Scheme-Based Alethic Realism: Agency, the Environment, and Truthmaking

MURAT BAÇ[1] and RENÉE ELIO[2]
[1]*Department of Philosophy, Bogaziçi University, 34342 Bebek, Istanbul, Turkey;*
E-mail: murat.bac@boun.edu.tr
[2]*Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada T6G 2H1;*
E-mail: ree@cs.ualberta.ca

Abstract. This paper presents a position called Scheme-based Alethic Realism, which reconciles a realist position on the nature of truth with a pluralistic Kantian perspective that allows for multiple "environments" in which truthmaking relationships are established. We argue that truthmaking functions are constrained by a stable phenomenal world and a stable cognitive architecture. This account takes truth as normatively distinct from epistemic justification while relativizing the truth conditions of our statements to what we call "Frameworks." The pluralistic aspect allows that these stable elements, while constraining representational and linguistic schemes, do not define a single framework for truthmaking relations. We strengthen this position by considering themes on situated rational agency from cognitive science and artificial intelligence, arguing that whatever enables or supports rational action within a particular environment must figure into some account of truth and truthmaking, and *vice versa*.

Key words: agency, artificial intelligence, cognitive science, conceptual schemes, framework, Kant, ontology, pluralism, rationality, realism, situated cognition, truth

1. Introduction

Alethic realism concerns the putative nature of propositional truth or truthmakers. In this paper, we present a position called Scheme-Based Alethic Realism, which aims to reconcile a realist position with a pluralistic Kantian perspective that allows for multiple "environments" in which truthmaking relationships are established. In brief, the position holds that truthmaking functions are constrained by a stable phenomenal world and a stable cognitive architecture. The pluralistic aspect allows that these stable elements, while constraining representational and linguistic schemes, do not define a single framework for truthmaking relations. Our goal in this paper is to develop and defend such a position and relate it to current views on agency and rational action that have emerged in the artificial intelligence and cognitive science areas.

There are three main interrelated philosophical points that comprise this position. First, we argue that there is nothing objectionable about the idea of positing an internal relation between truths and such semantic constructs as conceptual frameworks which mediate the makers of truths *and* allowing at the same time that there is some mind- and scheme-independent reality that sustains, and forms the

Minds and Machines **14:** 173–196, 2004.

basis of, those phenomenal truths. Second, we define and defend a certain criterion by which a given position on truth and truthmaking can be recognized as a form of alethic realism. Third, we present a form of alethic realism in which propositional truth is a joint function of a cognitive agent, its representational scheme (a term we will use interchangeably with 'conceptual framework') and its environment, i.e. a phenomenal reality. We maintain that there are no truths without these three components, and hence, truthmaking is neither completely external nor completely internal to cognitive agents.

We begin with some background remarks on realism, anti-realism, and why certain Kantian themes have crucial significance for our alethic realist account. We next review the interdisciplinary perspectives on agency and rational action from AI and cognitive science, which we believe offer a useful way to reframe and reconsider some of the philosophical issues about truth and truthmaking — for it seems that whatever enables or supports rational action must figure into some account of truth and truthmaking, and vice versa. With this background in mind, we consider more closely the various positions that attempt to address some of the inherent problems with a correspondence theory of truth. Our goal is to set the stage for a pluralistic sort of Kantianism, which tries to avoid the theoretical difficulties of both the realist and anti-realist views found in the epistemology/metaphysics literature. Our particular proposal in its essence embraces the notion that the truthmaking relation is determined neither purely ontologically nor in a fully pragmatic fashion. This account is similar to the position spelled out and developed by Hilary Putnam over a period of 15 years except that in our account, we aim to more precisely operationalize the correspondence relation within conceptual schemes. Moreover, our proposal is well in line, as we will try to show, with current conceptions of rational, situated agency as a function of an underlying architecture, representational frameworks that are enabled by that architecture, and the external environment.

2. Some Preliminaries: On Metaphysical and Alethic Realism

Broadly speaking, the position on truth and truthmaking that we present here is one that tries to offer a viable alternative to two epistemic–ontological extremes: absolutism and relativism. Such a treatment inevitably necessitates coming to grips with the debate between "realism" and "antirealism," labels that have proven fairly confusing and sometimes even theoretically counter-productive. These well-established pigeonholes are nonetheless useful in getting quickly into the heart of the philosophical issues we will take up here and, thus, we employ them as our point of departure.

In *Reason, Truth and History*, Hilary Putnam characterizes "metaphysical realism," famously, as follows:

> On this perspective, the world consists of some fixed totality of mind-independent objects. There is exactly one true and complete description of 'the

> way the world is'. Truth involves some sort of correspondence relation between words or thought-signs and external things and sets of things. (Putnam, 1986, p. 49)

Here Putnam is giving the name metaphysical realism to a position which is apparently not only about metaphysics — roughly, the philosophical investigation of the nature and structure of reality — but about our epistemic connection with the world via *truth* and also one *particular* sort of truth, viz., correspondence between symbols (e.g. words) and external reality. This position intermingles a stance on existence with a position on what counts as truth, which are separate, albeit related, matters. Such conglomeration of the different aspects of the matter is fairly typical of the way realism has been presented and discussed in the pertinent literature for the last few decades. However, it is something of a methodological hindrance and simply confusing to lump them together in analytic inquiries.

For our purposes here, we draw a distinction between two kinds or senses of realism: metaphysical and alethic. *Metaphysical Realism* (MR) concerns the existence of *objects* and, thus, it *prima facie* has no epistemic or semantic implications. The MRist argues (minimally) that the existence of the objects in our universe is independent of our beliefs, theories, and epistemic powers to form beliefs and theorize about the world. As such, $\mathrm{MR}_{\mathrm{MIN}}$ is the claim that there *are* things not depending on cognizers for their existence. Of course MRists need not, and typically do not, stop at such a modest level of ontological realism. Many of them defend what we may call $\mathrm{MR}_{\mathrm{MAX}}$, the thesis that there would still be objects existing *as such-and-such* (e.g. Mars existing *as a planet*) even if there were no human beings to perceive and think about them.

As we noted earlier, *Alethic Realism* (AR) is about the structure or nature of *truth* or *truthmakers*. Broadly speaking, the essential thesis of AR is that whether a statement is true or not is independent of our epistemic means such as justification, warrant, or evidence. Our evidence or beliefs about the truth of a statement — no matter how good, how consistent, how universally shared — is not what makes a statement true. Being understood this way, AR is silent about the possibility of knowledge or certainty. ARists often distinguish the "meaning of truth" from the "test of truth," hence generally postulating, implicitly or explicitly, some sort of logical gap between truth and evidence.[1] In this sense, an *anti*-AR position, as described and propounded by M. Dummett (1987), amounts to the *rejection* of such a logical gap and, therefore, of the idea that truth is logically/conceptually independent of evidence or justification.[2]

3. Agency, the Environment, and Truthmaking

The philosophical essence of the present paper can be articulated by reference to an underlying idea which finds its expression in a contemporary attempt to steer clear of two epistemic–ontological "pictures" that are often acknowledged to be

rather misguided.[3] There is, on the one hand, the objectivist world-view inherited presumably from the Plato–Aristotle tradition. In its most generic conception, this view holds that truthmakers are the fixed furniture of ontology and human agents strive to cognize the ready-made features of some external reality existing independently of them. This seems to imply *inter alia* a commitment to MR_{MAX} and an AR-ist commitment to some externalist correspondence theory of truth. Standing at the other extreme are various relativist views which aim to undermine the alleged gap between what is given to human cognition and whatever lies outside of it. Their most crucial (and controversial) claims are that there is no ontological realm beyond what is available to humans cognitively and conceptually and that there are no truth values beyond what we can "fix" by means of evidence and justification. While these are obviously distinct claims, anti-realists often deploy them in tandem in pointing out the weaknesses of realism.

In this epistemic-ontological context, Kant's original contribution is often acknowledged — by the relatively moderate — to be a highly interesting and significant one in that he attempts to maintain *both* that there is something beyond our cognitive reach *and* that what is epistemically available to us in general is objective, genuine knowledge. This insight, we submit, is not just path-breaking but *essentially correct*. Despite the fact that Kant advocated an idealist ontology, he was not, properly speaking, a subjectivist. This is the main reason why philosophers like Putnam sympathize with both the transcendental and empirical sides of his project while admitting, of course, that the universalist aspect of Kant's transcendental philosophy cannot be successfully defended in our era. In the preferred terminology of this paper, which will be developed fully later, it can be stated that the conditions of the possibility of empirical knowledge — *and* of the possibility of the objects of that knowledge — are determined not only by the stable aspects of our cognitive structure (architecture) and of reality but also by some representational (e.g. linguistic) framework supported by the cognitive structure.

This ontological-alethic picture can serve certain interdisciplinary as well as philosophical purposes. It seems reasonable to maintain that truth and truthmaking must be at the basis of rational action, or at least, at the basis of *successful* action by some cognitive agent in some environment. In the following two sections, we examine the inverse of this relationship, and ask how theoretical concepts of rational agency (or "success" in the world) inform our conceptions of truth and truthmaking. There is long standing discussion on what rational agency is and the standards against which it might be assessed, both for artifacts and for humans. This discussion continues and a detailed consideration of these themes is well beyond the scope of this article.[4] Here, we highlight a few key constructs from the artificial intelligence and cognitive psychology arenas that relate directly to themes that arise in a consideration of truth and truthmaking, semantic constructs, and the alethic realist position we later present.

3.1. ROBOTS MUST BE REALISTS

McCarthy (1974) long ago argued that the activity of AI should be "studying the structure of information and the structure of problem solving processes independently of applications and *independently of their realization in animals and humans*... [because] intellectual activity takes place in a world that has a certain physical and intellectual structure" (McCarthy, 1974, p. 64, emphasis ours). This view of intelligence is both Platonic (especially the italicized portion) and situated, at least insofar as its character is to be understood as governed in part by the structure of the environment. The world has certain properties that an artifact (or human) needs to discern if it is going to accomplish anything successfully in that world. It is a pragmatic prescription as well: understanding the structure of that world (however defined) will be the key to designing general and successful artificial cognitive agents. Now, the "world" for these artifacts may indeed be the "real" physical world (e.g. for robots) or it might be a software world. In brief, the environment is whatever phenomenal world the agent must sense and take actions in.

To build a successful agent, there must be a non-arbitrary mapping relationship between the input that the agent receives from the environment and its actions. Furthermore, many of the most interesting cognitive activities we might assign to such agents are not predicated on a direct mapping between sensed inputs and subsequent actions. Rather, there is some kind of internal manipulation of sensed inputs, and the more indirect the relationship is between the inputs and the resulting actions, the more complex that manipulation becomes. It is desired that the mapping function be sound: a derived truth must correspond to a "real" truth, at least to the extent that it allows the agent can succeed at whatever task it undertakes.

But soundness is not the end of the issue. Cognitive psychologists and AI researchers have long appreciated that the representational framework that serves as the concrete realization of this mapping function constitutes the strongest bias in what a reasoner or reasoning system can discern, infer or learn — easily or *at all* — about some external environment. These representational frameworks, even those that are governed by sound inference rules, are further understood to be only an imperfect account of what is "really going on" outside the agent. Pragmatically, AI system builders must ensure that the "right distinctions" are reflected, preserved, or discernable by the mapping function that moves from what the agent senses to what the agent does. Here, it is worth noting that an entire paradigm shift in AI was motivated by the assertion that the external world is its own best model (Brooks, 1991), and internal symbolic abstractions of it are going to be, by definition, impoverished and likely to fail in just those cases where the "right distinctions" did not get preserved by the abstraction. This is an important theme we will return to later. That said, even this situated cognition position does not get an agent designer out of the dilemma of devising non-arbitrary relations (based on *some* kind of abstraction) between the properties of an external environment and properties of

an agent's internal states that determine that agent's actions. It just does so under a different paradigmatic approach to the enterprise. The critical point to emphasize here is that even though cognitive/alethic success of the agents is of paramount importance, this does not yet undermine the contention that propositional *content* is generally indexed, in an essential way, to the states of affairs of a world that take place outside the cognitive or mental control of cognizers.

3.2. TRUTHMAKING IS NEITHER OUTSIDE NOR INSIDE A REASONING AGENT

The knowledge representation efforts in AI could be broadly construed as aiming to discern the best representations for preserving the truths about the external world within an admittedly imperfect internal model of that world held by an agent. It would seem, on the surface at least, that artificial intelligence embraces the realist perspective and correspondence theory of truth. But we have already gestured towards a problem with this simple characterization by contending that each different representational scheme constitutes a different bias for a cognitive agent, with respect to what it can sense, infer, or learn in an environment. In this section, we outline the more formal notions of rational agency from a cognitive science perspective that have a decidedly Kantian flavor and a pluralistic one at that.

The interplay between environment, cognitive architecture, and representational schemes has been at the heart of most theoretical conceptions of intelligence, behavior, or rationality, starting with Simon's (1982) notion of *bounded rationality*. In the context of a computational system, an architecture specifies invariant properties thought to define, constrain, and enable the ways in which the system achieves some particular functionality. Whether we are concerned with explanatory accounts of human cognition or prescriptive models for artificial cognition, it has long been appreciated that: (a) an architecture defines a set of primitive objects, processes, and organizational distinctions that are not further decomposed for explaining or producing the behavior of interest; (b) different primitives constitute different architectures; and (c) different architectures support or enable different algorithms, defining their "run time profile," e.g. how they might take to execute under different input conditions. Throughout the cognitive science literature, various analyses of agency have articulated and employed these ideas. Pylyshyn (1986, p. 31) introduced the term *cognitive impenetrability* to distinguish the invariant properties of the human cognitive architecture from the cognitive programs (behavioral models) that "run on" the architecture. Simply put, the idea is that what counts as being part of the human cognitive architecture is exactly that which operates outside the influence of what a person knows or believes, i.e. outside the influence of any program that is running on the architecture. This last point is crucial and has been considered in detail by numerous philosophical positions. A Kantian sort of realism with respect to this computational perspective accords well with the idea that the

programs which the mind can "run" are constrained by the very properties of the mind itself.

These themes resurface in formal considerations of agency and rational action adopted by the artificial intelligence community. To adopt the so-called agent paradigm, as currently and most prominently promoted by S. Russell and his colleagues,[5] is to embrace the notion of embedded or situated agency in a particular way that entails some prescriptive design methodologies for AI researchers. Specifically, this perspective defines the "agent function" to be a mapping function from some percepts — what the agent can sense about its environment — to actions that an agent can perform in that environment. With this abstract perspective in mind, the AI designer can then proceed to consider what the nature of that mapping function ought to be, *given* a particular environment with particular properties and a particular task or set of goals to achieve within that environment (Horvitz, 1989; Russell, 1997). This embedded agent view acknowledges the crucial role that the environment plays in defining agent performance, without requiring particular commitments on how that environment is internally modeled (or even if it is modeled at all). Rather, a particular configuration of environment properties (e.g. whether it is dynamic or stable, deterministic or indeterministic, completely or only partially accessible to the agent) coupled with a particular task specification does much to define properties of the mapping function (agent design) required to achieve some target behavior. One environment–task pair might require a complex internal model of the environment, complete with inferential capabilities and utility computations on possible worlds; another environment–task pair may require instead a sophisticated crafting of hardwired stimulus–response rules. A crucial element of this analysis is a performance measure — some way of evaluating the mapping function that is implemented to deliver the behavior of interest in that particular environment.

As Russell (1997) outlines, this general abstract perspective on agency leads to alternative formulations of rationality and intelligence. As we noted above, an underlying architecture or machine constrains many aspects of a particular program that might run on it. More formally, a particular agent can be defined as a particular program-architecture pairing. Russell defines the *bounded optimal* (agent) program as that program which, out of all the possible ones that *could* run on a given machine, maximizes the performance measure on a particular task in a particular environment. It is important to note that bounded optimality is not intended, in Russell's analysis, as property of abstract programs or machines, but of real programs running on real machines. Informally, the bounded optimal agent program performs as well as possible *vis-à-vis* a performance measure, given its particular computational resources, its task, and the environment in which it is situated. By this formulation, then, what counts as "the most intelligent" or "the most rational" agent program is a function of the architecture, the environment, a specific task, and a related performance measure. The agent program that is bounded optimal for a deterministic, completely observable environment is not likely to be the bounded

optimal program for a stochastic, partially observable environment. To paraphrase Goldman (1999), intelligence or rationality — recast as bounded optimality — is most definitely a success word for AI.

From this, it is a short jump to appropriating the notion of the bounded optimal agent function as a *bounded optimal truthmaking function* — where a "truthmaking function" is defined as one which takes as its arguments a cognitive (viz., architectural) structure, the programs that run on it, and a well-specified environment, and assigns truth values to bivalent sentences. The bounded optimal truthmaking function is the best actual truthmaking function that can be devised, *given* the underlying architecture, the environment-task pairing, and the success measure that is employed. This depicts truthmaking essentially in a Kantian fashion. It would be difficult to imagine how a concretely instantiated mapping function (program *plus* machine) could maximize a performance measure if it did not preserve some aspects of the environment in which the agent is operating. If we fully appreciate the environment as a parameter in determining what is bounded optimal, it seems clear that we must accept certain metaphysical and alethic aspects of realism.

We have pointed out a Kantian theme in current perspectives on situated rational agency within cognitive science disciplines. But is there a place for pluralism? It might seem that the bounded optimality position necessitates a *single* best agent function that maps environmental properties to performance within that environment, even when we relativize it to a particular performance metric for a particular task. In our view, that is not necessarily so. There could, in principle, be several agent functions that, once instantiated as particular programs on particular machines, that are indistinguishable from each other with respect to the performance measure. Each one of them might maximize performance, but in different ways that we cannot directly observe. We would still allow them to be equally good truthmaking functions in this respect. Two further points naturally arise here. First, we can ask just how different those truthmaking functions *could* be, if we hold the machine (architecture) constant, the task constant, and (naturally) the environment constant. Given the constraining relationship architectures impose on properties of algorithms, we might expect that several programs which are indistinguishable on a performance measure for some task–environment pairing to share certain fundamental characteristics. Thus, we could allow for *some* differences in the program part of the agent, while still expecting those differences to be negligible relative to what is common. The import of those differences *might* emerge (or not) given a *different* task within the same environment, or even a different performance measure on the *same* task. The critical point here is that the goodness or veridicality of truthmaking functions is completely relativized to task, environment, architecture, and some performance measure.

The second observation about plurality within this scheme is a more pragmatic one, and is manifest as an actual technical problem for AI researchers. Let us imagine these truthmaking-functions-instantiated-as-agent s are all populating the same environment, but each one constituted as a different program, possibly op-

timized for different tasks. How might these agents communicate with each other, if they are not running exactly the same program, i.e. based on the same mapping that determines what is sensed and what it means for what it is they must do? Depending on the differences, communication may be easy or it might be impossible. Much of the promise of agent research (as it is viewed today) rests on the notion of distributed, heterogeneous software entities — designed by different designers for different purposes — exchanging information, knowledge, and methods for problem solving. Knowledge sharing among such artificial agents entails a solution to what goes by the modest name of "the ontology problem" in multi-agent research. While various public and private enterprises have tried to specify, share, or provide general and domain specific ontologies, it is difficult to reach consensus on any particular way to carve up some domain or even on a methodology for deciding on which ontology might be best.[6] The pluralism of "equally good" truthmaking functions is not a problem when artificial agents — specially crafted to be successful at their own independent tasks — have no reason to exchange information. It is when they do. Crudely put, the multi-research community within AI *has* a pluralistic Kantian problem for itself, and continues to work towards some solutions.

By the above discussion, the mandate for AI researchers can be cast as finding bounded optimal programs for particular architectures that they might devise. Let us return to the matter of human cognition. While those who design artificial agents set their artifacts down in particular environments and apply a particular performance metric, the matter for human cognizers is different. As D. Cummins notes, unless you are a creationist, you do not believe that humans have been plunked down in a phenomenal world, with fully-fashioned cognitive machinery. Rather, you concede that the brain, like other aspects of biological creatures, has been shaped by natural selection. If you are a materialist, then you are committed (at least implicitly) to the view that the mind is what the brain does. She observes that if you accept both these premises, you are committed to the notion that the human cognitive architecture — the mind — was shaped by natural selection (Cummins, 2002, p. 134).

If we allow that the structure of the environment has shaped the mind, we allow that there are some external properties out there. But conversely, if we allow for this shaping process, then we allow that the mind in turn has adapted to deal with the structure of the environment. A stronger stance contends that the human mind has been "optimized somehow" to perform cognitive functions in an environment that has a particular structure (Anderson, 1991, p. 471). There are many premises and assumptions that are part and parcel of such positions when developed into positions on rationality (Stich, 1990; Stein, 1996). Philosophically, there is considerable controversy as to whether any evolutionary considerations contribute to epistemology (Campbell, 1974; Kim, 1988). Less controversially, we might concede that the human cognitive architecture has mechanisms, default representations, and properties that evolved to be *somehow* attuned to whatever environmental structure most directly impacted survival. Once that architectural

tuning was done, the kinds of programs that could "run" on it effectively and easily, such as linguistic representations, become constrained. Thisposition is sympathetic to the general view that evolutionary considerations do raise significant questions on matters of truth and on epistemology as a normative pursuit by humans, especially if we construe humans as merely a bounded optimal program for inquiry on these matters.

In sum, we have considered the AI activity of designing artificial cognitive agents because we can view that activity as essentially crafting a truthmaking function for artificial cognizers. The embedded agent perspective and the ensuing bounded optimality view is based on a central — but not exclusive — role of the phenomenal world in determining what counts as "the best" truthmaking function. These views bring us full circle to McCarthy's original characterization of intellectual activity as taking place in a world that has "a certain physical and intellectual structure," presumably independent of human cognizers. The implication of the bounded optimality view is that attempts to engineer artificial reasoning systems — and reverse engineer the human cognitive system — must be guided by considering both properties of that world structure and the kinds of truthmaking functions which enable reasoners to discern that structure in order to be successful in that world. As we have indicated earlier, truthmaking, conceived essentially in a Kantian fashion, is a joint function of several parameters: a cognitive agent's internal architecture, the representational frameworks and programs constrained by it, and the environment (or "reality") in which it is situated to achieve particular goals. Our view has the consequence that this function entails a complex constraining relation among those parameters and therefore truthmaking is neither completely external nor completely internal to a cognitive agent. This does not require a single truthmaking function or even one that is discernibly the best one; the latter is determined strictly by a particular measure of success on a particular task. Furthermore, if we are to make good theoretical sense of these parameters *and* the resultant cognitive product (viz., veridical representation), it seems sufficiently clear that there is no reasonable way to renounce certain metaphysical and alethic claims (or aspects) of realism. In what follows, we elaborate and recast these themes from a philosophical perspective.

4. Pluralistic Kantianism: Realism and Conceptual Frameworks

One crucial problem for the sort of realism Putnam attacked in his internal realist period is related to characterization of the ontology (or "environment") that is to constrain the truthmaking functions. This realism couples what we have dubbed "maximal metaphysical realism" with a magical theory of reference where the makers of truths are rendered external to the agents and, in Putnam's terminology, utterly dehumanized. In that case, the resultant alethic picture is obviously a chimerical and useless one for it fails to enlighten how such fully externalized

realities can play some actual role in constraining the non-arbitrary alethic relations that must be available to finite, situated, and pragmatically successful agents. A more viable alternative, of course, is to argue that the truthmaking functions (and the truth conditions of our empirical statements) are always determined and constrained by a stable cognitive structure and a stable phenomenal world — which is by definition open to cognition of the members of a linguistic community. These stable elements *inter alia* enable the cognitive agents to form various conceptual schemes, hence allowing them to make sense of their perceptual intake.

During the 1980s Putnam paved the way for a non-traditional sort of realism by maintaining that "*what objects does the world consist of?* is a question that it only makes sense to ask only *within* a theory or description" (Putnam, 1986, p. 49). In addition to this Kantian leitmotif, the later Putnam (1994) now admits that the semantic dimension of his internal realist account can allow for the idea of evidence-transcendent truth. This move obviously brings him closer to a frankly Kantian view of the truthmaking relation. Accordingly, the conditions of the possibility of veridical discourse are provided by our categories/schemes of interpretation and understanding. But *within such a scheme* the truthmaking relations are not formed in a haphazard manner. As we have stressed above, propositional truth is *indexed* (and not just accidentally) *to* representational systems and an environment in which such systems are employed to achieve particular goals. That common alethic success of actual agents is an incontrovertible fact seems to demand a philosophical account of propositional truth which pays due attention to the ontological, representational, and pragmatic aspects of the matter, as opposed to capitalizing on some of them to the exclusion of others.

What we thus obtain is a pluralistic Kantian position which preserves the transcendental spirit of Kant's philosophy while allowing that there may be a multiplicity of media or "environments" in which the truthmaking relations are established.[7] This kind of Kantianism holds not only that there is "something" beyond all human cognition and conceptualization (MR_{MIN}) and that truth conditions of our statements can only be formed and fashioned within the conceptual borders of human language, but also that there is no*singly* correct or veridical scheme in which truths are produced. The mind-independent reality cannot by and in itself manufacture objects and truthmakers of our world independently of cognitive agents' conceptual contributions. Once this point is granted, the door is opened for a more tenable account of truths, truthmakers, and the ontologies associated with them.[8] Moreover, pluralism about truth functions or truthmaking relation does not immediately invite some form of anti-AR into the picture. It is one thing to deny that the makers of our ordinary truths are licensed and fixed *sub specie aeternitatis*, yet another to reject that truth and epistemic justification are normatively distinct notions.

What needs to be specified and defended is how the phenomenal truthmakers are individuated and how they make our truthbearers true. We provide that specification in the section to follow. Before broaching that matter, however, we will briefly talk about the alethic side of realism and offer a criterion for it. In the

past, the gist of AR has been articulated in different ways in different contexts. B. Russell (1974), who attacked James's pragmatist ideas, put the matter in terms of the meaning-test distinction. Alethic realism has also been expressed by contrasting the truth conditions of our statements with the conditions of assertability for the speakers of particular languages. These different criteria seem to highlight different aspects of the same philosophical perspective and, thus, surely merit attention in our attempts to get an adequate picture of AR. One point to bear in mind is that contemporary philosophers approach the alethic issues from various angles with different intentions: some center their arguments around the concepts of meaning, understanding, and communication, whereas a considerable number of others do not regard them to be decisive or essential in a theory of truth. Consequently, it may be beneficial to have a criterion which broadly encompasses and demarcates most of the "realist" and "anti-realist" accounts that, one way or another, tackle the philosophical issues about propositional truth. We propose another criterion which, we believe, generally applies to the majority of the current perspectives:

> *The Discovery Criterion (DC) for AR*: A given philosophical perspective is a form of alethic realism if it endorses the idea that the truth or falsity of a truth-bearer is, generally speaking, a matter of discovery regardless of the ontological and epistemological commitments of that perspective. Accordingly, there may be a number of true or false propositions whose truth/falsity is not (or has not been, will not be, could not be, etc.) discovered by human cognizers.

The Discovery Criterion is a sufficient condition for AR, but we are not imposing it as a necessary one. It seems that this criterion can correctly identify the theories of both the ARists (e.g. Goldman, the later Putnam, W. P. Alston, M. Devitt) and the anti-ARists (e.g. R. Rorty, B. Allen), whose concern with AR or anti-AR lies chiefly within the boundaries of the highly disputed relation between propositional truth and epistemic justification or evidence.[9] The more interesting cases are, of course, those accounts which approach the matter from the perspective of meaning, understanding, interpretation, communication, and so on. Hence, Dummett — who describes AR in terms of our (alleged) understanding of the verification-independent truth conditions and finds it unintelligible — can be considered a strong anti-ARist in light of the above-given criterion. Furthermore, Davidson (1990) who defends truth-based semantics can also be regarded as a member of the anti-AR camp in that he is rather unlikely to endorse the idea that there may be some true statements which will never emerge in, or be brought to, discourse and interpretation. We also maintain that the Discovery Criterion is consistent with the bounded optimality perspective for agents, for the discovery of truth becomes a matter of rational (optimal, maximized) performance by an agent within some reality, regardless of what the actual instantiated agent program and architecture might be.

Given the initial plausibility of the Discovery Criterion, let us dwell on the essential idea behind it. Consider the following set of propositions:

$P_0(t_{\text{now}})$: There are no humans living on Earth at t_{now};
$P_1(t_{\text{now}})$: There is only one human living on Earth at t_{now};
$P_2(t_{\text{now}})$: There are two humans living on Earth at t_{now};
.
.
.
$P_N(t_{\text{now}})$: There are N humans living on Earth at t_{now}.

Take N to be a sufficiently large number such as 10^{15}. Our intuitions strongly suggest that at any given time, only *one* of $P_0, \ldots, P_N$ is true and all the rest are false. Besides, we also feel that this is a matter of *discovery* and, in this sense, truth is certainly "out there." As one can imagine, it may not be humanly possible ever to verify the truth of that particular P_i which happens to be true at a given t_{now}. Now, if the verificationist theories of truth are correct, we *cannot* say: "P_i *is* true now though it remains beyond our epistemic powers." The ARist, by contrast, believes that one of those N+1 propositions is true independently of our evidential capabilities. Stated in a slightly different way, there is something highly implausible about the idea that the true proposition above, $P_i(t_{\text{now}})$, suddenly becomes, or gets elevated to the state of being, true the moment we verify it. A more reasonable assumption would, of course, be to say that $P_i(t_{\text{now}})$ is true no matter what we can accomplish epistemologically.

Let us observe here that this particular question and the kind of answer we might discover for it are both constrained by a conceptual scheme and the interaction of that scheme and external reality. Thus, what counts as "alive" or as "human" — issues about which there are in fact serious debates — will be different under different conceptual schemes, and the answers returned by reality when probed with one or the other of these conceptual schemes will accordingly be different. But there is no doubt, as we elaborate below, that reality in some way or other is as much a determinant of the truth of $P_i(t_{\text{now}})$ as the conceptual scheme in which they are defined.

It is important to note here that by subscribing to the notion of non-epistemic truth, one is not thereby committed to the implausible idea that truth is totally independent of us. It has been argued by many realist philosophers that although truthmaking is inherently dependent on the existence of languages and language users — or more generally, on the existence of a representational system — this does not mean that individual truth values of our statements are also determined by human agents.[10] While our ways of categorizing the objects and events of the Cosmos invariably involve pragmatic elements and reflect our scientific and cognitive limitations (and, thus, could admittedly have been very different from what we currently have), we tend to think that a proposition like "Chlorofluorocarbon is harmful to the ozone layer in the earth's atmosphere" is either true or false independently of the pragmatic matters and considerations related to the contingent nature of ordinary human conceptualization.

5. Scheme-Based Alethic Realism: Frameworks and Situations

It is a desideratum that a scheme-based alethic theory lay out its theoretical tools as clearly as possible before embarking upon the task of developing a full-fledged account of phenomenal truthmaking. We will define to that end two terms, 'Framework' and 'Situation', and relate their role in the proposed alethic theory to our earlier discussions of situated agency. But let us first try to convey the underlying idea at a basic intuitive level.

Imagine an Ancient Greek (Thales), one of our contemporaries (Lisa), and an extra-terrestrial (ET), all looking at an ordinary telephone. Let's suppose that the cognitive functions of ET are so radically different from those of the humans that what it does "perceive" when it stares at a middle-sized dry object like a standard communication device bears no resemblance whatsoever to any of the perceptual objects available to human cognizers. According to the position we develop below, there can be no such truth as *q*: "there is a telephone in front of me" for ET. The same can also be said for Thales, but for reasons unrelated to the basic structure of his cognitive architecture. Even though Thales and Lisa are perceiving, physiologically speaking, the same "object," the ontological (or, rather, the phenomenal) maker of *q* is simply absent for the Greek philosopher. In this particular example, all three agents are assumed to have at least one thing in common: when they look at the telephone, they mentally isolate a certain part of the mind-independent reality and try to make sense of it in light of their epistemic/conceptual background. They all try, that is to say, to capture a meaningful "situation" out there by placing the object of their perception within a network of previous experiences and "world knowledge," understood in a broad manner. In a way, they try to situate that part of reality containing what we call a telephone within a "framework" fashioned and sustained by a finite set of cognitive functions and conceptual categories that are brought together as an interconnected and interdependent whole.

Let us now take a closer look at the alethic picture as seen from the perspectives of Lisa and Thales. When the Greek philosopher looks at a telephone, he confronts, strictly speaking, *not* the truthmaker of *q* but rather that of, e.g. *r*: "there is an unrecognizable black object in front of me" because the situation he actually encounters can be described or characterized — as far as "his world" is concerned — by *r*, not *q*. One can also say that the conceptual resources possessed by Thales allow or enable him to represent (or "frame") the content of his perception as *r* but never as *q*. Lisa's story, however, is altogether different: as an adult member of the modern world and a healthy cognizer, she can successfully recognize the truth conditions of *q* as she has not only the theoretical knowledge of the definition of a telephone but also the practical ability of *doing things* with that device. It is worth emphasizing that the latter function requires not only certain biological and cognitive capacities for action but also a social/cultural background that makes the usage of such a device possible. By our account, both sorts of factors contribute to her being able to take a certain portion of the mind-independent reality, to frame it as a more or

less definite occurrence of the phenomenal world that she lives in, and to recognize it as a situation which provides the phenomenal truthmaker for a proposition like q. In other words, Lisa's alethic success with respect to q springs ultimately from her practical success in properly recognizing and dealing with the objects of her world to achieve certain goals. From the embedded agency perspective on rational action, this is all that matters — the ultimate mapping of these percepts to truth values *and* successful actions (e.g. coming to the realization of the fact that q, making a phone call, or finding an object to use as a paperweight).

5.1. FRAMEWORKS AND SITUATIONS

We first spell out the definitions and then explain how they apply to actual cases in the phenomenal world of cognitive agents and how they illuminate truthmaking function (or relations) that take place in such a world.

> A *Framework* is a semantic abstraction from a given constellation (or network) of linguistic and non-linguistic communal practices of a given group of human cognizers such that (1) those practices enable the cognizers or practitioners belonging to the network to form and employ various tokens of symbols that are about parts or aspects of their world and also to communicate with each other, and (2) the abstraction represents or reflects the form(s) and/or style(s) of reasoning, communication and deliberate action of the practitioners.[11]
>
> An instance of *framing*, performed by an agent belonging to a linguistic community, can be defined as a (typically non-conscious, non-voluntary) act of employing, and imposing upon the external reality, a particular Framework by that agent.

We underscore at the outset that a linguistic system presumes (or can be viewed as) a kind of representational system. For our purposes here, we wish to equate a linguistic system with the more general notion of a representational system, although we have a special interest in dealing with linguistic ramifications. In particular, we indicate that framing is a process, and often an involuntary one at that — more about this later. This in turn implies that there is some default representational system at play, dictated largely by immutable features of the cognitive machinery.

> We now turn to the definition of 'Situation':
>
> A *Situation* is a set of framed circumstances that can be conceived by the members of a linguistic community as occurrences in their world — that is, occurrences that may affect their verbal and non-verbal actions. The representational Frameworks through which actual Situations are identified have a broader scope than mere assertoric or descriptive contents in a language. They operate upon a whole range of semantic and cultural elements that make up our lives.

In a nutshell, Situations are conceptually framed parts or aspects of reality, and Frameworks are our ways of getting into cognitive contact and dealing with it. Let us observe here that our exposure of the notions of Framework and Situation may give rise to questions about priority (probably in the form of some chicken-egg conundrum). But there is no puzzle here. From the ontological perspective, Situations are what constitute our phenomenal world. From the semantic-epistemic point of view, Frameworks are what we deploy in order to understand, and get knowledge about, the mind-independent reality.

5.2. TRUTH AND TRUTHMAKING

Another way to look at the notion of framing is to regard it as the involuntary imposition of an abstraction vocabulary on reality. An abstraction vocabulary consists of the "building blocks" of possible Situations. Those building blocks are typically (physical or non-physical) objects, properties, and relations that combinatorily make up our phenomenal world. Our concepts of a tree, a telephone, money, nostalgia, being white, being expensive are all elements of our abstraction vocabulary. These are all *semantic* abstractions, i.e. our conceptual means of framing parts/aspects of reality against our cognitive background. Representation of a Situation, on the other hand, is a second-order abstraction where the elements of the first-level are cognitively brought together to produce, for instance, judgments of the form *S* is *P*. The higher level operations of this sort are essentially *alethic* abstractions: the agent cognitively derives (or produces a representation of) an occurrence in the phenomenal world. Let us try to clarify these ideas.

An agent's success about first-order abstractions is related principally to having an adequate understanding or conception of the object (or property or relation) in question. In our example, Thales fails to recognize that there is a telephone in front him because his abstraction vocabulary does not contain that object in the first place. Lisa's system, on the other hand, is fully capable of doing the first-order abstraction; her understanding of the concept of a telephone is flawless. Nevertheless, this does not guarantee a success at the higher (viz., alethic) level. Suppose Lisa mistakes a black cat for a telephone and forms the propositional attitude that *q* whereas she is supposed to have the propositional attitude *v*: "there is a black cat in front of me." In this case, we would say that the *hypothesized Situation* (*q*) fails to match the *obtaining Situation* (*v*) and, thus, the "derived output" fails to be a truth. Put differently, the second-order cognitive operation gives rise to misrepresentation of reality or a false proposition.

It is in the above-described sense that a framework can be thought of as imposing an abstraction vocabulary out of which the phenomenal world is defined for the agent. This is where the discovery criterion for alethic realism presented earlier enters into the picture: the truth or falsity of a proposition is a matter of discovery (minimally, checking against some external reality). But an agent can

only ascertain the veracity of what she *can* phenomenally experience or frame. The truth associated with any particular situation is discover*able* only if the constituent elements pertaining to that situation are realized through the first order abstraction.

Thales's abstraction vocabulary or framework might allow second-level abstractions (that is, recognition of Situations) involving a black cat or a shadow, but not a telephone. When Thales frames what Lisa calls a telephone in his own peculiar way, i.e. as an "unrecognizable black object," he effectively imposes a fairly different abstraction vocabulary on reality and thus can formulate or recognize a fairly different Situation — which nonetheless bears *certain family resemblances* to that experienced by Lisa. It is important to note here that as far as *r* is an accurate description of Thales's belief-content and of the Situation he encounters, his hypothesized Situation perfectly matches the corresponding obtaining Situation; and the derived output (e.g. the thought or articulation of *r*) signifies Thales's alethic success.

According to the present account, the notion of a Situation enters into the equation twice: semantically/cognitively, in the *formation* of the hypothesized Situation (as a belief, statement, etc.) and, ontologically, in the phenomenal reality's actually "*containing*" — as viewed from the cognitive and conceptual perspective of the community to which the agent in question belongs — that hypothesized Situation. In the alethically happy cases, the truthmaking function is said to produce the intended result (truth) and the agent can be regarded as successful in that regard. We maintain that the truthmaking relation is both Kantian (in that the conditions of the possibility of knowledge are provided partly by our cognitive and conceptual resources) and pluralistic (in that those conditions are determined in a multiplicity of ways and must be indexed to the conceptual contributions of linguistic communities). One fundamental claim of pluralistic Kantianism is that the truth functions or satisfaction conditions of empirical propositions cannot be individuated in an initself reality independently of the cognitive and conceptual contributions of actual agents. Cognizers' ability to "isolate" (i.e. sense and perceive) some part of reality as an identifiable situation invariably plays an essential role in *determining* the makers of propositional truths — though this does *not* mean that cognizers "make" reality.

Of course, the kind of ability to perceive an identifiable situation is hardly a personal or individual feat. One must insist, in a Wittgensteinian spirit, that the *conceptual categorization* of the world by finite cognizers is always a communal affair, never a singly carried out individualistic act. (At least when those cognizers must or wish to communicate, otherwise they would have "the ontology problem" that confronts the artificial agents of today.) Consequently, without the speakers of particular languages there can be no truthmaking relations, for the absence of linguistic media implies the absence of some actual circumstances under which our statements or propositions can be said to have satisfaction conditions. Mutual understanding among speakers of particular languages is the task of communicating agents; and the speakers (their linguistic representation) are the instantiation

of some particular (truthmaking) mapping of that external reality to actions in that reality (e.g. utterances).

Would there be a telephone in the universe if all finite cognitive agents (including the above-mentioned trio) suddenly vanished? The answer may seem affirmative for those who are familiar with and sympathetic to Lisa's world, but one must exercise caution here. Seen from ET's — and presumably even Thales's — viewpoint, there would be a significant philosophical difference between saying "the telephone exists independently of finite cognizers" and "the metaphysical basis of what Lisa calls 'a telephone' exists independently of finite cognizers." A pluralistic Kantian perspective should maintain that we are justified to make claims of the latter sort but not the former. This is principally because the first statement implies that the objects and truthmaking relations can establish *themselves* in the external reality irrespective of the intensional aspects of the world inhabited by finite agents such as ourselves. Such a position is, as we have argued above, a patently untenable one.

The account of Frameworks and Situations offered here is a natural upshot of the contention that actual cognizers are, and must be, "situated." The truthmaking function — which yields truth values as output — basically necessitates the presence of a mind-independent reality and certain "ways" of making sense of that reality, i.e. abstraction vocabularies through which an agent comes in contact with this reality. The former of these two components can never generate propositional truths on its own. The latter component, on the other hand, consists of a cognitive architecture and "programs" running on and constrained by it. But what is the philosophical connection between different Frameworks and the mind-independent reality? We maintain that the answer to this question comes from "Kantian" considerations and relates to our earlier considerations of a common, evolved cognitive architecture. We might say that the common architecture of human cognizers is the foundational zero-th order framework (abstraction vocabulary) imposed on the environment. Thus, the agents whose cognitive structure is essentially the same due to belonging to the same species cannot in general have radically different (first and second order) abstractions of external reality. From a naturalist point of view, the cognitive agents and the mind-independent reality in which they find themselves *cannot* fail to inform or affect one another. Any particular framework that can yield propositions and their status as truthbearers is determined as a function of their veracity *vis-à-vis* external reality. A Situation like snow's being white is a joint product of the external reality, human cognition, the linguistic (representational) tools of human agents, and the specific task to which such a state of affairs is being considered. Therefore, both the traditional correspondence theory and its customary rivals (idealism or antirealism) are mistaken theses. The correct response to the question of where the phenomenal truthmakers are located is that they really do not "reside" anywhere determinate (such as the external reality or human mind).[12] Situations are neither merely subjective representations nor fully autonomous entities. As Putnam would say, the linguistic and conceptual resources

of agents and the external reality together give philosophical birth to Situations and, hence, more narrowly, to truthmakers. To restate it from our perspective here, the default Framework constituted by the (human) cognitive architecture coupled with external reality allows a certain realm or set of Situations, and indirectly, defines truthmaking as a joint function of a cognitive agent, its representational system, and particular tasks paired with a phenomenal world in which it is situated.

5.3. ON STATES OF AFFAIRS AND VOLUNTARY FRAMING

As must have been clear thus far, we prefer the concept of Situation to that of "state of affairs" since the latter has a long philosophical history of being associated narrowly with the notion of sentence-like ontological entities like facts or propositions. Broadly construed, there can be observational (or empirical), scientific, religious, poetical, and spiritual Situations. Consequently, the common notion of a fact or state of affairs is considerably more restricted in scope than what we call a Situation. The concept of a Situation is arguably more flexible to encompass vague, indeterminate, and heavily contextual cases (as opposed to "hard facts" such as "snow is white" or "there are three books on Brown's desk at time *t*") and, hence, is more likely to serve better as a conceptual tool in a *general* theory of truth. In particular, we doubt that the notion of a fact or state of affairs can handle the amazing variety of circumstance types that make up our world(s). (Think about a scientist's describing an astronomic phenomenon, a poet's characterization of autumn rain, our ordinary discourse about the absence of an object in a given context, and statements/circumstances/practices related to human nature, God, bad luck, inflation, overpopulation, broken promises, meditation, and so on.) After all, strictly propositional "truthmaker Situations" — which relate to distinct, unambiguous, objectively observable empirical contexts and which often get the honorific title of "facts" — are just a subclass of the whole set of Situations we can capture or conceive in our world.

Let us point out here an important distinction between *voluntary* and *involuntary* framing. The former kind, which requires what the cognitive scientists call "effortful cognitive processes," is operative when, for instance, a physicist decides to employ corpuscular, rather than wave, mechanics in order to solve a given problem about the propagation of light. More generally, we are not concerned with framing in the sense of theory proposal and revision, i.e. effortful construction of theoretical constructions to explain a system or set of events. Here, we have been primarily concerned with the notion of involuntary framing, the means by which we come in contact with the external world, and how that external world itself has shaped the very means by which that contact is made and interpreted. We do take it as an important consequence that such involuntary framing influences aspects of voluntary framing, i.e. that the evolved cognitive architecture constrains the subsequent abstractions that we might voluntarily impose.

According to AR, truth is a robust and "objective" notion. However, it is a fact that a posteriori statements exhibit an amazing range of variety and diversity across the epistemo-semantic spectrum. Employing a variant on Dummett's common example, "Smith is a terrorist" is a statement whose truth value must be "discovered" *and* also a statement which needs a great amount of interpretation — more than, say, a statement such as "There are six apples in this basket." Both statements, however, are *made true* (or false) by a world of Situations which is generated out of a collaboration of the external reality and the Frameworks that various cognitive agents impose on it. There is, in other words, no plausibility in the idea that "mere consensus" or similar pragmatic considerations would suffice to create truths. Thus, certain negative thoughts and feelings that one might feel towards Smith cannot *make* him a terrorist if he does not have the properties of being a terrorist *as defined by a particular Framework coupled with a particular reality with respect to some task and some metric by which that task's success is judged.*[13] Needless to say, there will always be cases where the truth value of a given statement will be decided chiefly epistemically (e.g. in assessing the artistic value of a sculpture). Yet, a general theory of truth ought to aim at illuminating the onto-semantic structure of *most of* our truthbearers, admitting that it can explicate better those cases residing on the more "robust" side of the alethic spectrum — such as directly observable and intersubjectively verifiable circumstances.

6. Conclusion

During the last decade, some philosophers with Kantian *and* Wittgensteinian tendencies have tried — with some success, in terms of credibility — to improve realism's track record by combining and supplementing their alethic views with certain contextualist and pragmatist elements. The present paper is meant to constitute another attempt in the same direction, viz., in showing how these two notions can go hand in hand in a congruent manner. The way we see it, a pluralistic Kantian theory is a linguistic (or relativistic) version of Kant's transcendental account that is, nonetheless, aligned with a nonepistemic conception of propositional truth. We believe that there are several sympathetic and consistent themes from the cognitive science literature, in particular the prevailing views on rational, situated agency. During the 1950s and 60s cognitive psychology emerged essentially as a "Kantian" alternative to the then dominant behaviorist trends. Therefore, it is just natural that the contemporary cognitive scientists and AI researchers build their theoretical frameworks upon certain Kantian (Alethic Realist and scheme-based) premises. We have tried to highlight some of these themes, emphasizing especially how the realist's intuitions are shared to a large extent by the cognitive scientists, and vice versa. One philosophical linchpin of this account is that truth of a statement is a matter of its veridically describing the relevant Situation that is constructed in reality and recognized within a representational Framework. There are truths (or, rather, truth values) which we do not know yet and which we may

fail to discover forever. Our process of "discovering" the truth value — given that Framework — is the process of formulating and posing the question to reality. Having done so, we have (voluntarily *and* involuntarily) imposed our Framework on reality and "read off" the answer to our questions accordingly. It is important to stress that while there is a mind-independent reality in the absence of humans or their representational systems, it is not possible to identify any *particular* Situation without reference to a *particular* Framework. (To claim otherwise would be tantamount to saying that the mind-independent reality can somehow manage to produce truthmakers for finite agents irrespectively of their actual or possible conceptual contribution.)

The Framework/Situations perspective outlined above introduces semantic constructs into the consideration of truth and truthmaking and hence is prone to attract criticisms of a Davidsonian kind. However, like other contemporary pluralistic or linguistic Kantians, we find the Davidsonian portrayal of conceptual schemes seriously misleading.[14] Lynch (1998) suggests that our schemes of concepts can be understood as exhibiting certain resemblances and differences between one another. The same can be stated about Situations and Frameworks. An act of framing where, say, an agent carefully observes the clouds and tries to forecast the weather, cannot be vastly different across different cultures. And the way we now frame the phenomenon of rain must have great similarities with that of the ancients. Just like we cannot be cut off from the mind-independent reality, we are never in our conceptual "islands" vis-à-vis various ways of framing that reality and representing the phenomenal Situations to ourselves.

Two rather crucial differences between the theory sketched here and the earlier attempts to specify the semantic-ontological "dynamics" of the truthmaking relation are that (1) our account points out the irreducible interrelations among the notions of environment, agency, and cognitive architecture, and (2) it provides a relatively detailed explanatory picture of the makers of our phenomenal truths and the semantic medium that renders the truthmaking relation possible. We implicitly adopt Goldman's notion of truth as a "success word" but have tried to couch it in quasi-evolutionary terms. In juxtaposing discussions of rational agency, particularly from the bounded optimality perspective, and truthmaking from a philosophical perspective, we have tried to close the gap between these two notions. To be an agent is *to be* a concrete instantiation of some truthmaking function. To be a successful agent (by some metric applied to some purpose) that operates with an external reality requires a truthmaking function which involves that external reality. The account offered here makes no assertions about the "nature" of the truthmaking function, only that it consists of external reality, some cognitive machinery, and a particular conceptual Framework supported by that cognitive machinery. The Frameworks are, in some fundamental way, constrained by the cognitive machinery. The cognitive machinery in turn has been tuned (in some fundamental, if not necessarily optimal, way) by external reality. Consequently, propositional truth is both non-arbitrary and non-noumenal.

According to the pluralist perspective, there may be numerous "true stories" of the external world. Two different Frameworks would enable two different truth-making functions, as we have defined the notion here, because they are defined by different abstraction vocabularies. They can be equally acceptable if the truth-bearers that result are orthogonal to each other, or are silent on the very matters that might distinguish them, or are described at different levels of abstraction such that distinguishing features cannot be discerned or perhaps do not matter. We have talked about this possibility in the context of bounded optimal agent programs being indistinguishable from each other on a given performance measure — hence they all appear to be equally "right" in their ways of coming in contact with and hence operating within reality to achieve some particular purpose. Still, we can allow, from an operational or pragmatic viewpoint, that there might be one "true" story (in the sense of ultimately proving to be the "best of the bunch") *when*, and only when, it is relativized to a certain task and evaluated by a certain success criterion. The pragmatics of the matter is as crucial as the semantic aspect; but it would be a fundamental mistake to conflate the two in our attempts to articulate a viable account of the finite agents' alethic undertakings in a world they find themselves situated in.

Acknowledgements

Scheme-based Alethic Realism and its relation to other philosophical views were developed by the first author and are presented in more detail in Baç (2001). This paper greatly benefited from the comments and criticisms of Bruce Hunter and Bernard Linsky. Murat Baç would also like to acknowledge the research funding provided by the Killam Foundation, Canada and Renée Elio acknowledges funding provided by the Natural Sciences and Engineering Research Council of Canada.

Notes

[1] See, e.g. B. Russell (1974).

[2] This classification is not meant to reflect exactly what all contemporary philosophers have in mind when they talk about those realisms. For example, the Metaphysical Realism defined above is different from that described and attacked by Putnam who takes such realism to be not only about the external objects and facts but also about our alethic connection with them. Similarly, our depiction of AR differs from the account William Alston offered in his *A Realist Conception of Truth*, in that Alston's alethic realism comes with certain implications about the world. See especially his (1999, pp. 49 and 84).

[3] Putnam (1994, p. 446) calls these two positions "reactionary metaphysics" and "irresponsible relativism," respectively.

[4] The interested reader is directed to Stein (1996), Stanovich (1999) and Elio (2002).

[5] The basic principles of this position that we sketch here, after percolating within the AI community, emerged as the foundational theme for a leading AI textbook. See Russell and Norvig (1995).

[6]One methodology is for humans to introspect on what the ontological distinctions. Another method — philosophically sympathetic to a linguistic Kantian position — is to derive the ontology *from* language use. Schubert and colleagues regard general knowledge in texts as defining relationships implied to be possible in the world, and employ statistical analysis techniques to derive those possible relationships as the foundation of an ontology that supports common sense reasoning. See Schubert (2002).

[7]The idea of semantic pluralism has been defended recently by, e.g. Lynch (1998), Wright (2001), Horgan (2001).

[8]See Baç (2003) for a pluralistic Kantian account of the ontological basis of truthmaking relations.

[9]See Goldman (1986, 1999), Putnam (1994), Alston (1996), Devitt (1991), Rorty (1991), Allen (1995).

[10]Several realist philosophers have pointed out the kind of distinction drawn here. See, e.g. Goldman (1986, pp. 155–156), Putnam (1987, p. 20), Lynch (1998, pp. 137–138), Searle (1995, p. 166), Goldman (1999, p. 20), Baç (1999, pp. 204-205).

[11] The expression 'styles of reasoning' was used, e.g. by Hacking (1998).

[12]And this is the major ontological difference between our Situations and, for example, Austin's and Warnock's use of the term. See Warnock (1962) and Johnson (1992).

[13]Of course, one can lead people around to thinking that he is a terrorist; but, again, this would not be sufficient to make him so, given a particular Framework coupled with a given reality out of which Situations are conceived.

[14]This is an issue we cannot broach within the confines of this paper. For convincing critiques of Davidson's minimalist semantics, see Malpas (1992), Hacking (1998), Lynch (1998). See also McDowell's remarks (1999, p. 91).

References

Allen, B. (1995), *Truth in Philosophy*, Massachusetts: Harvard University Press.

Alston, W. P. (1996), *A Realist Conception of Truth*, London: Cornell University Press.

Anderson, J. R. (1991), 'Is Human Cognition Adaptive?', *Behavioral and Brain Sciences* 14, pp. 471–485.

Baç, M. (1999), 'Propositional Knowledge and the Enigma of Realism', *Philosophia* 2(1–2), pp. 199–223.

Baç, M. (2001), 'Alethic Actualism: A Quasi-Realist Theory of Truth and Knowledge', unpublished doctoral dissertation, University of Alberta, Edmonton.

Baç, M. (2003), 'The Ontological Status of Truthmakers: An Alternative to Tractarianism and Metaphysical Anti-Realism', forthcoming in *Metaphysica*.

Brooks, R.A. (1991), 'Intelligence Without Representation', *Artificial Intelligence* 47, pp. 139–159.

Campbell, D. T. (1974), 'Evolutionary Epistemology', in P. A. Schilpp, ed., *The Philosophy of Karl R. Popper*, Illinois: Illinois Open Court.

Cummins, D. D. (2002), 'The Evolutionary Roots of Intelligence', in R. Elio, ed., *Common Sense, Reasoning, and Rationality*, New York: Oxford University Press.

Davidson, D. (1990), 'The Structure and Content of Truth', *Journal of Philosophy* xxxvii, pp. 279–328.

Devitt, M. (1991), *Realism and Truth*, Oxford: Basil Blackwell.

Dummett, M. (1987), *Truth and Other Enigmas*, London: Duckworth.

Elio, R. (2002), *Common Sense, Reasoning, and Rationality*, New York: Oxford University Press.

Goldman, A. I. (1986), *Epistemology and Cognition*, Cambridge: Harvard University Press.

Goldman, A. I. (1999), *Knowledge in a Social World*, Oxford: Clarendon Press.

Hacking, I. (1998), 'Language, Truth, and Reason', in L. M. Alcoff, ed., *Epistemology: The Big Question*, Malden: Blackwell Publishers.

Horgan, T. (2001), 'Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence', in M. Lynch, ed., *The Nature of Truth*, Cambridge: MIT Press.

Horvitz, E. (1989), 'Reasoning About Beliefs and Actions Under Computational Resource Constraints', in L.N. Kanal, T. S. Levill and J. F., eds, *Uncertainty in Artificial Intelligence 3*, Amsterdam: North Holland, Elsevier Science Publishers.

Johnson, L. E. (1992), *Focusing on Truth*, New York: Routledge.

Kim, J. (1988), 'What is "Naturalized Epistemology"?', *Philosophical Perspectives 2, Epistemology*, pp. 381–405.

Lynch, M. P. (1998), *Truth in Context: An Essay on Pluralism and Objectivity*, Cambridge: MIT Press.

Malpas, J. E. (1992), *Donald Davidson and the Mirror of Meaning*, Cambridge: Cambridge University Press.

McCarthy, J. (1974), 'Review of Artificial Intelligence: A General Survey', *Artificial Intelligence 5*, pp. 317–322. Reprinted in V. Lifschitz, ed., (1990), *Formalizing Common Sense: Papers by John McCarthy*, New Jersey: Ablex. Page references to reprint.

McDowell, J. (1999), 'Scheme-Content Dualism and Empiricism', in L. E. Hahn, ed., *The Philosophy of Donald Davidson*, Chicago: Open Court.

Putnam, H. (1986), *Reason, Truth, and History*, London: Cambridge University Press.

Putnam, H. (1987), *The Many Faces of Realism*, LaSalle: Open Court.

Putnam, H. (1994), 'Sense, Nonsense, and the Senses: An Inquiry into the Powers of the Human Mind', *The Journal of Philosophy* xci, pp. 445–517.

Pylyshyn, Z. (1986), *Computation and Cognition: Toward a Foundation for Cognitive Science*, Cambridge: MIT Press/Bradford Books.

Rorty, R. (1991), *Objectivity, Relativism, and Truth*, Cambridge: Cambridge University Press.

Russell, B. (1974), *The Problems of Philosophy*, Oxford: Oxford University Press.

Russell, S. J. (1997), 'Rationality and intelligence', *Artificial Intelligence* 94, pp. 57–77.

Russell, S. J., and Norvig, P. (1995), *Artificial Intelligence: A Modern Approach*, Englewood Cliffs, NJ: Prentice Hall.

Schubert, L. K. (2002), *Can We Derive General World Knowledge From Texts?*, Human Language Technology Conference (HLT 2002), San Diego, CA, 24–27, March.

Searle, J. (1995), *The Construction of Social Reality*, New York: Free Press.

Simon, H. (1982), *Models of Bounded Rationality*, Cambridge: MIT Press.

Stanovich, K. E. (1999), *Who is Rational?: Studies of Individual Differences in Reasoning*, New Jersey: Lawrence Erlbaum.

Stein, E. (1996), *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*, Oxford: Clarendon Press.

Stich, S. (1990), *The Fragmentation of Reason*, Cambridge: MIT Press.

Warnock, G. J. (1962), 'Truth and Correspondence', in C. D. Rollins, ed., *Knowledge and Experience*, Pittsburgh: University of Pittsburgh Press.

Wright, C. (2001), 'Minimalism, Deflationism, Pragmatism, Pluralism', in M. Lynch, ed., *The Nature of Truth*, Cambridge: MIT Press.