

Does Altruism Exist?

Introduction

ALTRUISM AND EVOLUTION

ALTRUISM IS a concern for the welfare of others as an end in itself. Improving the welfare of others often requires a cost in terms of time, energy, and risk. Even the simple act of opening a door for someone requires a tiny expenditure of time and energy. At the opposite extreme, saving a life often requires a substantial risk to one's own.

Seemingly altruistic acts take place all around us, from simple courtesies to heroic self-sacrifice. The question “Does altruism exist?” might therefore seem like a silly topic for a book. Yet the claim that altruism does *not* exist has a long tradition in philosophical, political, economic, and biological thought. Those who challenge the existence of altruism do not deny that there are seemingly altruistic *acts* but question whether they are based upon altruistic *motives*. Pure egoists might help others as a means to pursue their own selfish ends, but they don't qualify as altruists because they don't care about the welfare of others as an end in itself.

The idea that altruism truly doesn't exist—that *all* people care *only* about themselves—might seem too extreme to take seriously.

Consider, however, that the word “altruism” didn’t exist until 1851, when it was coined by the French philosopher Auguste Comte. If people are altruistic, then why doesn’t the word (or its equivalent) have a more ancient pedigree?

The plot thickens when we consider the role that altruism plays in religious thought. About a decade ago, the John Templeton Foundation commissioned two eminent theologians and scholars of religion, Jacob Neusner and Bruce Chilton, to organize a conference on altruism in world religions. They began by asking one of their colleagues, William Scott Green, to outline a general framework for studying altruism that could be applied to particular religions by the other conference participants. Green reviewed definitions of altruism and provided this concise version: “Intentional action ultimately for the welfare of others that entails at least the possibility of either no benefit or a loss to the actor.” According to the conference participants—each an expert scholar on a given religion—this concept is foreign to the imagination of *all* of the world’s major religious traditions.¹

I will return to Comte and religious belief systems in chapter 6. For now, suffice it to say that the question “Does altruism exist?” is not trivial and takes us into deep intellectual waters. In this book I use evolutionary theory as our navigational guide. Altruism occupies center stage in Darwinian thought because it appears difficult to explain as a product of natural selection. If natural selection favors traits that cause individuals to survive and reproduce better than other individuals, and if altruistic acts increase the survival and reproduction of others at a cost to the altruist, then how can altruistic traits evolve? This question has been debated from Darwin to the present and might seem as unresolved among evolutionists as among philosophers.² On the contrary, I believe that a resolution is at hand, which I attempt to convey in this book.

This book initiates a series of short books on big questions organized by the Templeton Press and Yale University Press. The idea of a short book on a big question is both enabling and constraining. It must stick to the basics and provide an accessible introduction to a large and diverse audience. It should act as a portal to the academic literature without getting bogged down in the details.³ Ideally, it should offer something to the expert in addition to the novice. I welcome the opportunity to write a short book on the big question of altruism for three reasons.

First, some of the biggest controversies in science end up appearing obvious in retrospect, even though decades were required for their resolution. Examples include the Copernican view of the solar system, Darwin's theory of evolution, and the theory of continental drift. We find it hard to understand in hindsight why smart people took so long to reach a consensus on these bodies of knowledge. I believe that the question of how altruism evolves is such a controversy that is just entering its resolution phase. Despite a half century of disagreement, people will look back and wonder what all the fuss was about. This book operates from a postresolution perspective, which enables me to explain how altruism evolves in a short space.

Second, most of the thinking on human altruism is not from an evolutionary perspective. Foundational ideas were established by thinkers who lived before Darwin or by contemporaries who did not fully absorb the implications of his theory. Even the modern academic literature is largely not from an evolutionary perspective, for reasons that I will recount as we go along. This literature often fails to carve the subject of altruism at the right joints, to use a venerable philosophical phrase. In other words, it often fails to make distinctions that are foundational from an evolutionary perspective, such as between proximate and ultimate causation.

When we base our analysis on evolution, some questions that appeared central become peripheral and other previously ignored questions become central. As an analogy, suppose that someone owes you money and offers to pay by cash or check. You might have a mild preference for one payment method, but your main concern is to be paid. As we shall see, some criteria for defining human altruism that have occupied center stage are like payment methods. They might exist, but we shouldn't care that much about them compared to other ways of thinking and feeling that produce equivalent results. More generally, I look forward to showing how evolutionary theory can carve the topic of altruism at the right joints, or at least *different* joints, in a short space.

Third, understanding altruism from an evolutionary perspective has enormous practical import in addition to its intellectual charm. My own inquiries took a practical turn in 2006 when I began to study altruism in the context of everyday life in my hometown of Binghamton, New York. A year later I helped to create the Evolution Institute, the first think tank to formulate public policy from an evolutionary perspective.⁴ These endeavors have enabled me to explore topic areas such as economics and business, where the prevailing assumption is that people are motivated entirely by self-interest. Practices that follow from this assumption have had ruinous consequences, as we shall see. Yet well-meaning efforts to promote altruism can also have pathological consequences when they do not take basic evolutionary considerations into account, as strange as that might seem. By the end of this short book, if I have done my job well, the reader will understand why evolutionary theory is essential for accomplishing the altruistic goal of making the world a better place.

Groups That Work

IMAGINE THE purest of altruists, who thinks only about making the world a better place and cares nothing for her own welfare, other than as part of the common good. Now imagine convincing this person that the world of her dreams is possible. There is only one catch. The people capable of achieving and maintaining her utopia do not necessarily think or feel about themselves and others in the same way that she does. She might need to learn how to think and feel differently to accomplish her altruistic goals.

Would a pure altruist accept this bargain? It seems to me that she must. If she insisted on clinging to her habits of thought, she would interfere with her own altruistic impulse to make the world a better place.

This thought experiment highlights two meanings of altruism that need to be distinguished from each other. The first meaning refers to how people *act*. If our pure altruist were asked to describe her perfect world, she would spend more time on what people do than how they think or feel. At a local scale, they would refrain from robbery, rape, and murder. A helping hand would be offered to those who fall upon hard times. At a global scale, there would be

an absence of warfare. The division between rich and poor nations would disappear. The great mass of humanity would somehow manage its affairs to avoid polluting and overheating the earth. It is impossible to talk about making the world a better place without referring to *actions* of this sort.

The second meaning of altruism refers to the *thoughts and feelings* that cause people to act as they do. Some thoughts and feelings are oriented toward the welfare of others and society as a whole, while others are more self-oriented. It might seem that the two meanings are straightforwardly related to each other. Get people to think and feel more altruistically, and the world will become a better place. If everyone could become as pure as our pure altruist, utopia can be achieved.

A little reflection reveals a more complicated story. Any given action can be motivated by a diversity of thoughts and feelings. I might help you because (a) I think it's the right thing, (b) I take pleasure in your pleasure, (c) I regard it as my ticket to heaven, (d) I am trying to improve my reputation, (e) I'm trying to put you in my debt, or (f) I'm being paid to do it. It's not obvious which of these thoughts and feelings count as most altruistic (a vs. b, for example) or if the most altruistic thoughts and feelings result in the most helping. Perhaps reputational mechanisms (d) work better than voluntary altruism (a–c), for example. Or perhaps market forces driven by self-interest result in societies that work best, as some economists contend.

These complications force us to recognize a one-to-many relationship between any given action and the mental events that can cause the action. They also force us to recognize that *our preference for some thoughts and feelings over others is based primarily on the actions they produce*. There is no other reason to privilege thoughts and

feelings that count as altruistic over those that count as selfish.

In this chapter, we begin answering the question “Does altruism exist?” by investigating altruism at the level of action, deferring altruism at the level of thoughts and feelings until later. When Ted benefits Martha at a cost to himself, that’s altruistic, regardless of how he thinks or feels about it. Altruism at the level of action is closely related to group-level functional organization. The reason that people need to provide services for each other is because we are a social species and mutual aid is required to accomplish together what cannot be accomplished alone. The list of activities that required mutual aid before the advent of agriculture included child care, hunting and gathering, defense against predators, and defense and offense against other human groups. The advent of agriculture led to an autocatalytic cycle between the production of resources and larger societies with division of labor, leading to modern megasocieties. Today we are so dependent upon the actions of others that we could no more survive on our own than an ant separated from its colony.

Something is functionally organized when its parts work together in a coordinated fashion to achieve a given end. The organelles of a cell and the organs of a multicellular organism are miracles of functional organization designed by natural selection to enhance survival and reproduction. The parts of a can opener are functionally organized to open cans. They are not functionally organized for chopping wood. Anyone can make this determination by attempting to use a can opener for both tasks.

When a group of organisms is functionally organized, its members coordinate their activities for a common purpose, just like the organs of an organism and the parts of a can opener. A group that is extremely well organized could even be called a *superorganism*, a

metaphor with a venerable history in philosophical, religious, and political thought.

Do human groups ever achieve a degree of functional organization that deserves the term superorganism? How about animal groups? In short, do functionally organized groups exist? This is not exactly the same question as “Does altruism exist?” but it is highly relevant and a lot easier to answer, as we shall see.

To begin, consider an aircraft carrier. The ship itself is a marvel of functional design for the purpose of providing a mobile airport. Nobody who studies its features can doubt its function any more than doubting the function of the heart or the toothed wheel of a can opener. The social organization of the crew is also a marvel of functional design, as anthropologist and cognitive psychologist Edwin Hutchins describes in his book *Cognition in the Wild*.¹ Several hundred in number, the crew is organized into teams that are specialized to perform various tasks. Hutchins focuses on the team responsible for the crucial task of navigation and especially locating the position of the ship when it is close to shore. The process of measuring the compass directions of landmarks on the shore, marking them on a chart, and using triangulation to identify the position of the ship is easy to understand and can be performed by a single person given enough time, but it must be performed by a team of people to ensure accuracy and to recalculate the ship’s position every few minutes. The social organization of the navigational team has been refined over the course of nautical history and has even coevolved with the physical architecture of the chartroom. It includes the training of team members over the long term and enough redundancy so that some members of the team can be called away without compromising the performance of the group too severely. One reason that Hutchins chose to study navigational teams is that their social interactions can be

understood in purely computational terms—the perception and transformation of information, leading to physical action—in just the same way that a single brain can be understood in computational terms. This particular group of people engaged in this particular task is functioning like a single brain. I say more about this comparison as we proceed.

Another example of group-level functional organization serves us throughout the book. Elinor Ostrom was awarded the Nobel Prize in economics in 2009 for showing that groups of people are capable of managing their own resources, but only if they possess certain design features.² She and her associates did this in part by assembling a worldwide database of groups that attempt to manage resources such as irrigation systems, ground water, fisheries, forests, and pastures. These are called *common-pool resources* (CPR) because they cannot easily be privatized. They are vulnerable to a social dilemma that ecologist Garrett Hardin made famous in his 1968 article titled “The Tragedy of the Commons.”³ Hardin imagined a village with a field upon which any townspeople could graze his or her cows. The field can only support so many cows, but each townspeople has an incentive to exceed this capacity by adding more of his or her own cows, resulting in the tragedy of overuse. Prior to Ostrom’s work, economic wisdom held that the only solution to this problem was to privatize the resource (if possible) or to apply top-down regulation. Ostrom’s demonstration that CPR groups are capable of managing their own affairs was so new and important that it warranted the Nobel Prize.

Ostrom’s work was equally noteworthy for showing that groups must possess certain design principles to effectively manage their resources. An efficacious group needs these design principles in the same way that we need our organs, a can opener needs its parts, and the navigation team aboard an aircraft carrier needs to

be structured as it is. Ostrom identified eight core design principles required for the CPR groups in her worldwide database to effectively manage their affairs:

1. *Strong group identity and understanding of purpose.* The identity of the group, the boundaries of the shared resource, and the need to manage the resource must be clearly delineated.

2. *Proportional equivalence between benefits and costs.* Members of the group must negotiate a system that rewards members for their contributions. High status or other disproportionate benefits must be earned. Unfair inequality poisons collective efforts.

3. *Collective-choice arrangements.* People hate being told what to do but will work hard for group goals to which they have agreed. Decisionmaking should be by consensus or another process that group members recognize as fair.

4. *Monitoring.* A commons is inherently vulnerable to free-riding and active exploitation. Unless these undermining strategies can be detected at relatively low cost by norm-abiding members of the group, the tragedy of the commons will occur.

5. *Graduated sanctions.* Transgressions need not require heavy-handed punishment, at least initially. Often gossip or a gentle reminder is sufficient, but more severe forms of punishment must also be waiting in the wings for use when necessary.

6. *Conflict resolution mechanisms.* It must be possible to resolve conflicts quickly and in ways that group members perceive as fair.

7. *Minimal recognition of rights to organize.* Groups must have the authority to conduct their own affairs. Externally imposed rules are unlikely to be adapted to local circumstances and violate principle 3.

8. *For groups that are part of larger social systems, there must be appropriate coordination among relevant groups.* Every sphere of activity has an optimal scale. Large-scale governance requires finding the

optimal scale for each sphere of activity and appropriately coordinating the activities, a concept called *polycentric governance*.⁴

These core design principles were described in Ostrom's most influential work, *Governing the Commons*, published in 1990. A recent review of ninety studies that have accumulated since the original study provides strong empirical support for the efficacy of the core design principles, with a few suggested modifications that differentiate between the resource system and those authorized to use it.⁵ I was privileged to work with Lin (as Ostrom encouraged everyone to call her) and her postdoctoral associate, Michael Cox, for several years prior to her death in 2012, resulting in an article entitled "Generalizing the Core Design Principles for the Efficacy of Groups," which I draw upon throughout this book.⁶

These two examples establish a point that I hope will be obvious to most readers: at least some human groups are impressively functionally organized, inviting comparison to a single organism. How about animal groups? Nonhuman species afford many examples of group-level functional organization, including but not restricted to social insect colonies such as ants, bees, wasps, and termites.⁷ These colonies have been celebrated for their groupishness since long before science existed as a cultural institution. Beehives are pictured on the road signs in Utah because the Mormons, following a long Christian tradition, admire the industry of individual bees on behalf of their groups.

Worker honeybees forgo reproduction and are willing to die in defense of their colony. These acts appear highly altruistic in behavioral terms (although the mental world of bees is no more scrutable than that of humans), but I focus here on a more subtle aspect of colony-level functional organization: the ability to make collective decisions. Bee colonies make daily decisions on where to forage, how to allocate the worker force, and so on. In addition,

when a colony splits by swarming, the swarm is faced with the critical task of finding a new nest cavity. Social insect biologist Tom Seeley and his colleagues have studied this process in minute detail.⁸ For some of their experiments, they transport beehives to an island off the coast of Maine that lacks natural nest cavities. The researchers present the swarms with artificial nest cavities that differ in their size, height, distance from the swarm, compass orientation, exposure to the sun, size of the opening, physical position of the opening, and other parameters. In this fashion, the researchers have determined that honeybee swarms give the nest cavities a thorough inspection and reliably pick the best over the worst, just like a discerning human house hunter. How do they do it?

Seeley and his colleagues can answer this question because all of the scout bees have been marked as individuals and monitored on their travels from the nest cavities to the swarm. Most scouts do not visit more than one nest cavity and therefore do not have enough information to make an individual choice. Instead, the collective decision is made on the basis of social interactions that take place on the surface of the swarm. The bees that make up the swarm, including the queen, take no part in the decisionmaking process, other than providing a surface for the scouts to interact with each other. The interactions can be regarded as a contest among the scouts that is decided on the basis of their persistence, which in turn is decided by the quality of the cavity that has been inspected. Once the winning faction reaches a critical threshold, a new phase begins that involves arousing the entire swarm and guiding it to its new home. Even to a veteran evolutionist such as myself, it seems miraculous that such a coordinated social process among mere insects could have evolved by natural selection. But then, you and I are nothing more than a coordinated social process among mere cells!

How does the decisionmaking process of a honeybee colony compare to that of an individual such as you or me? Detailed research on rhesus monkeys involves having them watch a number of dots moving either left or right on a screen.⁹ The monkeys are trained to move their head in the direction that most of the dots are moving, by rewarding them with a few drops of sweet liquid if they succeed and few drops of salty liquid if they fail. As they are making their decision, the researchers record the activity of neurons in their brains. Remarkably, the interactions among the neurons resemble the interactions among the bees. Some neurons fire at the sight of right-moving dots and others fire at the sight of left-moving dots. Their firing rates are unequal because more dots are moving in one direction than the other. As soon as the stronger “faction” reaches a threshold, the decision is made and the monkey turns its head in the appropriate direction. Once we regard an individual organism as a population of socially interacting cells, the possibility of a group mind, comparable to an individual mind, becomes less strange.

Biologists expect social insect colonies to have group minds because they function as superorganisms in so many other respects. The phenomenon of group-level cognition in nonhuman species is more general, however. Adult female African buffaloes (but not juveniles or males) “vote” on where to graze at night by standing up and pointing their heads in their preferred direction with a distinctive posture while the herd is resting in the afternoon. This information is somehow integrated and, when evening falls, the entire herd gets up and moves in the chosen direction “as if conducted by a single master.”¹⁰

I have studied group cognition in American toad tadpoles with my former graduate student Charles Sontag.¹¹ We created food patches that varied in quality by mixing different amounts of

aquarium fish flakes, ground to a fine powder, in an agar matrix. We placed the food patches in the tadpoles' natural environment and videotaped their response. Within minutes, the tadpoles started streaming toward the food patches and discriminated the better from the worse patch. They even formed columns that bore an eerie resemblance to ant columns. In laboratory experiments, we determined that the tadpoles communicate through surface waves and that individual tadpoles appear unable to locate a food patch by themselves, even from a short distance.

In addition to empirical examples such as these, a growing body of theory suggests that group-level cognition (along with other group-level adaptations) should be found in many social species, not just social insect colonies.¹² In addition, animal groups can be functionally organized in some respects but not others. They might cooperate to locate the best food patches but compete once they get there, for example.

Let's take stock of the progress that we have made in this chapter. First, we have made a crucial distinction between altruism at the level of *action* and altruism at the level of *thoughts and feelings*. Second, we have established that altruism at the level of thoughts and feelings is worth wanting only insofar as it leads to actions that make the world a better place. Third, we have established that making the world a better place requires *group-level functional organization*—people coordinating their activities in just the right way to achieve a common goal. Fourth, we have established that group-level functional organization exists in both human and non-human societies. At least some of the time, members of groups coordinate their activities so well that they invite comparison to the cells and organs of a single organism.

The examples of group-level functional organization that I have provided do not seem very altruistic in the conventional sense. A

navigational team charting the position of a ship is different from a captain going down with the ship or a soldier falling upon a grenade to protect his comrades. A group of scout bees deciding upon the best tree cavity seems different from a worker bee delivering a suicidal sting to an intruder. The key difference involves the cost of providing services to other members of one's group. The greater the cost, the more altruistic the action appears. But this is a matter of degree, not kind, as we shall see in the next chapter.

I have purposely established the major points in this chapter with minimal reference to evolution. Evolution isn't needed to make the distinction between altruism based on action and altruism based on thoughts and feelings. It isn't needed to assert the one-to-many relationship between any given action and the mental events that can cause the action. It isn't needed to assert that our preference for thoughts and actions should be based on what they cause people to do. Evolution isn't even needed to establish the functional organization of groups, any more than it is needed to establish the functional organization of a can opener. The added value of an explicit evolutionary perspective becomes apparent as we proceed.