The Belief Cloud: Framework and Applications*

Ricardo Dahis

May 30, 2020

I Introduction

This essay attempts to describe a framework that is taken for granted in philosophy, science and public debate but, to the best of my knowledge, has not yet been succinctly spelled out in text from first principles. At minimum it serves as an useful synthesis, but, hopefully, it also clears space for advancing new ideas.

The belief cloud is a conceptual visualization of every claim humans have ever come up with, how they relate logically, and the degree of belief in each of them. It divides claims into facts, and those others that are still open for debate. The cloud represents the current state of our knowledge, and it also helps visualizing how it grows. Moreover, for example, we can apply this idea to communication. If each person is represented by a subjective belief cloud, then attempts at progress in communication are about "intersecting" clouds and aligning what is in the collective uncertainty space. All these ideas will be explained in more detail below.[1]

I have two goals with this essay. First, on a personal level, I hope it contributes to a clearer understanding of what knowledge is, where it starts and where it ends. Tracing an explicit line around what we know will simultaneously clarify what we are ignorant of. And ignorance is no crime. In fact, we would all benefit from a healthy dose of skepticism, humility and curiosity. On a moral level, we should strive to maximize the flourishing, meaning and well-being of conscious beings. But that pursuit

*I would like to thank Amanda de Albuquerque, Mariana Carvalho, Travis Chamberlain and Tushmit Hasan for valuable comments. Email: rdahis@u.northwestern.edu

[1]The idea is hinted at in an anecdote about Richard Feynman after receiving the Nobel Prize in physics: At dinner with the queen of Sweden [...] she asked me what I won the Nobel Prize for. When I said "Quantum Physics", she said, "Oh, we can't talk about that, because nobody understands it." and I said "On the contrary, we know quite a lot about quantum physics, and *that's* why we can't talk about it. It's everything else we *don't* know about - like how to solve poverty, and lower crime, and stop drugs, that we *can* talk about!"

should be compatible with what we know about the Universe. We do not need to lie to ourselves with unjustified faith or superstition when contemplating existence or facing off nihilism. May this essay be a small step in that direction.

Second, the framework presented below helps society to collectively contextualize ideas and speech into a coherent structure. We increasingly need to clear space for better communication, particularly in times of political polarization, religious sectarianism, and mutual intolerance (Levitsky and Ziblatt, 2018). Ideas matter, and the only alternative to violence is conversation. Sound government policy and, ultimately, peace, depend on it.

A few caveats before we start. First, the essay takes an abstract approach to human affairs and behavior. This is for the sake of clarity: I want to emphasize an ideal for which we should strive without necessarily accounting for all the complexities an intellectually honest thinker would demand. Second, it purposefully avoids long-lasting debates in metaphysics and assumes the default positions (Searle, 1999).[2]

II Framework

We start from first principles and build from there. The framework depends only on basic set theory and standard concepts in formal logic.

Definition 1. A *claim* is an assertion with propositional value.

I represent a claim as c. It may expressed in any language (words or symbols).

Definition 2. The *claim space* C is the set of all potential claims.

The claim space may, in principle, be infinite. For the purposes of this essay, it only needs to be large. They can be arrived at by guessing, by logical deduction or induction, or by any other method. As discussed in Dahis (2018), claims may be *verifiable*, *falsifiable*, *realistic*, *complete*, among others. The claim space is divided into *affirmative* claims (e.g. "plants perform photosynthesis"), and *entailment* claims (with a $c_1 \implies c_2$ form, e.g. "if it rains today, then we will stay at home"). It is also divided into claims about the *external* world and claims about *internal* states and experience.

Definition 3. A *theory* is a set of claims $T \subseteq C$.

Definition 4. A claim $c \in C$ is *true* if it corresponds to reality.

[2]Among others, default positions include the following. There is a real world that exists independently of us. We have direct perceptual access to that world through our senses. Causation is a real relation among objects and evens in the world.

Definition 5. The *belief*, *likelihood*, or *probability*, that a single claim c is true is expressed as a function $p : C \to [0, 1]$.

First, notice that beliefs p are in the mind of the observer, as in Bayesian statistics. Being perfectly ignorant of a claim's truth value is equivalent to an *uninformed prior* of $p(c) = 0.5$. Having knowledge about its truth value is having $p(c) \to 0$ or $p(c) \to 1$. Abusing notation, the vector $P = p(C)$ represents the state of our beliefs.

We cannot objectively *know* whether claims are true or not.[3][4] Therefore, in human affairs, we arbitrarily pick thresholds $\overline{p}$ and $\underline{p}$ such that claims with $p(c) \geq \overline{p}$ are treated as true and claims with $p(c) \leq \underline{p}$ are treated as false.[5][6] Interestingly, this notation also allows for changing beliefs over time, such that claims considered true before are later proved false (or incomplete, or insufficient).[7]

Definition 6. The *evidence space* $E \subseteq C$ is the set of claims expressing all measurements, observations, experiments, and data ever recorded.

A big question in philosophy is what *justifies*, or *supports*, belief. Briefly, standard answers are the following. First, belief may be justified by objective reasons, i.e. direct verifiable evidence, not being refuted by tests and experiments, or being a logical consequence of other objectively justified claims. Other sources of justification, albeit not in the standards of science, are revelation, introspection or intuition. Additionally, since Plato and Aristotle, a standard definition of knowledge is of *justified true belief*.[8]

Definition 7. A set of *facts* F is composed of claims of sufficiently high probability, namely $F = \{c \in C | p(c) \geq \overline{p}\}$.

Facts F include claims directly expressing evidence, i.e. $E \subseteq F$ (e.g. "water is made of CO2 molecules")[9], and may contain claims assumed or indirectly considered true.

Definition 8. A set of *assumptions* A is composed of claims taken as fact but not necessarily justified by evidence.

[3] This claim bypasses two thousand years of philosophy and epistemology about realism, skepticism, the problem of induction, among others.

[4] This discussion applies basically only to external claims. To know whether internal claims are true or false, one only needs to introspect and to see clearly (Harris, 2014).

[5] For simplicity I choose not to discuss (1) how $\bar{p}$ is chosen in practice, and (2) whether such a threshold always exists.

[6] It is not necessarily the case that $\overline{p} = 1 - \underline{p}$.

[7] See Popper (1959), Kuhn (1962) and Feyerabend (1975).

[8] For a primer on epistemology, see Henry Folse's page at `http://people.loyno.edu/~folse/episbasic.html`.

[9] For simplicity I choose not to allow for varying beliefs over evidence. In other words, I suppose $p(e) = 1, \forall e \in E$. In real life, however, not all evidence is made equal. Measurements may be imprecise, biased, or made by faulty instruments.

Making assumptions are part of everyday human experience, be it explicitly or implicitly. As any other beliefs, assumptions may be justified by various reasons. And some may be more or less useful than others.[10] Finally, they may be interpreted as artificially enlarging F.

Next, I define basic terminology to describe how claims relate to each other through basic rules of logic. Define $r : C \to C$ as the *refutation* correspondence, mapping a claim to the set of claims falsified by it, i.e. claims $d \in C$ such that $c \implies \neg d$. Abusing notation, for any given set of claims $S \subseteq C$, define $r(S)$ as all claims jointly refuted by it. The next concept is central to the essay.

Definition 9. The *uncertainty space* U is the set of claims consistent with F, namely $U = \{c \in C | c \notin r(S), \forall S \subseteq F\}$.

In words, this represents all claims that are not refuted by facts. This is the set of claims that humans can ever debate and disagree upon. Notice a few points. First, U weakly shrinks as F grows. As knowledge accumulates, less and less claims remain open for debate.[11] For example, as we accept General Relativity as a fact describing the Universe, we rule out that the Earth is flat. Second, we can think of facts F discretely or continuously. The former refers to only accepting claims as facts when their probability rises above $\bar{p}$. The latter refers to saying claims are "more a fact" with any small increase in their probability (and not necessarily rising above $\bar{p}$).

Given the discussion above, we can imagine a *belief cloud* made of claims C, their logical relationships and belief properties.[12] Claims form entailment clusters, with arrows flowing from sets S which imply other claims, and refutation clusters too. We can visualize the cloud's properties, such as a belief structure, with each claim's $p(c)$, and some classified as facts. Additionally, we can picture the uncertainty space as the shrinking cluster of claims not jointly refuted by the set of facts. This cloud represents the current state of our knowledge. It also helps visualizing knowledge growth: it grows when beliefs over claims change away from 0.5 (based on new evidence or other reasons), when the set of facts F expands, and when the uncertainty set U shrinks.

[10] Assumptions summarize attitudes and beliefs in a variety of contexts. Established religions and their holy books support faith in (contradicting) sets of claims, e.g. "Jesus was born out of a virgin", "Moses received the ten commandments from God on top of Mount Sinai", or "Muhammad rose to heaven at Al Aqsa". In science, assumptions may be made in order to study different scenarios analytically, or when there is good reason to do so. In forensic investigations, where hard evidence is often scant, assumptions can greatly advance investigations.

[11] This idea has been most famously challenged by Kuhn (1962) when discussing the structure of scientific revolutions. While I do not rule the possibility that claims believed to be true at one point may be believed otherwise later on, I am laying out a more abstract view of the growth of knowledge. This accommodates paradigm changes and revolutions.

[12] In mathematics, the cloud would be represented by a directed graph (DG). Claims are nodes, and their entailment relationships are edges.

Lastly, despite humanity's achievements in science, we may never have a complete picture to describe the Universe. We may, however, make the following assumption.

Assumption 1. The Universe is described by a finite set of claims R.

Notice two points. First, it is not necessarily the case that $R \subseteq C$. We may have never stated claims that truly describe the Universe, either for lack of vocabulary or lack of imagination. Second, there is no *a priori* relationship between R and F. What we believe to be true may be a set smaller, larger or incompatible with R. Our attempt at knowledge is approximating F to R. Acquiring knowledge is a selection problem: we must choose what claims to believe in and what not to.

III Applications

This Section builds on the basic definitions above, and applies them to different aspects of human affairs.

III.A Communication

The *fundamental problem of knowledge* is that the objects $\{p(\cdot), \overline{p}, \underline{p}\}$ do not exist objectively, i.e. they exist only subjectively inside observers. In other words, there are no all-knowing entities that rationally compute beliefs, take all evidence into consideration and establish which claims are facts and which are not. In reality, there are only observers, with their limited minds, personalities, and assumptions, trying to communicate and collectively move knowledge further.

Definition 10. A *human*, or *person*, is a set of claims $C_h \subseteq C$, a vector of beliefs $p_h : C_h \to [0,1]$, and truth cutoffs $\{\overline{p}_h, \underline{p}_h\}$.

Notice a few points. First, a person holds only a limited number of claims in mind (with varying cognitive capacity). This includes how much evidence E_h a person has in mind for consideration.[13] It also summarizes aspects of a person's personality, her worldviews and core beliefs.[14] Second, a person has subjective beliefs p_h, which need not coincide with other people's. Each person may judge different claims as true or

[13] I interpret C_h broadly. It may represent the set of claims a person can hold in mind, or that she has ever considered, or that she can understand.

[14] It is not the purpose of this essay, but this framework is consistent with the modern field of Cognitive Behavioral Psychology. As first conceptualized by Beck (1976) and summarized by Wallace (2016), we can classify a person's beliefs as core beliefs, assumptions and automatic thoughts.

false, and may update beliefs differently as new evidence accumulates (which need not coincide with Bayes rule).[15][16][17] Third, a person has subjective truth cutoffs to judge what claims are facts or not. Some people are easily convinced by certain arguments, while others are not. Fourth, this framework naturally generates a person-specific set of facts F_h and uncertainty space U_h.

We can also represent each person as a separate belief cloud. Communication is, then, an attempt to map, "align" and change each person's belief cloud. Through conversation, people can clarify what each person believes, find areas of agreement, introduce new ideas (nodes), and modify beliefs (persuasion).s

Next, I outline several stages of communication and intellectual exchange. First, there is *factual alignment*: when people can agree on the facts and evidence. If we can all agree on what counts as hard facts, then we can delineate what is the collective uncertainty space U, within which we can discuss. Factual alignment should be a lower bound any civilized society should strive for, but it is still often challenging to achieve.

Second, there is *honest disagreement*: when individuals can clearly discern what extra assumptions each one makes, can compute their logical implications, and understand where any disagreement starts from. It may be described as "agreeing to disagree". This is often difficult to achieve because people don't reveal (or don't know) what their underlying assumptions and beliefs are. It is common for debaters to take hours to arrive at the fundamental source of disagreement, the assumption that both cannot agree to grant.

A third situation is *consensus*: when everyone accepts the same set of beliefs and facts. Obviously, this would make everyone cognitively and emotionally equal, which is not only unrealistic, but probably undesirable.

Moreover, society is pushed away from factual alignment and honest disagreement by various forces. First, nowadays different people consume different news sources and live in polarized social media echo chambers. This makes people hold non-converging beliefs and believe "different facts" (which may outright destroy social fabric and mutual toleration).

[15] I choose to simplify the framework and not allow for varying intelligence, or processing power capabilities, across humans. Some people are more skilled in interpreting evidence in favor of some claims vs. others, in "connecting dots" and generating new ideas. I also do not discuss *cognitive dissonance*, when humans hold contradicting beliefs simultaneously in mind.

[16] Truth cutoffs could also be different across different domains. Some people may never change their minds in political matters, but accept new facts, say, in cooking.

[17] For the sake of generality, I do not discuss further how beliefs are formed and updated. Whole fields within Psychology, Neuroscience, Economics and Politics study how individuals form opinions and beliefs.

Second, beliefs are often not amenable to contradicting evidence or logic. Be it caused by religious dogma, by tribalism, or simply by a psychological disconnection from reality, some people simply have no answer to the question "what evidence could ever change your mind?" (Haidt, 2012).[18][19] Third, we are often guilty of employing logical fallacies[20] and being misled by cognitive biases (Kahneman, 2013) and motivated reasoning (Epley and Gilovich, 2016). Finally, there is always widespread asymmetry of information regarding each person's and institution's beliefs and credibility. In a context where people cannot differentiate facts from opinions, or when they must yield power to authority (e.g. "experts", media pundits), bad arguments may never be purged away with proper refutation.

III.B Policy Making, Democracy and Business

When decisions need to be made, the role of scientists and analysts is to provide facts and estimates of costs and benefits of different scenarios. In other words, their role is to map the uncertainty space U. Moreover, analysts need to be honest about what the data can rule out, and what it cannot. Statistics can only take us so far. Data are seldom perfect, and inference is only as credible as its supporting assumptions (Manski, 2007). Within the uncertainty space, the decision makers bear the responsibility for choices made. They may want to make extra assumptions (expanding their F_h), which would steer decisions in a particular way, but it should be clear that those are decisions under uncertainty.

The previous paragraph applies both in policy making and in business. When voting for different policies, ideally, scientists first delineate U by presenting the facts and trade-offs. Then, with that in hand, democracy can do its job. Voters and representatives weigh the different options according to their preferences, and make decisions.[21] Analogously, in business, decisions need to be made fast, and only rarely there is fully compelling evidence to support them. At that point, to make decisions, good executives rely on intuition, experience, over-confidence, and other signals available to them. It is no wonder successful business-people get big financial compensation.

[18] As perhaps too acidly expressed by Dawkins (2001), any person who does not believe established facts or theories are either ignorant, stupid, insane or wicked.

[19] Beyond communication itself, a more serious problem arises when groups are not willing to tolerate the existence of different beliefs, practices and culture. Some level of political liberalism and mutual tolerance is necessary for societies to sustain norms of honest disagreement.

[20] Standard examples are the strawman, ad hominem, personal incredulity, appeal to emotion, and appeal to nature.

[21] Mapping U has a major implicit benefit: it puts a clear boundary on what facts end and preferences start.

III.C Science

The framework in Section II naturally encapsulates science and the scientific method.

Definition 11. *Science* is a set of testable theories, a set of observations and evidence supporting these theories, and a set of systematic practices and protocols to generate measurements, categorize data, formulate new theories, and perform experiments.

In other words, we can say the following. Science is constantly generating new theories (expanding C), generating new evidence (expanding E), testing theories (changing beliefs $p(\cdot)$), and establishing new facts (expanding F). By doing so, as it generates more knowledge, it progressively shrinks the uncertainty space U.

Being constrained by the facts, and deliberately working on problems inside U is at the core of the scientific endeavor. Science adopts an open-ended correction mechanism, with institutional incentives in place, of always trying to disprove claims, as opposed to outright proclaiming some to be true. This practice, plus a large amount of human ingenuity and curiosity, is what made science so successful and trustworthy.[22]

[22] Obviously, this is not to deny the existence of problems in day-to-day work in science and research, e.g. publication bias, p-hacking, lack of transparency, academic politics, etc.

References

Beck, A. T. (1976). *Cognitive Therapy and the Emotional Disorders*. International Universities Press.

Dahis, R. (2018). Is Economics a Science? Well, Not Yet. *SSRN Electronic Journal*, pages 1–22.

Dawkins, R. (2001). Ignorance Is No Crime. *Free Inquiry*.

Epley, N. and Gilovich, T. (2016). The Mechanics of Motivated Reasoning. *Journal of Economic Perspectives*, 30(3):133–140.

Feyerabend, P. (1975). *Against Method*. New Left Books.

Haidt, J. (2012). *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Vintage, 1st edition.

Harris, S. (2014). *Waking Up: A Guide to Spirituality Without Religion*. Simon & Schuster, 1st edition.

Kahneman, D. (2013). *Thinking, Fast and Slow*. FSG Adult, 1st edition.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Levitsky, S. and Ziblatt, D. (2018). *How Democracies Die*. Crown, 1st edition.

Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.

Popper, K. (1959). *The Logic of Scientific Discovery*. Routledge, 2nd edition.

Searle, J. (1999). *Mind, Language and Society: Philosophy In The Real World*. Basic Books.

Wallace, L. (2016). *Cognitive Behavioural Therapy: 7 Ways to Freedom from Anxiety, Depression, and Intrusive Thoughts*. Independently Published.