The Why is Not the Same as the How: Levels of Analysis and Scientific Progress in Psychology

Laith Al-Shawaf
Jul 20, 2020 — 10 min read

The emerald-colored jewel wasp (*Ampulex compressa*) parasitizes the common cockroach with a sophistication that beggars the imagination. Here is how the morbid scene unfolds. The wasp begins by injecting a venemous compound into the roach's body, temporarily paralyzing its victim. Next, it seeks out two specific

areas of the roach's brain, injecting them with a custom-made neurochemical cocktail. This injection has a surgically precise effect: it leaves the roach's motor abilities intact, but robs it of its will to escape—an astonishing combination that will become important later. Finally, the wasp delivers its *coup de grâce*: it drags the zombified roach to its grave, deposits an egg on it and buries it alive. Why does the wasp bother to do these specific things in this exact sequence?

Sixty years ago, the great ethologist Niko Tinbergen argued that in order to achieve a complete understanding of any biological trait or system, we have to answer four separate questions: (1) how the system develops during an organism's lifespan (known as *ontogeny* or development); (2) how the system works in the immediate present (*mechanism* or immediate causation); (3) how the trait evolved over time (*phylogeny*); and (4) why the trait evolved (its *function* or adaptive value). These four questions—ontogeny, mechanism, phylogeny, and function —represent different "levels of analysis." Tinbergen's key insight was that if you want to understand something about an organism, you need all four levels of analysis.

These four levels can be collapsed into two broader categories known as the proximate and ultimate levels of analysis (see also Mayr 1961). The *proximate level* consists of ontogeny and mechanism, which deal with how the trait developed during an organism's life and how it works in the immediate present. The *ultimate level* consists of phylogeny and function, which are about how and why the trait evolved. Proximate explanations explain *how* something works—for example, how the liver works, or what causes the sucking reflex in infants. Ultimate explanations explain *why* the system works that way, or why it exists in the first place—for example, why we have a liver (to filter blood, detoxify poisons and facilitate digestion) or why the sucking reflex evolved (to facilitate breastfeeding). The proximate and ultimate levels of analysis answer different questions about the same phenomenon, and both are crucial to a complete understanding of the phenomenon.

It wouldn't be much of an exaggeration to say that the conflation of the proximate and ultimate levels of analysis is one of the biggest problems in psychology today, on par with issues like the replication crisis (see also here and here), the absence of an overarching theoretical framework for the discipline and the heavy reliance on WEIRD samples. Addressing this issue is key to advancing the field.

The grisly scene described above explains *how* the wasp parasitizes its host—this is the proximate level of analysis. A complete proximate analysis would explain the physiological basis of the behavior, how the mechanisms responsible for the behavior develop during the wasp's lifetime, and more.

The ultimate level of analysis is different. It explains the *why* of these steps—why does the wasp engage in these specific behaviors? The answer: the first injection serves to temporarily paralyze the roach, giving the wasp enough time to locate the precise areas of the roach's brain that it must target next.

Why does the second injection leave the roach's motor abilities intact, but remove its will to escape? Because the tiny wasp faces a problem: it needs to drag the roach to its burial site, but it is too small to manage this if the roach is paralyzed. So the wasp's injection *nullifies the roach's ability to initiate movement, but does not actually paralyze the roach's legs*—this way, when the wasp begins to drag the roach to its nest, the roach's legs will automatically move in concert with the rest of the body, helping the process along. Other solutions would not have worked: if the roach had been paralyzed, the wasp would have been unable to drag it, and if the roach had retained its ability to initiate its own movements, it would have escaped. The wasp's solution—eradicate the roach's ability to initiate its own movement, but leave the mobility of its legs unimpaired—brilliantly avoids both problems.

Finally, at the end of the process, why does the wasp deposit an egg on the roach and then bury it alive? Because when the wasp larva hatches, it will need a *live* victim to feed on. In fact, the wasp does one more remarkable thing to facilitate this: its injection slows down the roach's metabolism, keeping its victim alive for longer. This ensures that the roach will still be fresh when the wasp's offspring hatches, hungry for a nightmarish meal of horrors.

To properly understand what the wasp is doing, we must answer two different questions: the proximate one (how does the wasp accomplish this?) and the ultimate one (*why* does the wasp do those things?). Any answer that does not address both levels of analysis is incomplete.

This is a well-established principle in biology, and it would be difficult to overstate how far-reaching its implications are. It serves as the background theoretical framework for all research on animal behavior. Its originator, Niko Tinbergen, was

co-recipient of the only Nobel Prize ever awarded for the study of animal behavior (in 1973). There is not a single biologist studying behavior who isn't familiar with the four questions and their grouping into the two larger categories of proximate and ultimate. In my home discipline of psychology, the situation is very different: the principle has the same immense scientific value and far-reaching implications, but only a minority of psychologists are familiar with it.

For more than a century, psychology has been dominated by an almost exclusive focus on the proximate level of analysis. All of the most familiar types of explanation in psychology are proximate, including cognitive, neurophysiological, sociocultural, learning and developmental explanations. They are all important, of course, and all necessary for a comprehensive science of the mind. But they leave out an entire level of analysis.

Ultimate Explanations Are Necessary for a Complete Science of the Mind

Why would the explanatory partitioning of phenomena into different levels of analysis apply only to biology, and not to psychology? Just like the heart and the liver, aspects of the mind are subject to the same four questions: how they develop during the organism's life (*ontogeny or development*); how they work in the present moment (*mechanism*); how they evolved over time (*phylogeny*); and why they evolved (*function*).

Scientists have long known that they cannot skip either the proximate or ultimate level of analysis if they want a complete understanding of our bodily organs. The same goes for our *mental organs*—if we want a complete understanding of, say, attention, memory and emotion—we will need to address these aspects of the mind at both the proximate and ultimate levels of analysis.

This does not imply that every aspect of our minds has an evolved function. As evolutionary psychologists will tell you, our minds contain plenty of byproducts (side effects) that have no evolved function. But even these functionless byproducts require the ultimate level of analysis: they have evolved over time (so they require the phylogenetic level of analysis) and are byproducts of adaptations that have a biological function (so they require the functional level of analysis). There is simply no way around the conclusion that the ultimate level of analysis applies to the mind and how it works.

Psychologists who focus solely on proximate questions can still make great strides in discovering the way the mind works. The history of psychology in the twentieth century is testament to that. But progress becomes more rapid when you incorporate the ultimate level of analysis. More importantly, it's impossible to build a complete science of the mind if you ignore one of the two basic questions that pertain to all living organisms. Evolutionary explanations are not an optional add-on in psychology, as many seem to think; they are indispensable. If we want a richer and more complete science of the mind, we cannot omit half of the explanatory equation.

Ultimate Explanations Shed New Light

Ultimate explanations shed new light—offering insights distinct from those offered by the proximate level. Consider the following question: in most sexually reproducing species, why does the sex ratio hover around 1:1? In species like humans, with an XY sex-determination system, the proximate answer is well-known: roughly, each offspring will inherit an X from its mother and either an X or a Y from its father, with a 50% chance of each. So each new zygote has a roughly 50% chance of being XY (male) and 50% chance of being XX (female), making the population sex ratio approximately 1:1. Compelling though this explanation is, it is only half the answer—the proximate half.

What more could an evolutionary answer contribute? It might seem as though there's nothing left *to* explain. The ultimate explanation belies this expectation. It goes something like this: imagine that female births are less common than male births in a population, so there are fewer adult females in the population. Because females are rarer, on average they have higher reproductive success than males. Because females have higher reproductive success, individuals who are genetically pre-disposed toward producing daughters end up having more grand-offspring, on average. This makes those genes that predispose toward producing daughters increase in frequency in the population. This, in turn, makes female births more common. As female births become more common, the increased reproductive success associated with being female is reduced, eventually tapering off to nothing as the sex ratio reaches 1:1. In this way, a population with an initial bias toward fewer females will eventually gravitate toward a 1:1 sex ratio. The same argument holds if you substitute males for females in the example. In other words, a

population that starts with fewer male births will also tend to converge on a 1:1 sex ratio.

The key thing to realize about this ultimate explanation is that it is not an *alternative* to the proximate genetic answer; it is a *complement* to it.

Notice also that the ultimate explanation sheds new light on the 1:1 sex ratio. When you first discover the ultimate explanation for a phenomenon in physiology, psychology or behavior, you learn something genuinely new, even if you already knew the proximate explanation for that same phenomenon. And, crucially, your new layer of understanding is neither redundant, nor in conflict with, your existing proximate understanding. Instead, it complements what you know, completing the picture. The ultimate explanation also makes it clear that any sense of epistemic completeness we had earlier was an illusion: in reality, we had only half of the explanation.

Ultimate analyses offer additional explanatory and predictive benefits, too. For example, this analysis of the 1:1 sex ratio yields predictions about which species will not exhibit this ratio, and why. Ultimate explanations in psychology have the same scientific virtues: they shed new light, explain existing findings and predict new ones. This is true for a wide variety of psychological phenomena, including anger, pride, shame, moral cognition, mate preferences, mental health, personality, politics, social learning, epistemic vigilance, spatial and numerical thinking, core knowledge, illness, education, war, hunger, disgust, psychopathology, reasoning, aging, cancer, psychopathy, religious belief, animal cognition, machine learning and artificial intelligence, child development, parenting, status management, sex differences, birth control, eating disorders, reputation, punishment, revenge, altruism, empathy, emotion, anxiety, friendship, gratitude, grief, recursive thinking, cognitive biases, indirect speech, common knowledge and our ability to detect violators of social contracts.

The Proximate–Ultimate Distinction can Turn Needless Conflict into Productive Exchange

Accepting the proximate-ultimate distinction has another salutary effect: it can turn unnecessary conflict into productive exchange. For example, one of the most common false dichotomies in psychology is the evolution vs. learning fallacy. It maps directly onto the ultimate–proximate distinction.

The conventional wisdom in the social sciences is that if a behavior or psychological trait is learned, then it is not evolved, and vice versa. This is an inaccurate and deeply misleading way of conceptualizing the issue. The fallacy can be traced to the ultimate–proximate distinction: evolutionary explanations are at the ultimate level of analysis, whereas learning is at the proximate level. They are therefore not in conflict. In fact, not only are the two perfectly compatible, but in a great many cases *we will need to invoke both* to comprehensively explain the facts of psychology and behavior.

The proximate–ultimate distinction therefore helps dissolve the evolution vs. learning fallacy, a longstanding impediment to progress in psychology. In the older nativist vs. empiricist iterations of this debate, the key mistakes of the respective camps were (1) the erroneous notion that learning is relatively unimportant, and organisms are born with all the knowledge they need, and (2) the equally incorrect notion that organisms learn everything they know, and evolution isn't needed to explain learned behaviors. Adopting just one conceptual tool— the proximate–ultimate distinction—steers us clear of that explanatory Scylla and Charybdis. And once we've gotten rid of those inaccurate claims, we're free to replace them with the following mutually compatible ideas: (1) learning matters greatly, (2) learning mechanisms exist because they evolved and (3) their patterns of operation reflect the kinds of problems they evolved to solve. Studying how the mind works becomes easier and more fruitful without conceptually muddled dichotomies clouding our thinking. The proximate–ultimate distinction helps to achieve this.

It is worth pausing to note that, like learning-based explanations, sociocultural ones are also at the proximate level of analysis. This means that evolutionary and sociocultural explanations are *also* often compatible. It is common to think that there is some kind of inherent tension between these two kinds of explanations, but there isn't—at least not a necessary one (see this paper for more detail on when proximate and ultimate explanations *can* conflict). The more widespread this realization becomes in the social sciences, the more we can foster collaboration between researchers with widely different backgrounds and theoretical approaches.

Properly understood, then, this conceptual tool could be a catalyst that helps us build bridges between different social and life sciences. The proximate–ultimate

distinction helps nudge us in this direction, toward a union of disparately trained scientists studying evolution and sociocultural influences in a collaborative manner.

Conclusion

The proximate–ultimate distinction is key to progress in psychology for at least three reasons. First, psychology cannot hope to be a comprehensive science of the mind if it continues to ignore half of the explanatory equation. Second, ultimate explanations provide an additional layer of understanding, distinct from that offered by the proximate. And third, accepting the distinction between the two levels of analysis can turn unnecessary conflict into productive exchange.

William James—one of the founders of modern psychology—was often remarkably prescient. He stressed the importance of evolutionary thinking in psychology as early as the late 1800s. More than a century later, we are making real progress in fulfilling James's vision. The ultimate goal is to build a mature science of the mind that explains not only how our psychological mechanisms work, but also *why* they work that way and why they exist in the first place.

So the next time you encounter an evolutionary explanation of a phenomenon for which you've already heard a plausible proximate explanation, please remember that the two are typically not alternatives. They are *complementary* explanations that answer different questions at different levels of analysis. The proximate explanation that you're familiar with probably focuses on learning, culture, hormones, the brain, or a combination thereof. The ultimate explanation might instead focus on how the behavior evolved from an ancestral form, or why it evolved in the first place. Not only is there no inherent conflict at play, but it is only when we combine the two levels of analysis—how the system works and why it evolved that way—that we can finally approach a complete understanding of the phenomenon in question.