversies in the Philosophy of

Science. Chicago, IL: University of Chicago Press, 1990.
LEVI, I. *The Fixation of Belief and its Undoing:* Changing Beliefs Through Inquiry. Cambridge, UK: Cambridge University Press, 1991.
MARTIN, E.; OSHERON, D. N. Belief Revision in the Service of Scientific Discovery. In: *Mathematical Social Sciences*, v. 36, n. 1, pp. 57–68, 1998.
SOBER, E. R. *Ockham's Razors:* a User's Manual. Cambridge, UK: Cambridge University Press, 2015.
TESTA, R. R. *Revisão de Crenças Paraconsistente Baseada em um Operador Formal de Consistência.* Doctoral thesis at University of Campinas, 2014.
______. The Cost of Consistency: Informational Economy in Paraconsistent Belief Revision. In: *South American Journal of Logic*, v. 1, n. 2, pp. 461-480, 2015.
TESTA, R. R.; CONIGLIO, M. E.; RIBEIRO, M. M. AGM-Like Paraconsistent Belief Change. In: *Logic Journal of the IGPL*, v. 25, n. 4, pp. 632-672, 2017. doi:10.1093/jigpal/jzx010
TESTA, R. R.; CONIGLIO, M. E.; RIBEIRO, M. M. *Paraconsistent Belief Revision Based on a Formal Consistency Operator.* CLE e-prints, v. 15 n. 8, 2015.
WHITEHEAD, A. N. *Science and the Modern World.* New York, NY: Free Press, 1925.

6. Acknowledgements

This paper was developed with semi-continuous assistance of professor Rafael Testa, and I have greatly benefited from reading his most recent works (most of which were written with Marcelo Coniglio and Márcio Ribeiro) and from our afternoon discussions. I see the present paper as an elaboration of discussions developed in his doctoral dissertation and subsequent papers. I must also thank professor Marcelo Coniglio, who was Rafael Testa's doctoral advisor and who first taught me and my colleagues the theory of belief revision. He also assisted professor Testa in writing the papers from which I have learned so much. I hope my informal, and perhaps naïve, discussions in this paper can help discussions in the field of belief revision snowball into a comprehensive inquiry on the relation between paraconsistency and scientific thought, even if the ideas outlined here should be ultimately revised or overthrown — once a *critical mass* of counter-objections has accreted in our collective epistemic state, of course.